IBM Systems

IBM Director
Systems Management Guide

Version 5.10

IBM Systems

IBM Director
Systems Management Guide

Version 5.10

Note

Before using this information and the product it supports, read the information in Appendix N, "Notices."

Fourth Edition (October 2005)

This edition applies to version 5.10 of IBM Director and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Appendix A. IBM Director commands 529

Figures

Tables

About this book

This book provides instructions for using IBM® Director 4.20 for systems-management tasks. IBM Director consists of the following tools to meet your systems-management needs:

- IBM Director Console is the graphical user interface (GUI) for IBM Director Server. Using IBM Director Console, you can conduct comprehensive systems management using either a drop-and-drag action or a single click.
- IBM Director command-line interface (dircli) is the command-line interface for IBM Director Server. You can use a command-line prompt to access, control, and gather information from IBM Director Server.

This documentation also provides planning and implementation information for event management.

Conventions and terminology

These notices are designed to highlight key information:

Note: These notices provide important tips, guidance, or advice.

Important: These notices provide information or advice that might help you avoid inconvenient or difficult situations.

Attention: These notices indicate possible damage to programs, devices, or data. An attention notice appears before the instruction or situation in which damage can occur.

Related information

This topic provides links to additional information related to IBM Director.

IBM Director resources on the World Wide Web

The following Web pages provide resources for understanding, using, and troubleshooting IBM Director and other systems-management tools.

IBM Director information center

publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html

Updated periodically, the IBM Director information center contains the most up-to-date documentation available on a wide range of topics.

IBM Director Web site on ibm.com®

www.ibm.com/servers/eserver/xseries/systems_management/ibm_director/

The IBM Director Web site on ibm.com has links to downloads and documentation for all currently supported versions of IBM Director. Information on this site includes:

- IBM Director 5.10 - downloads and documentation
- IBM Director 4.22 - downloads and documentation

- IBM Director 4.22 Upward Integration Modules (UIMs) - downloads and documentation
- IBM Director 4.21 - downloads and documentation
- IBM Director 4.20 - downloads and documentation
- IBM Director Hardware and Software Compatibility document - lists supported @server and IBM xSeries® systems, as well as all supported operating systems. It is updated every 6 to 8 weeks.
- Printable documentation for IBM Director - available in Portable Document Format (PDF) in several languages

IBM Systems Software information center

www.ibm.com/servers/library/infocenter/

This Web page provides information about IBM Virtualization Engine™, IBM Director, and other topics.

IBM ServerProven® page

www.ibm.com/pc/us/compat/index.html

This Web page provides information about IBM xSeries, BladeCenter®, and IntelliStation® hardware compatibility with IBM Director.

IBM Systems Management Software: Download/Electronic Support page

www.ibm.com/servers/eserver/xseries/systems_management/ibm_director/

Use this Web page to download IBM systems-management software, including IBM Director. Check this Web page regularly for new IBM Director releases and updates.

IBM Servers

www.ibm.com/servers/

This Web page on ibm.com links to information, downloads, and IBM Director extensions such as Remote Deployment Manager, Capacity Manager, Systems Availability and Software Distribution (Premium Edition) for IBM servers:

- IBM BladeCenter
- IBM iSeries™
- IBM pSeries®
- IBM xSeries
- IBM zSeries®

IBM Redbooks™

www.ibm.com/redbooks/

You can download the following documents from the IBM Redbooks Web page. You also might want to search this Web page for documents that focus on specific IBM hardware; such documents often contain systems-management material.

Note: Be sure to note the date of publication and to determine the level of IBM Director software to which the Redbooks publication refers.

- *Creating a Report of the Tables in the IBM Director 4.1 Database* (TIPS0185)
- *IBM Director Security* (REDP-0417-00)
- *IBM eServer™ BladeCenter Systems Management with IBM Director V4.1 and Remote Deployment Manager V4.1* (REDP-3776-00)
- *Implementing Systems Management Solutions using IBM Director* (SG24-6188)

- *Integrating IBM Director with Enterprise Management Solutions* (SG24-5388)
- *Managing IBM TotalStorage® NAS with IBM Director* (SG24-6830)
- *Monitoring Redundant Uninterruptible Power Supplies Using IBM Director* (REDP-3827-00)

Remote Supervisor Adapter

Remote Supervisor Adapter overview
www.ibm.com/support/docview.wss?uid=psg1MIGR-4UKSML

This Web page includes links to the *Remote Supervisor Adapter User's Guide* and *Remote Supervisor Adapter Installation Guide*.

Remote Supervisor Adapter II overview
www.ibm.com/support/docview.wss?uid=psg1MIGR-50116

This Web page includes information about the Remote Supervisor Adapter II.

Other documents

For planning purposes, the following documents might be of interest:

- *Planning and installation guide - IBM eServer BladeCenter (Type 8677)*
- *IBM Management Processor Command-Line Utility User's Guide version 3.00*

How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information. If you have any comments about this book or any other IBM Director publication, use the form for reader's comments is provided at the back of this publication. If the form has been removed, you may address your comments to:

International Business Machines Corporation
Design & Information Development
Department CGFA
PO Box 12195
Research Triangle Park, NC 27709-9990
U.S.A.

What’s new in release 5.10

This topic provides information about new features and enhancements in IBM Director 5.10.

Improvements to how you work in IBM Director

Enhanced user interface

IBM Director 5.10 contains the following features designed to improve usability:

- Enhanced user interface allows more intuitive hierarchical viewing of managed objects in a single-pane view, the ″classic″ three-pane view, or a two-pane combination view
- New customizable details view of managed objects
- Toolbar is now customizable for users

Event Action Plan wizard

The Event Action Plan wizard now can be launched from IBM Director Console and used to edit existing event action plans. In addition, the wizard has been redesigned to improve usability and render it more powerful. You can specify additional event filters, the systems to which you want to apply the event action plan, and schedule when the event filters are applied.

Improved accessibility

IBM Director 5.10 meets the accessibility standards for Section 508 of the US Rehabilitation Act. The major changes to the product include the following:

- IBM Director Console can be navigated using the keyboard only.
- IBM Director Console includes ″Accessibility Preferences″ that enable users to customize such display attributes as color, font size, and contrast.
- IBM Director Console implements the Java™ accessibility API which supports interaction with assistive technology.
- The IBM Director documentation is delivered in a Web-based Information Center.

New command-line interface

The **dircmd** command-line interface is deprecated in favor of a new command-line interface: **dircli**. The **dircli** command-line interface supports existing **dircmd** bundles, plus offers a new set of commands for accomplishing common system-management tasks.

Server Configuration Manager

New task to create or update server configuration profiles. Configuration includes the service processors in @server xSeries servers.

Software Health

New task to check for outdated firmware, drivers, and director agents on managed objects.

Unattended installation

IBM Director Server now can be installed in unattended mode.

Web-based (Information Center) product documentation

New in version 5.10, the IBM Director information center is a comprehensive, browser-based information system that provides easy access to the most up-to-date product information available. Updated periodically, the IBM Director information center contains:

- Assistance for the tasks that users must perform
- Conceptual information
- Reference for commands, extensions, icons, security, and many other topics
- Usage scenarios for IBM Director

To find information, users can search, browse the contents, follow links from one topic to related topics, and print the topics they want to read offline. The IBM Director information center is available at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Inventory enhancements

Inventory collection has been improved with the following new features in IBM Director 5.10:

- Filter queries can be designed for inventory data that has not been collected
- Enhanced tree navigation in the Inventory Query Browser
- Inventory change monitoring
- Custom collections of inventory tables
- Optional events on inventory completion or errors
- Improved control over inventory collection through additional preferences:
  - Enable/disable background inventory service
  - Specify the maximum number of agents to perform inventory collection at the same time, to control resource usage
  - Specify the default collection type for the three agent levels

Upward integration enhancements

IBM Director 5.10 includes the following enhancements to the upward integration modules (UIM):

- Support for Microsoft® Operations Manager (MOM)

Support for more systems in IBM Director

SMI-S storage devices

- Support for SMI-S 1.1 compliant Storage Managed Systems, including the IBM System Storage DS300 and DS400 devices
- Information displayed in the Hardware Status task and events for all supported SMI-S storage devices

ServeRAID™ hardware and software

IBM Director 5.10 improves support for IBM ServeRAID controllers:

- Configuration Management Station on Linux®
- Support for VMware ESX Server, versions 2.1, 2.5, and 2.51, Console
- Support for VMware ESX Server, versions 2.1, 2.5, and 2.51, guest operating systems
- Support for Red Hat Enterprise Linux AS, ES, and WS, version 4.0

- Support for ServeRAID Server and Console installations on EM64T and AMD64
- Hardware support for the IBM ServeRAID-8i controller
- Separate installation for the ServeRAID Manager extension

For complete support information, refer to the IBM ServerProven page at www.ibm.com/pc/us/compat/index.html.

Architectural support for more systems

IBM Director 5.10 contains a significant change in the product architecture. IBM Director Server now can manage three different types of managed systems:

Level-0 (″agentless″) managed systems

IBM Director manages these systems through the network services that are native to the operating system: SMB/CIFS/DCE-RPC protocols for Windows® systems, or Secure Shell (SSH) for other systems. No IBM Director software is installed. You can perform the following tasks on these managed systems from the IBM Director Console:

- Collect inventory that is available from the operating system
- Install IBM Director Core Services (Level 1) or IBM Director Agent (Level 2)
- Reboot the operating system (Windows or Linux)
- Use Remote Session task to execute command-line programs (only if SSH is present)
- Shutdown/power-off systems (Windows)

Level-1 managed systems

IBM Director Core Services must be present. In addition to the tasks supported by Level-0 managed systems, you can perform the following tasks on these managed systems from the IBM Director Console:

- Collect platform-specific inventory
- Install IBM Director Agent (promote to Level-2 managed system)
- Manage events using event action plans, event subscription, and the event log
- Monitor hardware status
- Reboot or shutdown the managed system
- Use Remote Session task to execute command-line programs (only if SSH is present)
- Distribute system-level update packages

Level-2 managed systems

IBM Director Agent must be installed. You can perform the full complement of IBM Director tasks on the managed system.

Additional systems supported for IBM Director Server and IBM Director Console installation

xSeries servers

- Red Hat Enterprise Linux AS and ES, version 4.0, for AMD64 and EM64T
- Red Hat Enterprise Linux AS and ES, version 4.0, for Intel® x86
- SUSE LINUX Enterprise Server 9 for AMD64 and EM64T

- Windows Server 2003, Enterprise, Standard, and Web x64 Editions

iSeries™ servers

- AIX 5L, Version 5.3
- Red Hat Enterprise Linux AS, version 4.0, for IBM POWER
- SUSE LINUX Enterprise Server 9 for IBM POWER

System p5 and pSeries servers

- AIX 5L, Version 5.3
- Red Hat Enterprise Linux AS, version 4.0, for IBM POWER
- SUSE LINUX Enterprise Server 9 for IBM POWER

System z9 and zSeries servers

- Red Hat Enterprise Linux AS, version 4.0, for IBM System z9, zSeries and S/390
- SUSE LINUX Enterprise Server 9 for IBM System z9, zSeries and S/390

Additional systems supported for IBM Director Agent installation

xSeries servers and Intel-compatible systems (32-bit operating systems)

- Novell NetWare, version 6.5
- VMware ESX Server, version 2.5, with the following guest operating systems:
  - Red Hat Enterprise Linux AS, ES, and WS, version 3.0 (Update 3 required)
  - SUSE LINUX Enterprise Server 8 for x86 (Service Pack 3 required)
  - SUSE LINUX Enterprise Server 9 for x86
  - Windows 2000, Advanced Server and Server Editions (Service Pack 3 or later required)
  - Windows Server 2003, Enterprise, Standard, and Web Editions (Service Pack 1 required)
  - Windows XP Professional Edition (Service Packs 1 and 2 required)
- VMware ESX Server, version 2.51, with the following guest operating systems:
  - Red Hat Enterprise Linux AS, ES, and WS, version 3.0 (Update 4 required)
  - SUSE LINUX Enterprise Server 8 for x86 (Service Pack 3 required)
  - SUSE LINUX Enterprise Server 9 for x86 (Service Pack 1 required)
  - Windows 2000, Advanced Server and Server Editions (Service Pack 3 or later required)
  - Windows Server 2003, Enterprise, Standard, and Web Editions (Service Pack 1 required)
  - Windows XP Professional Edition (Service Packs 1 and 2 required)
- Microsoft Virtual Server 2005 with the following guest operating systems:
  - Windows 2000, Advanced Server and Server Editions (Service Pack 3 or 4 required)
  - Windows Server 2003, Enterprise, Standard, and Web Editions
- Microsoft Virtual Server 2005 (Service Pack 1) with the following guest operating systems:

  - Windows 2000, Advanced Server and Server Editions (Service Pack 3 or 4 required)
  - Windows Server 2003, Enterprise, Standard, and Web Editions
  - Windows Server 2003, Enterprise, Standard, and Web x64 Editions
  - Windows XP Professional Edition (Service Pack 2 required)
  - Windows XP Professional x64 Edition

xSeries servers and Intel-compatible systems (64-bit operating systems)

- Red Hat Enterprise Linux AS, version 4.0, for Intel Itanium
- Windows Server 2003, Datacenter, Enterprise, Standard, and Web x64 Editions
- Windows Server 2003, Datacenter and Enterprise 64-bit Itanium Editions

iSeries servers

- Red Hat Enterprise Linux AS, version 4.0, for IBM POWER

iSeries servers with xSeries options

iSeries server installations can use the following xSeries options:

- Integrated xSeries Server (ISX)
- xSeries servers that are attached to the iSeries servers via the Integrated xSeries Adapter (IXA)

Using these xSeries options, you can install IBM Director Agent and IBM Director Core Services on the following operating systems:

- Red Hat Enterprise Linux AS and ES, version 3.0, for Intel x86
- Red Hat Enterprise Linux AS and ES, version 4.0, for Intel x86
- SUSE LINUX Enterprise Server 8 for x86
- SUSE LINUX Enterprise Server 9 for x86
- Windows 2000, Advanced Server and Server Editions
- Windows Server 2003, Enterprise, Standard, and Web Editions

Note: Whether these operating systems are supported in your iSeries environment depends on the following criteria:

- The Integrated xSeries Server (ISX) installed in the iSeries server
- The xSeries server that is attached to the iSeries server via the Integrated xSeries Adapter (IXA)
- The release of i5/OS or OS/400 installed on the iSeries server

For more information, see *IBM Director Hardware and Software Compatibility*. You can download this document from www.ibm.com/servers/eserver/xseries/systems_management/ibm_director/.

System p5 and pSeries servers

- Red Hat Enterprise Linux AS, version 3.3, for IBM POWER
- Red Hat Enterprise Linux AS, version 4.0, for IBM POWER

System z9 and zSeries servers

- Red Hat Enterprise Linux AS, version 4.0, for IBM System z9, zSeries and S/390
- SUSE LINUX Enterprise Server 9 for IBM System z9, zSeries and S/390

z/VM® Center management

z/VM Center is a new extension for provisioning Linux systems on virtual hardware that is based on real IBM System z9 and zSeries hardware and the z/VM hypervisor. z/VM Center provides two tasks:

Virtual Server Deployment

- Create and maintain templates for virtual hardware.
- Create and maintain templates for Linux systems.
- Create and delete virtual hardware.
- Create and delete instances of the Linux operating system.

Server Complexes

- Use the templates from Virtual Server Deployment to create virtual hardware and Linux instances on this hardware in a single step.
- Manage configurations of Linux instances and virtual hardware. Configuration domains include network settings, Linux configuration scripts, disk access, and VM Resource Manager (VMRM) performance goals.
- Apply configuration changes across multiple Linux instances.

Security enhancements

Security

IBM Director security has been improved with the following changes:

- AES support for UDP encryption
- Auditing on the server
- For new installations, security settings, including console-server SSL, are on or selected by default
- On Windows only, group administration of privileges
- PAM authentication support on UNIX®
- User-authenticated **dircli** command-line interface to replace **dircmd**

Other enhancements

IBM Director 5.10 includes the following general enhancements:

- Apache Derby is now the default IBM Director database that is bundled with the product. It is supported on all the operating system on which IBM Director Server can be installed, with the exception of i5/OS.
- Changed the name of the default group ″All Systems and Devices″ to ″All Managed Objects″; this group now has a default association, ″System Membership,″ that associates systems with their platforms.
- Replace the timestamp on the status bar with the number of managed objects displayed in the Group Contents pane
- Support for IBM Java Runtime Environment (JRE) 1.4.2, server release 2
- User-selected associations are persisted per group

Discontinued features in release 5.10

BladeCenter management

The BladeCenter Assistant task has been replaced with the BladeCenter Management task.

DMI Browser

The DMI Browser task has been removed.

Management Processor Assistant task

The Management Processor Assistant task has been replaced by the new Server Configuration Manager task.

Microsoft Management Console (MMC)

Microsoft Management Console (MMC) is no longer supported as of release 5.10.

Server Plus Pack

The Server Plus Pack has been withdrawn; however, some of its components are still available:

- Capacity Manager is separately available for purchase for IBM @server xSeries systems
- Rack Manager is now part of the base installation of IBM Director
- System Availability is available from the IBM Web site as a separate, installable extension

Active PCI Manager and Software Rejuvenation are not supported in release 5.10, and if installed for a previous version of IBM Director, they will be uninstalled when IBM Director is upgraded to version 5.10.

Web-based Access

Web-based Access has been removed from the base installation of IBM Director. It is available from the IBM Web site as a separate, installable extension.

Chapter 1. Getting started

Introducing IBM Director

This topic provides an overview of IBM Director.

IBM Director is an integrated, easy-to-use suite of tools that provide you with comprehensive systems-management capabilities to help realize maximum system availability and lower IT costs. Its open, industry-standard design enables heterogeneous-hardware management and broad operating-system support, including most Intel microprocessor-based systems and certain IBM @server System p5®, iSeries, pSeries, System z9®, and zSeries servers.

IBM Director automates many of the processes that are required to manage systems proactively, including capacity planning, asset tracking, preventive maintenance, diagnostic monitoring, troubleshooting, and more. It has a graphical user interface that provides easy access to both local and remote systems.

IBM Director can be used in environments with multiple operating systems and integrated with robust workgroup and enterprise management software from IBM (such as Tivoli® software), Computer Associates, Hewlett-Packard, Microsoft, NetIQ, and BMC Software.

IBM Director environment

IBM Director is designed to manage a complex environment that contains numerous servers, desktop computers, workstations, mobile computers (notebook computers), and assorted devices. IBM Director can manage up to 5000 Level-2 systems.

An IBM Director environment contains the following groups of hardware:

- One or more servers on which IBM Director Server is installed. Such servers are called *management servers*.
- Servers, workstations, desktop computers, and mobile computers that are managed by IBM Director. Such systems are called *managed systems*.
- Network devices, printers, or computers that have Simple Network Management Protocol (SNMP) agents installed or embedded. Such devices are called *SNMP devices*.
- Additional managed objects such as platforms and chassis. Collectively, all managed systems, devices, and objects are referred to as *managed objects*.

Figure 1 on page 2 shows the hardware in an IBM Director environment.

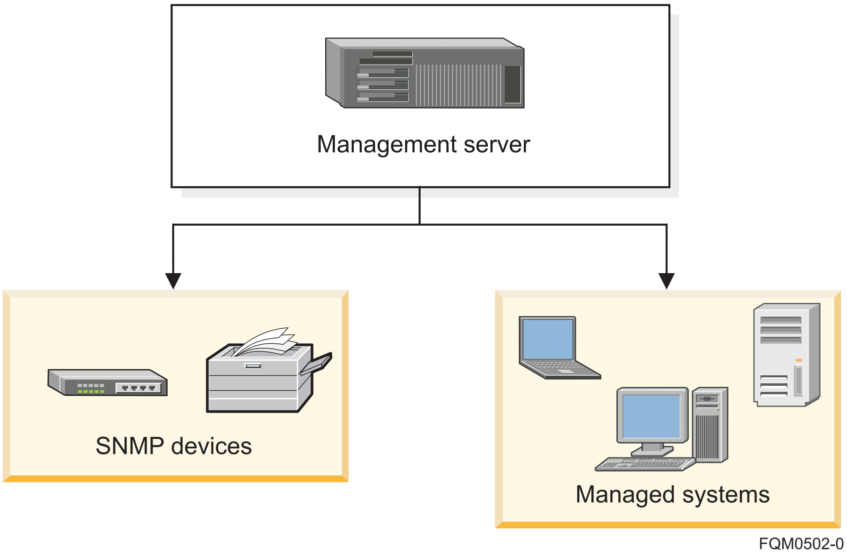

Figure 1. Hardware in an IBM Director environment

Using IBM Director Console

You can use IBM Director Console to group managed objects, view associations, start tasks, and set IBM Director options and preferences.

The IBM Director Console interface

This topic describes the main components of the IBM Director Console interface.

Before you begin using IBM Director Console, review the layout of its interface.

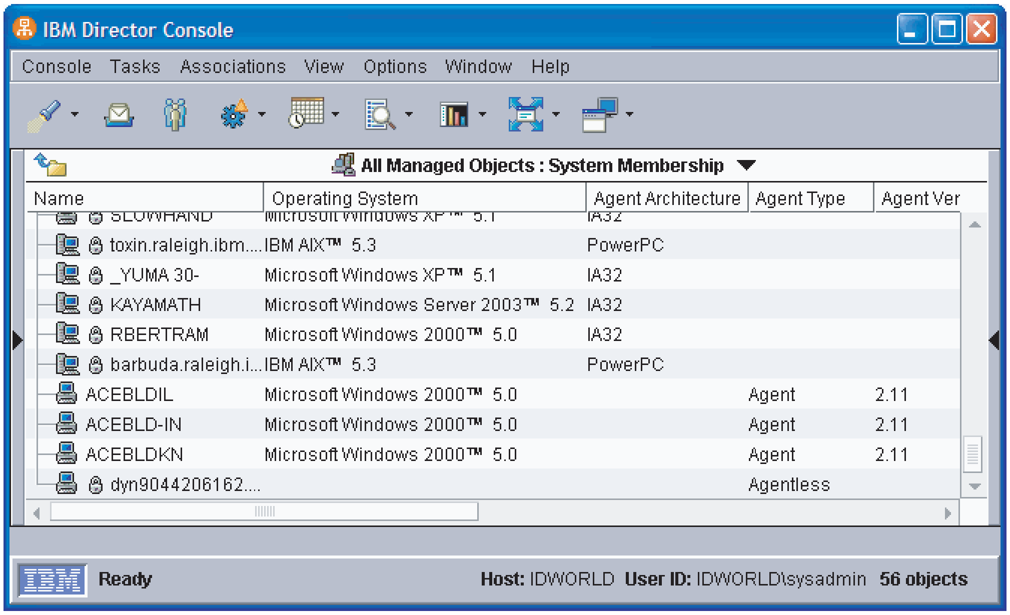

Name	Operating System	Agent Architecture	Agent Type	Agent Ver
SLOWHAND	Microsoft Windows XP™ 5.1	IA32		
toxin.raleigh.ibm....	IBM AIX™ 5.3	PowerPC		
_YUMA 30-	Microsoft Windows XP™ 5.1	IA32		
KAYAMATH	Microsoft Windows Server 2003™ 5.2	IA32		
RBERTRAM	Microsoft Windows 2000™ 5.0	IA32		
barbuda.raleigh.i...	IBM AIX™ 5.3	PowerPC		
ACEBLDIL	Microsoft Windows 2000™ 5.0		Agent	2.11
ACEBLD-IN	Microsoft Windows 2000™ 5.0		Agent	2.11
ACEBLDKN	Microsoft Windows 2000™ 5.0		Agent	2.11
dyn9044206162....			Agentless	

Ready **Host:** IDWORLD **User ID:** IDWORLD\sysadmin **56 objects**

Figure 2. IBM Director Console

Along the top of the IBM Director Console interface is both a menu and a toolbar.

Below the menu and toolbar, one, two, or three panes will be visible. From left to right, these are:

- Groups
- Group contents (pane title indicates selected group)
- Tasks

Below the panes is the marquee area and hardware-status alert display. The ticker-tape messages scroll across the marquee area. The hardware-status alert display is located in the bottom-right corner of the interface.

At the bottom of the IBM Director Console interface is a status bar indicating the ready or busy status of IBM Director, the host and login information for IBM Director Server, and the number of managed objects in the Group Contents pane.

IBM Director Console is usually accessed on the management server. On Windows management servers, an additional, and separate, Server Status icon displays the ready or busy status of IBM Director Server in the Windows system tray.

Note: Icons in the Windows tray, including the IBM Director Server Status icon, may disappear when a user selects the Windows High Contrast display. To enable display of the Server Status icon after selecting the high contrast display, select **All Programs** → **Startup** → **Server Status** from the Start menu.

IBM Director Console toolbar

This topic describes each of the icons on the IBM Director Console.

Along the top of the IBM Director Console interface is a toolbar containing nine icons.

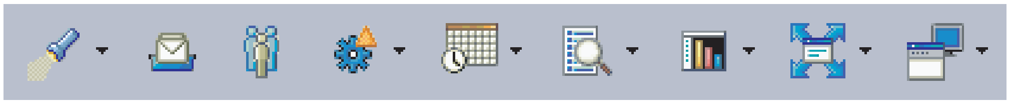

Figure 3. IBM Director Console toolbar

From left to right, the icons are briefly described below:

Discover
Clicking the button for this icon starts discovery of all systems and devices. Clicking the menu arrow for this icon allows you to select a type of system or device to discover.

Message Browser
Clicking the button for this icon opens the Message Browser window.

User Administration
Clicking the button for this icon opens the User Administration window.

Event Action Plans
Clicking the button for this icon opens the Event Action Plan Builder window. Clicking the menu arrow for this icon allows you to select the Event Action Plan Wizard, event logging options, or help for event action plans.

Scheduler
Clicking the button for this icon opens the Scheduler window. Clicking the menu arrow for this icon allows you to open help for scheduler.

Inventory
Clicking the button for this icon opens the Inventory Query Browser window. Clicking the menu arrow for this icon allows you to build a custom query, edit the software dictionary, collect inventory, create custom collections, create or view inventory monitors, or open help for inventory.

Resource Monitors
Clicking the button for this icon opens the Resource Monitors window. Clicking the menu arrow for this icon allows you to import a plan file, open the All Available Recordings or All Available Thresholds window, or open help for resource monitors.

Software Distribution
Clicking the button for this icon opens the Software Distribution Manager window. Clicking the menu arrow for this icon allows you to manage file distribution servers, view package history, create a package category, open a software distribution package, or open help for software distribution.

Remote Control
Clicking the button for this icon opens the Remote Control window. Clicking the menu arrow for this icon allows you to open help for remote control.

Panes in IBM Director Console

The IBM Director Console interface includes three panes: Groups, Group Contents, and Tasks.

You can resize the panes by dragging the borders between them. You can hide either the Groups or Tasks pane by clicking on the border between that pane and the Group Contents pane.

Note: The Group Contents pane may not be hidden.

Groups
The Groups pane lists all the groups available, including the default groups and any groups you have defined.

- Clicking a group selects that group for certain tasks performed from the toolbar or the menu. It also selects that group in the Group Contents pane.
- Right-clicking some groups also displays a context menu allowing tasks to be performed on the group.

Group Contents
The Group Contents pane lists the managed objects included in the group selected in the Groups pane. The title of the Group Contents pane indicates which group is selected.

- Clicking the title opens a menu from which you can select a group to display.
- Clicking a listed object selects that object for certain tasks performed from the toolbar or the menu.
- Right-clicking an object selects that object and displays a context menu allowing tasks to be performed on the object.

The icon for each managed object indicates both the type of managed object and its online status: icons for online objects appear in color, while icons for offline objects appear in gray.

A padlock icon beside a managed object indicates that the object is secured and inventory information about the object cannot be collected. To request access to the object, right-click the managed object and click **Request Access**. By providing a valid user name that has local administrative rights to that managed object and password, you can access the system.

Notes:

1. (BladeCenter chassis and physical platforms only) The padlock icon is displayed if a valid login profile does not exist for the service processor. You can access the system using the **Request Access** action as above.
2. (ISMP systems only) You cannot log in to an ISMP directly, as it lacks a userid and password. Instead, connect out-of-band to an ISMP installed on an ASM interconnect network through a Remote Supervisor Adapter or Remote Supervisor Adapter II serving as the ASM gateway.
3. (ASM processor systems only) Use the Management Processor Assistant to configure an out-of-band path to the ASM processor system, then change the userid and password to request access the physical platform using IBM Director Console.

Tasks The Tasks pane lists tasks which can be performed in IBM Director.

Note: Although the list of tasks in the Task pane is static, not all tasks are available for all groups or managed objects.

Right-clicking blank space in any pane displays a context menu from which you can change the pane's appearance or sorting, or perform tasks specific to that pane. For example, in the Group Contents pane you can create new managed objects manually, find and view objects, or perform actions on the selected group.

Working with tables

This topic describes general procedures for viewing tabular information in IBM Director Console.

In IBM Director Console, information is often displayed in tables. You can customize the display of data in many of these tables in several ways.

Note: Not all of these actions may be available for all tabular views. These actions must be performed with a mouse or other pointing device, and are not available through the keyboard.

Sort table data on a column

Click a column header to sort the data in the table by the values in that column. Click the header again to change the sort order.

- An upward-pointing triangle symbol indicates the column is sorted in ascending order
- A downward-pointing triangle symbol indicates the column is sorted in descending order
- Some columns may have additional sort options that are displayed as parenthetical text in the column header

Resize table columns

Drag the border of a table column heading to resize it.

Rearrange table columns

Drag a table column heading left or right to a new column location to rearrange table columns.

In addition, the details view of the Group contents pane in the main IBM Director Console window may be customized by selecting what columns are displayed. Right-click a column heading and select Customize columns. The Console Preferences window opens to the Details View Preferences page. Select the columns you wish to view, then click OK.

Starting tasks

This topic describes how to start tasks in IBM Director.

You can start most tasks in IBM Director in four ways:

- Dragging a task from the tasks pane onto a managed object (or a managed group, in some cases)
- Dragging a managed object (or a managed group, in some cases) onto a task in the tasks pane
- Right-clicking a managed object (or managed group, in some cases)
- Selecting the managed object or group, then selecting a task from the menu bar

Throughout this documentation, only dragging a task onto a managed object or group is explained as the method of starting tasks, although you can use any of the methods.

Some IBM Director functions, such as the Event Action Plan Builder and Scheduler, may be started either from the menu bar or from the toolbar.

Note: When IBM Director Console is processing a task, the hourglass is displayed for that window and you cannot use the mouse to work with the window. Although it might be possible to work with the window using key strokes, do *not* do so.

Concepts

This section discusses concepts that will help you understand how IBM Director works. Becoming familiar with the IBM Director components and understanding the concepts in this section enables you to use IBM Director most effectively.

Accessibility

This topic describes the accessibility features in IBM Director.

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. These are the major accessibility features in IBM Director:

- You can use screen-reader software to hear what is displayed on the screen.
- You can operate all features using the keyboard instead of the mouse.
- You can choose from a variety of high-contrast color schemes and large font sizes in the IBM Director Console.

Keyboard shortcuts

You can use keys or key combinations to perform operations that can also be done through mouse actions.

Keyboard shortcuts for windows, frames, panes, and icons:

You can use keys or key combinations to navigate windows, frames, panes, and icons in the IBM Director Console interface.

Window

Action	Keyboard shortcut
Activate the default button.	Enter

Option pane

Action	Keyboard shortcut
Navigate in or out of the option pane.	Alt+F6
Hide a dialog.	Esc
Active the default button (if defined).	Enter

Dialog

Action	Keyboard shortcut
Navigate out of the dialog.	Alt+F6
Hide the dialog.	Esc
Active the default button (if defined).	Enter

Scroll pane

Action	Keyboard shortcut
Navigate forward out of the scroll pane.	Tab
Navigate backward out of the scroll pane.	Shift+Tab
Move up or down.	Up arrow or down arrow
Move left or right.	Left arrow or right arrow.
Navigate to the beginning or end of data.	Ctrl+Home or Ctrl+End
Navigate up or down one block.	PgUp or PgDn
Navigate to the left or right.	Ctrl+PgUp or Ctrl+PgDn

Split pane

Action	Keyboard shortcut
Navigate forward out of the split pane.	Tab or Ctrl+Tab
Navigate backward out of the split pane.	Shift+Tab or Ctrl+Shift+Tab
Navigate between split panes.	Tab or F6
Navigate to the splitter bar.	F8
Toggle the focus between two split bars (for windows with three split panes).	F8
Resize the split pane vertically.	Up arrow or down arrow
Resize the split pane horizontally.	Left arrow or right arrow

Action	Keyboard shortcut
Maximize the size of the split pane .	Home
Minimize the size of the split pane.	End

Notebook (tabbed pane)

Action	Keyboard shortcut
Navigate into the tabbed pane.	Tab
Navigate out of the tabbed pane.	Ctrl+Tab
Navigate to the left or right tab.	Left arrow or right arrow
Navigate to the tab above or below.	Up arrow or down arrow
Navigate from the tab to the page.	Enter or Ctrl+Down
Navigate from the page to the tab.	Ctrl+Up
Navigate to the previous or next page.	Ctrl+PgUp or Ctrl+PgDn

Frame

Action	Keyboard shortcut
Display a window menu.	Alt+Spacebar
Active the default button (if defined).	Enter

Internal frame

Action	Keyboard shortcut
Open or restore the frame.	Ctrl+F5, Alt+F5, or Enter
Close the frame.	Ctrl+F4 or Alt+F5
Move the frame.	Ctrl+F7 or Alt+F7
Resize the frame.	Ctrl+F8 or Alt+F8
Minimize the frame size.	Ctrl+F9 or Alt+F9
Display a window menu.	Alt+Spacebar
Active the default button (if defined).	Enter

Keyboard shortcuts for the menu bar and toolbar:

You can use keys or key combinations to navigate standard controls in the IBM Director Console interface.

Menu bar

Action	Keyboard shortcut
Jump to the menu bar.	Alt or F10
Navigate out of the menu bar.	Esc or Alt
Navigate within the menu bar.	Arrow keys
Select the next or previous menu item.	Right arrow or left arrow
Activate the default or selected item.	Enter

Action	Keyboard shortcut
Display a menu.	Use one of these keyboard shortcuts: • Up arrow • Down arrow • Enter • Spacebar • Alt+Character accelerator key (if defined)
Hide a menu.	Esc or Alt

Menu

Action	Keyboard shortcut
Display a menu.	Enter or F10
Display a submenu.	Right arrow
Navigate to the next item or wrap to the top.	Down arrow
Navigate to the previous item or wrap to the bottom.	Up arrow
Hide the menu.	Esc
Hide the submenu.	Left arrow
Active the default or selected item.	Enter

Menu items

Action	Keyboard shortcut
Navigate in or out of a menu.	Arrow keys
Activate an item.	Enter, spacebar, or Alt+Character accelerator key (if defined)
Display a submenu.	Right arrow
Hide a submenu.	Left arrow or Esc

Check box menu items

Action	Keyboard shortcut
Navigate in or out of the check box menu.	Arrow keys
Select or clear a check box menu item.	Enter
Hide a check box menu.	Enter

Radio button menu items

Action	Keyboard shortcut
Navigate in or out of a radio button menu.	Arrow keys
Select or clear a radio button menu item.	Enter
Hide a radio button menu.	Enter

Pop-up menus

Action	Keyboard shortcut
Display a pop-up menu.	Shift+F10
Display a pop-up submenu.	Right arrow
Hide a pop-up menu.	Esc
Hide a submenu.	Left arrow
Navigate within a pop-up menu.	Up arrow or down arrow
Activate a pop-up menu item.	Enter or spacebar

Toolbar

Action	Keyboard shortcut
Navigate forward out of the toolbar.	Tab
Navigate backward out of the toolbar.	Shift+Tab
Navigate within the toolbar.	Arrow keys
Active a toolbar item.	Enter
Display the Customized Toolbar menu (when focus is on an icon on the main IBM Director Console window toolbar).	Shift+10

Tool tips

Action	Keyboard shortcut
Display a tool tip.	Ctrl+F1
Hide a tool tip.	Esc or Ctrl+F1

Keyboard shortcuts for standard interface controls:

You can use keys or key combinations to navigate standard controls in the IBM Director Console interface.

Buttons

Action	Keyboard shortcut
Navigate forward.	Tab
Navigate backward.	Shift+Tab
Activate the default button.	Enter
Activate any button	Spacebar or Alt+Character accelerator key (if defined).
Activate Cancel or Close.	Esc

Check boxes

Action	Keyboard shortcut
Navigate forward.	Tab
Navigate backward.	Shift+Tab
Navigate within a group.	Arrow keys
Select or clear a check box.	Spacebar

Radio buttons

Action	Keyboard shortcut
Navigate forward.	Tab
Navigate backward.	Shift+Tab
Navigate within a group.	Arrow keys **Note:** To select the radio button, navigate to it.
Select or clear a radio button.	Spacebar

Combination boxes

Action	Keyboard shortcut
Navigate forward out of the combination box.	Tab
Navigate backward out of the combination box.	Shift+Tab
Display the drop-down list.	Alt+Down arrow
Hide the drop-down list.	Esc or Alt+Up arrow
Active the selected menu item.	Enter
Navigate up or down the drop-down list.	Alt+Up arrow or Alt+Down arrow
Navigate to a list item without selecting it.	Initial character of the list item
Move up or down the drop-down list.	Up arrow or down arrow

Lists

Action	Keyboard shortcut
Navigate forward out of the list.	Tab
Navigate backward out of the list.	Shift+Tab
Activate the selected list item.	Enter
Navigate within the list.	Up arrow or down arrow
Navigate to the beginning or end of the list.	Ctrl+Home or Ctrl+End
Select all list items.	Ctrl+A
Select a single list item	Spacebar **Note:** Using the spacebar clears the previous selection.
Select an additional list item.	Ctrl+Spacebar

Action	Keyboard shortcut
Select a range of list items.	Shift+Spacebar
Extend the selection up or down one item.	Shift+Up arrow or Shift+Down arrow
Extend the selection to the top or bottom of the list.	Shift+Home or Shift+End
Extend the selection up or down one block.	Shift+PgUp or Shift+PgDn
Navigate up or down a block.	PgUp or PgDn

Sliders

Action	Keyboard shortcut
Navigate forward out of the slider.	Tab
Navigate backward out of the slider.	Shift+Tab
Increase the value	Up arrow or right arrow.
Decrease the value	Down arrow or left arrow.
Set the maximum value.	Home
Set the minimum value.	End
Increase the value by a set range.	PgUp
Decrease the value by a set range.	PgDn

Tables

Action	Keyboard shortcut
Navigate forward out of the table.	Ctrl+Tab
Navigate backward out of the table.	Ctrl+Shift+Tab
Navigate to the next cell.	Tab or right arrow
Navigate to the previous cell.	Shift+Tab or left arrow
Navigate to the next row from the last column.	Tab or right arrow
Navigate to the previous row from the first column.	Shift+Tab or left arrow
Navigate vertically to the next or previous block.	PgUp or PgDn
Navigate horizontally to the left or right one block.	Ctrl+PgUp or Ctrl+PgDn
Navigate to the first or last cell in the row.	Home or End
Navigate to the first or last cell in the table.	Ctrl+Home or Ctrl+End
Select all cells in the table.	Ctrl+A
Clear the current selection.	Use one of these keyboard shortcuts: • Up arrow or down arrow • Ctrl+Up arrow or Ctrl+Down arrow • PgUp or PgDn • Ctrl+PgUp or Ctrl+PgUp • Home or End • Ctrl+Home or Ctrl+End
Extend the selection up or down one row.	Shift+Up arrow or Shift+Down arrow

Action	Keyboard shortcut
Extend the selection to the right or left one column.	Shift+Left arrow or Shift+Right arrow
Extend the selection to the beginning or end of the row.	Shift+Home or Shift+End
Extend the selection up or down one block.	Shift+PgUp or Shift+PgDn
Extend the selection left or right one block.	Ctrl+Shift+PgUp or Ctrl+Shift+PgDn
Extend the selection to the beginning or end of the column.	Ctrl+Shift+Home or Ctrl+Shift+End
Edit the cell without overriding the existing text.	F2
Delete the cell text before editing.	Esc

Trees

Action	Keyboard shortcut
Navigate forward out of the tree.	Tab
Navigate backward out of the tree.	Shift+Tab
Expand the entry	Right arrow or Enter (if collapsed).
Collapse the entry	Left arrow or Enter (if expanded).
Navigate up or down one entry.	Up arrow or down arrow
Navigate to the first entry in the tree.	Home
Navigate to the last visible entry in the tree.	End
Navigate vertically up or down one block.	PgUp or PgDn
Navigate to the left or right one block.	Ctrl+PgUp or Ctrl+PgDn
Select all entries.	Ctrl+A or Ctrl+Slash
Clear the selection.	Ctrl+\
Select a single entry.	Ctrl+Spacebar
Select a range of entries.	Shift+Spacebar
Extend the selection up or down one block.	Shift+PgUp or Shift+PgDn
Extend the selection to the top of the tree.	Shift+Home
Extend the selection to the bottom of the tree.	Shift+End

Keyboard shortcuts for text components:

You can use keys or key combinations to navigate text components in the IBM Director Console interface.

Text fields

Action	Keyboard shortcut
Navigate into the text field	Alt+Character accelerator key (if defined).
Navigate forward out of the text field.	Tab
Navigate backward out of the text field.	Shift+Tab

Action	Keyboard shortcut
Navigate to the previous or next character.	Left arrow or right arrow
Navigate to the previous or next word.	Ctrl+Left arrow or Ctrl+Right arrow
Navigate to the beginning or end of a field.	Home or End
Submit an entry.	Enter
Select all text in the field.	Ctrl+A
Clear the selection.	Arrow keys
Extend the selection to the left or right one character.	Shift+Left arrow or Shift+Right arrow
Extend the selection to the beginning or end of the field.	Shift+Home or Shift+End
Extend the selection to the next or previous word.	Ctrl+Shift+Left arrow or Ctrl+Shift+Right arrow
Copy the selected text.	Ctrl+C
Cut the selected text.	Ctrl+X
Paste from the clipboard.	Ctrl+V
Delete the previous or next character	Backspace or Delete

Text panes

Action	Keyboard shortcut
Navigate into the text pane	Tab or Alt+Character accelerator key (if defined).
Navigate forward out of the text pane.	Ctrl+Tab
Navigate backward out of the text pane.	Ctrl+Shift+Tab
Navigate vertically up or down one block.	PgUp or PgDn
Navigate up or down one line.	Up arrow or down arrow
Navigate to the left or right one component or character.	Left arrow or right arrow
Navigate to the beginning or end of a line.	Home or End
Navigate to the previous or next word.	Ctrl+Left arrow or Ctrl+Right arrow
Navigate to the beginning or end of the text pane.	Ctrl+Home or Ctrl+End
Navigate up or down one block.	PgUp or PgDn
Navigate to the left or right one block.	Ctrl+PgUp or Ctrl+PgDn
Navigate to the next or previous HTML link or other focusable element.	Ctrl+T or Ctrl+Shift+T
Navigate out of a focusable element that accepts a tab.	Ctrl+Tab or Ctrl+Shift+Tab
Activate a hyperlink.	Ctrl+Spacebar
Extend the selection up or down one block.	Shift+PgUp or Shift+PgDn
Extend the selection to the left or right one block.	Ctrl+Shift+PgUp or Ctrl+Shift+PgDn
Extend the selection up or down one line.	Shift+Up arrow or Shift+Down arrow
Extend the selection to the left or right.	Shift+Left arrow or Shift+Right arrow

Action	Keyboard shortcut
Extend the selection to the beginning or end of the line.	Shift+Home or Shift+End
Extend the selection to the beginning or end of the text pane.	Ctrl+Shift+Home or Ctrl+Shift+End
Extend the selection to the previous or next word.	Ctrl+Shift+Left arrow or Ctrl+Shift+Right arrow
Extend the selection vertically up or down one block.	Shift+PgUp or Shift+PgDn
Extend the selection to the left or right one block.	Ctrl+Shift+PgUp or Ctrl+Shift+PgDn
Select all text in the text pane.	Ctrl+A
Clear the selection.	Arrow keys
Copy the selected text.	Ctrl+C
Cut the selected text.	Ctrl+X
Paste from the clipboard.	Ctrl+V
Delete the previous or next component or character.	Backspace or Delete
Insert a line break.	Enter
Insert a tab.	Tab

Configuring IBM Director Console appearance

The Appearance Preferences page enables you to customize the look of your IBM Director Console. You can set the color for text, backgrounds, and links. You can choose a background image and decide whether to show a shadow.

Complete the following steps to customize the background for IBM Director Console:

1. On the IBM Director Console window, select **Options** → **Console Preferences**.
2. On the Console Preferences window, click the **Appearance** tab.
3. Select the text and background colors for IBM Director Console GUI components. You can also select shadow settings and a background image.
4. Click **OK**.

Configuring IBM Director Console colors and fonts

The Accessibility Preferences page allows you to customize your IBM Director Console colors and fonts.

Note:

- If you change the **Accessibility Preferences** while other windows are visible, the window might not be displayed correctly after the change. If this occurs, close and reopen the particular window to fix the problem.
- Operating System changes in font, color and size can be reflected in both the title bar and the client area of the application. IBM Director Console settings for font size and color affect only the client area of the application.

To customize your IBM Director Console colors and fonts, complete the following steps:

1. On the IBM Director Console window, select **Options** → **Console Preferences**.
2. On the Console Preferences window, click **Accessibility Preferences**.
3. Select the color and font settings that you want.

4. Click **OK**.

Associations

Associations change the organization of a group of managed objects that is displayed in the Group Contents pane.

You can apply an association type to the group that is currently displayed in the Group Contents pane. When you apply an association, the association persists the next time you display that group.

If the group that is currently displayed in the Group Contents pane contains managed objects that do not apply to the selected association, those objects appear in blue type under the “Not Associated” node.

You also can display additional information about the managed objects that are displayed in the Group Contents pane by selecting one or more association options from the bottom half of the **Associations** menu. For example, you can view managed objects to which event action plans have been applied. If a managed object has an event action plan applied to it, the managed object is displayed as a tree structure that you can expand to view which event action plans have been applied to the object.

Association type

The following associations are available:

None

System Membership
Shows the relationship between Level-0, Level-1, and Level-2 managed systems and logical and physical platforms, with systems the top level and the associated platforms as child nodes. This is the default association for the All Managed Objects group.

Object Type
Shows the managed objects based on object type (such as managed systems, SNMP devices, and chassis).

TCP/IP Addresses
Shows the managed objects based on TCP/IP address.

TCP/IP Host Names
Shows the managed objects based on TCP/IP host names.

IPX Network IDs
Shows the managed objects based on network IDs.

Domains/Workgroups
Shows the managed objects based on domains and workgroups.

Chassis Membership
Shows all the blade servers in a BladeCenter chassis. This is the default association for the Chassis and Chassis Members group.

Cluster Membership
This is the default association for the Clusters and Cluster Members group.

Physical Platform–Remote I/O Enclosures
Shows the managed objects based on remote input/output (I/O) enclosures.

Platform Membership
Shows the relationship between managed systems and platforms, with platforms the top level and the associated systems as child nodes. This association is useful if you have a single system that represents multiple managed objects. Depending on the IBM Director task you want to perform, the managed object that you target might differ. This is the default association for the Platforms and Platform Members group.

Rack Membership
Shows all the managed objects in a rack. This is the default association for the Racks with Members group.

Scalable Partitions Membership
This is the default association for the Scalable Partitions group.

Scalable Systems Membership
This is the default association for the Scalable Systems and Members group.

TCP/IP Routers/DNS
Shows the managed objects based on TCP/IP routers or domain name space (DNS).

Status Shows the managed objects based on status.

SNMP System Object ID
Shows the managed objects based on SNMP system object ID.

HMC Membership
This is the default association for the HMC and HMC Members group.

z/VM Server Complexes Membership
Shows all z/VM systems with their server complexes, with included Linux guest systems. At the top level, the association shows the z/VM systems. Under each z/VM system, it shows the server complexes. Under each server complex, the association shows the tiers in that server complex. Under each tier, the association shows the Linux guest systems in that tier. The association also includes a “Free systems” node (at the second level) for Linux guest systems in the z/VM system that are not in any server complex.

Linux on System z9 and zSeries Platform Membership
Shows all discovered Level-0, Level-1, and Level-2 managed Linux systems that run on a System z9 or zSeries mainframe. The association tree shows which Linux systems run natively in a logical partition (LPAR) and which of the other Linux systems run under which z/VM system. Linux systems that run under an unknown z/VM system are grouped accordingly.

The association also shows z/VM manageability access points. A *z/VM manageability access point* is a Linux system that has been set up to enable the z/VM Center task for a particular z/VM system. In the association tree, z/VM manageability access points appear twice, as a Linux system under a z/VM system and as the z/VM manageability access point for that z/VM system.

Systems that are not below the LPAR or z/VM subtrees cannot be associated with an LPAR or a z/VM system, possibly because they are not Linux on System z9 and zSeries systems or because they are locked.

Default associations

The None association is the default association for most groups. This table shows the groups that have default associations other than None:

Table 1. Default associations other than None

Group	Default association
All Systems and Devices	System Membership
Chassis and Chassis Members	Chassis Membership
Clusters and Cluster Members	Cluster Membership
HMC and HMC Members	HMC Membership
Platforms and Platform Members	Platform Membership
Racks with Members	Rack Membership
Scalable Partition	Scalable Partitions Membership
Scalable Systems and Members	Scalable Systems Membership

Association options

The following association options are available:

Software Packages
Shows which software packages, if any, have been delivered to the managed objects in the group using the Software Distribution task.

Jobs Shows all tasks, if any, that are scheduled to be run against the managed objects in the group.

Activations
Shows all tasks, if any, that have already been run against each managed object in the group.

Resource Monitors
Shows the resource monitors, if any, that have been applied to the managed objects in the group.

Event Action Plans
Shows the event action plans, if any, that have been applied to the managed objects in the group.

Common Information Model

The Common Information Model (CIM) is a language-independent programming model that defines the properties, operations, and relationships of objects in enterprise and Internet environments. Using the CIM, IBM Director has a single model for communicating with these different resources. IBM Director uses the CIM to access data on Level-1 and Level-2 managed systems.

The CIM and Web Based Enterprise Management (WBEM) are standards that are developed by a consortium of major hardware and software vendors (including IBM) called the Distributed Management Task Force (DMTF). The CIM provides the framework by which a system can be managed by using common building blocks rather than proprietary software. If a device is CIM-compliant, software that is also CIM-compliant, such as IBM Director, can manage that device.

The infrastructure used by IBM Director for CIM instrumentation consists of the following:

CIM Client

The CIM Client is a management application that uses CIM to manage devices. A CIM Client can reside anywhere in the network, because it uses HTTP to talk to CIM Object Managers and Agents.

CIM Managed Object

A CIM Managed Object is a hardware or software component that can be managed by a management application by using CIM. When Level 2: IBM Director Agent or Level 1: Core Services is installed on a system, the applicable CIM software is installed and that system becomes a CIM Managed Object.

CIM Object Manager

The CIM Object Manager (CIMOM), also known as a CIM server, is the software entity that receives, validates, and authenticates the CIM requests from the CIM Client. It then directs the requests to the appropriate component or device provider.

IBM Director locates the CIMOM through discovery. When Level 1: Core Services is installed on a system, the CIMOM registers itself to the SLP and supplies its location, IP address, port number, and the type of service that it provides (management.software.IBM:director-agent). When IBM Director performs an SLP discovery, the IBM Director SLP service is identified and the system running that service is displayed as a Level-1 managed system. With this information, IBM Director can directly communicate with the CIMOM. Director discovers the CIMOM on Level-2 managed system using a proprietary protocol.

Using CIM as the framework, IBM Director can perform tasks on the managed system. As requests arrive, the CIMOM validates and authenticates each request. It then directs the requests to the appropriate functional component of the CIMOM or to a device provider. The provider makes calls to a device-unique programming interface on behalf of the CIMOM to satisfy the IBM Director requests.

Shortcuts to CIM classes and methods

By creating *shortcuts*, or subtasks, you can bypass navigating through the class tree to reach a specific class or method. You can define two types of shortcuts:

- A *user-selected class* that displays the instances, properties, and methods that are associated with a specified class on the selected managed system.
- A *user-selected method* that runs on the selected managed system.

Default CIM subscriptions for IBM Director Core Services

This topic provides conceptual information about the default CIM subscriptions for IBM Director Core Services.

The Level-1 managed system event architecture is governed by the WBEM suite of standards as implemented by the Pegasus CIM Object Manager version 2.5. This event architecture includes indication providers, indication handlers, and indication consumers.

Indication providers

Software modules managed by the CIMOM that monitor a resource in a computer system and send indications when some threshold has been exceeded for that resource.

Indication handlers

Software modules managed by the CIMOM which can export the indication to an interested subscriber.

Indication consumers
Software modules that subscribe to the CIM indications in which they are interested.

No indications are forwarded by the CIMOM unless there is an interested subscriber.

Level-1 IBM Director Core Services installs the following types of indication providers, indication handlers, and indication consumers:

- Indication providers that monitor hardware components in the managed system and send indications when there is a problem
- An indication handler that exports indications with outstanding subscriptions
- A set of indication consumers which can subscribe or unsubscribe to the indications

IBM Director Core Services also installs a helper service called a CIM listener. The listener receives events sent by the indication handler and ensures they are routed to the correct consumers.

The installation also creates a default set of subscriptions between the consumers and the indications.

Table 2. Default subscriptions for CIM indication consumers

CIM indication consumer	Default subscriptions
IBM Director events (CIM > System events)	All indications of all severities *after* the Level-1 managed system is discovered and unlocked by a management server
Hardware Status task (also the IBM Director Console Group Contents pane)	All indications of all severities except for Lease Expiration and Warranty Expiration
Microsoft Windows Event Log (event log event ID)	All indications of all severities
Local message window	None
Microsoft System Management Server (SMS)	All indications of all severities
SNMP (Tivoli NetView®, HP OpenView, CA Unicenter, Tivoli Enterprise Console®)	All indications of all severities
Microsoft Operations Manager 2005 (alerts)	All indications of all severities
Tivoli Enterprise Console (native events)	All indications of all severities

Additional reference information is available for IBM Director Core Services default subscriptions, for CIM indications, and for the cimsubscribe command used to modify subscriptions.

Discovery

This topic provides information about the IBM Director discovery process.

Discovery is the process by which IBM Director Server identifies and establishes connections with systems and devices that it can manage. The management server sends out a discovery request and waits for responses from managed systems. The managed systems listen for this request and respond to the management server that sent the request.

Before IBM Director can manage a system, that system must be discovered by IBM Director Server. After a system or device has been discovered, an icon that represents the object is displayed in the IBM Director Console window when the applicable group is selected.

Note: (Windows 2000, Server Edition only) The initial discovery performed by the management server is resource intensive. After the initial discovery is completed, the resource utilization returns to normal.

The type of discovery method that IBM Director uses to connect with systems and devices varies based on the type of device being discovered and the type of network protocols used by that device.

Managed systems and devices are classified into the following three levels:

- Level 0: Any system or device without IBM Director Agent or IBM Director Core Services installed and has SSH or DCOM/SMB running.
- Level 1: Any system with IBM CIM instrumentation installed. CIM instrumentation can be preloaded, as it is with AIX® or i5/OS, or it can be installed using the Level 1: IBM Director Core Services package.
- Level 2: Any system with IBM Director Agent installed, including System p5 and pSeries running AIX, IBM iSeries running IBM i5/OS, systems running Linux 32/64, systems running Windows 32/64, and Novell NetWare. This level is limited to 5,000 licenses.

During discovery, the management server searches for Level-0, Level-1, and Level-2 managed systems. The management server then stores addresses of those systems to the IBM Director database. To discover managed systems, IBM Director can either request information from the systems and devices that are accessible on the network, or the systems and devices can be configured to send an event to IBM Director, which will cause the management server to add that system to its database. This capability is only available for Level-2 managed systems and requires that the **Auto-add unknown agents which contact server** check box to be selected in the Level-2 Discovery Preferences window.

Depending on the complexity of your network, and the needs of your organization, you will need to configure the discovery process. You will need detailed information about the layout of your network, specifically the subnet, port, and LAN information.

The base-management server includes support for the following manageable entities or managed objects:

- System Endpoint
- Storage Devices that are optimized to support the SMI-S CIM standard for storage devices
- Physical or logical platforms (for example a physical machine or hypervisor) that can host an operating system
- SNMP (or generic) Devices
- Base Clusters
- Hypervisor services objects, such as the zSeries z/VM MO that create and delete Linux for System z9 and zSeries guest operating systems
- Hardware Control Points, such as the eServer BladeCenter chassis or the Power HMC, which are the management gateway to other physical platform objects

Level-0 and Level-1 managed systems can be added manually, much like IBM Director Systems. To manually add any Level-0, Level-1 or Level-2 managed system, right-click in the systems pane and select **New** → **System**. There are only two fields, network address and system name. The system name is optional. If you do not fill it in, the host name of the system is used. After you click OK, IBM Director Server attempts to discover what type of system (Level-0, Level-1, or Level-2) is represented by the address given in the Systems window. First, it checks for the Level-2 protocol (ports 14247 on Windows, 14248 on Linux). If the system does not respond, it checks for the Level-1 protocol (SLP port 427). If the system still does not respond, it checks for the Level-0 protocols (ssh port 22 or Windows RPC ports 137-139, 145). Once the server determines the system type, it creates a managed object. If none are detected, an error is reported.

Director supports agent-initiate and server-initiated discovery.

Agent initiated discovery

In IBM Director, *agent-initiated discovery* occurs when managed systems contact the IBM Director Server rather than IBM Director Server searching for managed systems. This is a push-based discovery.

This implementation has several advantages. An agent-initiated discovery will always succeed as long as there is a TCP/IP connection from the management server to the managed systems. Also, the network traffic due to discovery requests will be negligible, compared to a server-initiated discovery.

The two available agent-initiated discovery algorithms are:

- Agent-initiated discovery using the "Add known server" entry in a unattended install of IBM Director.
- Using the **genevent** command in a batch file, which sends an event to the management server with the managed systems name and IP address.

Server initiated discovery

In IBM Director, *server-initiated discovery* occurs when IBM Director Server searching the network for Level-0, Level-1, and Level-2 managed systems. Typically, this is referred to simply as discovery. The advantage of this type of discovery is ease of configuration. You set the discovery settings once at the management server. IBM Director offers several ways to perform server-initiated discovery:

- Broadcast discovery
- Multicast discovery
- Broadcast relay discovery
- Unicast discovery
- Service Location Protocol discovery

Broadcast discovery is the default discovery method. From IBM Director Console, you can change the discovery preferences for the management server according to your business needs and network requirements.

Broadcast discovery:

When you use broadcast discovery, the management server transmits a single request to the entire subnetwork on which the management server is located, in an effort to discover manageable objects on the network.

During a broadcast discovery, IBM Director Server sends out a broadcast request to all the IP addresses that are in a specified subnet. Typically, the subnet is one where the management server is installed, but you can also send a broadcast to other subnets if broadcasts are not filtered by your network infrastructure. (By default, most gateways do not permit broadcasts to pass over subnets.)

IBM Director Server broadcasts requests using each protocol supported by IBM Director (for example, SNMP, SSH, DCOM, SLP, and IPC). You also can use broadcast discovery on other subnets if your network gateway is configured to permit broadcast messages. Your network might be configured to allow broadcast messages within a subnet, but prevent them from passing over to other subnets. The IP addresses that successfully respond to the broadcast request are saved to IBM Director along with the applicable protocol information. IBM Director Console displays the objects that respond to the broadcast discovery and that support one or more of the protocols that IBM Director uses.

Note: By default, Level-0 managed systems are not discovered using broadcast discovery. Level-0 managed systems are discovered using unicast discovery, and an IP address range must be specified for the unicast.

Attention: Broadcast discovery consumes network resources. If IBM Director is configured to perform broadcast discoveries frequently, your network resources might be inefficiently used. As a best practice, contact your network administrator to determine the best discovery method for your organization.

Multicast discovery:

In multicast discovery, the management server sends a request to a specified IP address, called the *multicast group*. The multicast group that is used by IBM Director Server by default is 224.0.1.118.

Multicast discovery is used to identify Level-1 and Level-2 managed systems. This discovery method is useful on networks that are configured to filter broadcast requests but not multicast requests.

One of the attributes of a multicast request is the maximum time-to-live (TTL), which is the number of times a request is passed between subnets. After the TTL expires, the packet is discarded.

Note: You can use multicast discovery to discover systems across multiple subnets without configuring specific network information for each subnet. However, some networks are configured to prevent multicast requests from passing between subnets. As a best practice, contact your network administrator to determine the best discovery method for your organization.

Broadcast relay discovery:

In broadcast relay discovery, the management server sends a discovery request message to a specific Level-2 managed system, instructing the managed system to perform a discovery on the local subnet using a general broadcast. To perform the broadcast relay discovery, the system that performs the general broadcast must have already been discovered by IBM Director.

Broadcast relay discovery is used to identify Level-2 managed systems.

This discovery method is useful when the management server and the managed systems belong to different subnets and the network is configured to filter broadcast requests across those subnets.

Note: Typically, the Level-2 managed system that performs the general broadcast discovery is on a different subnet from the management server; however, it is not required.

When managed systems on the same subnet as the Level-2 managed system receive the discovery request, they reply directly to the management server that made the original request. This type of discovery generates less network traffic than a unicast discovery and avoids many of the problems associated with broadcast discovery when the network is configured to filter or prevent broadcasts. You might want to consider broadcast relay discovery if you have multiple physical locations in which managed systems reside, with lower-bandwidth network infrastructure (such as T1 or frame relay) between these physical sites.

Unicast discovery:

In unicast discovery, the management server sends requests directly to an exact IP address or range of IP addresses. Each address that you specify is contacted individually.

Unicast discovery is used to discover Level-0, Level-1, and Level-2 managed systems.

You can use this discovery method if your network filters both broadcast and multicast requests.

The disadvantage of a unicast discovery is that an IP packet must be sent for each individual IP address, which increases network traffic.

Service Location Protocol discovery:

In Service Location Protocol (SLP) discovery, the management server sends a request message for the IBM Director Agent SLP service type. An SLP Service Agent that replies to the request is identified in IBM Director Console as a Level-1 managed system.

SLP is an open-source Internet-standards track protocol that allows network applications, such as IBM Director, to discover the location and configuration of network services in a network. In an SLP implementation, an agent is a software entity that processes SLP protocol messages. There are three types of SLP agents:

User Agent (UA)
The SLP User Agent is a software entity that is looking for the location of one or more services. In an IBM Director environment, IBM Director Server acts as the user agent when it performs an SLP discovery.

Service Agent (SA)
The SLP Service Agent is a software entity that advertises the location of one or more services. In an IBM Director environment, Level 1: Core Services acts as the service agent. These Level-1 managed systems can advertise through the use of multicast messages and unicast responses to queries.

Directory Agent (DA)
The SLP Directory Agent is a software entity that acts as a centralized

repository for service location information. If your network administrator has configured a directory agent, you can configure IBM Director to use the directory agent to discover service agents.

When you install Level 1: IBM Director Core Services on a system, SLP and the IBM Director Agent SLP service type (management.software.IBM:director-agent) are installed on that system. Common Information Model (CIM) also is installed with Level 1: IBM Director Core Services.

SLP discovery is used to identify only Level-1 managed systems.

You can configure IBM Director to send an SLP discovery request as a unicast, multicast, or broadcast message. Some older versions of service agents do not support multicasting and might have to be discovered by using a broadcast.

Event management

An *event* is an occurrence of significance to a task, system, or managed object, such as the completion or failure of an operation. In a system-management environment, IBM Director Server receives events, traps, and notifications from many sources.

These sources include, but are not limited to, the following programs and protocols:

- IBM Director native events generated by IBM Director Agent
- CIM indications from the CIMOM that is installed as part of IBM Director Agent and IBM Director Core Services
- Microsoft Windows event log
- Windows Management Instrumentation (WMI)
- SNMP traps through out-of-band communication
- Platform Event Traps (PET) through out-of-band communication from Alert Standard Format (ASF)-capable systems and Intelligent Platform Management Interface (IPMI)- capable systems
- IBM service processors notifications through out-of-band communication

When IBM Director Server receives these events or notifications, it converts them into IBM Director events. For example, when IBM Director Server receives a CIM indication, it converts the CIM indication into an IBM Director event of the type CIM. When you view the Event Filter Builder tree, the CIM events are displayed under the CIM node in the tree.

Note: IBM Director can convert CIM indications into other event types, including event types that are used by enterprise-level system-management programs, such as SNMP events. Using these event types, IBM Director can provide system data to the by enterprise-level system-management programs through the IBM Director Upward Integration Modules. For more information, see the ″CIM indications in IBM Director″ section of the *IBM Director Events Reference*.

However, these SNMP events are not the same as SNMP traps that IBM Director Server receives out-of-band (that is, not through IBM Director Agent or IBM Director Core Services). Out-of-band SNMP traps are generated by hardware products and other software programs. They are displayed under the SNMP node in the Event Filter Builder tree, but beneath a different subnode.

You can use the events in the Event Filter Builder tree when working with managed objects. To monitor one or more events, you must create an event filter that contains an event type from one of these sources, use the event filter as part of an event-action plan, and then apply the event-action plan to a managed object. Events from the Windows event log are displayed in the Windows event log tree in the Event Type Filter Builder. Events from WMI are displayed in the Common Information Model (CIM) tree.

Alerts and resolutions

In IBM Director, an event can be in one of the following categories: alert and resolution. Typically, an *alert* is the occurrence of a problem relating to a managed object. A *resolution* is the occurrence of a correction or solution to a problem.

Note: In the IBM Director product, there are tasks and features that use the word *alert* in place of the word event. Also, ServeRAID Manager uses the word *notification* instead of event.

Monitoring operating-system specific events

If you want to monitor Windows- or i5/OS-specific events in the IBM Director environment, you must create an event-action plan in order for IBM Director to process these events. The predefined active event-action plan in IBM Director, Log All Events, does not monitor these operating-system specific events.

Managed objects running Windows or i5/OS can generate the following operating-specific events:

Window-specific event types	• Windows event log • (Optional) A subset of the following CIM events: – Windows event log – Windows services – Windows registry
i5/OS specific event types	• Msgq

Even though these events are generated by their respective operating systems (or an optional layer that is installed on the operating system), IBM Director does not process these events unless you create an event-action plan to do so. When you install IBM Director, it has one predefined active event-action plan: Log All Events. However, this event-action plan does not log these Windows- or i5/OS-specific events. You must create an event-action plan with a simple-event filter that contains the event types for one or more of these events. Then, you must apply this event-action plan to the managed object running Windows or i5/OS.

When IBM Director Agent starts on a managed object running Windows, the twgescli.exe program starts, too. This program listens for IBM Director Server to send a message to IBM Director Agent that an event-action plan has been applied to that managed object. If the event-action plan includes a simple-event filter that contains the event types for any of the Windows-specific events, IBM Director appropriates these events for its own use. This is called *event subscription*. The twgescli.exe program subscribes to the event types that are specified in the event-action plan and translates the Windows-specific events into an IBM Director event type. Then, the program forwards the events to the management server from which the event-action plan was applied.

When IBM Director Agent starts on a managed object running i5/OS, the process is the same with comparable code to twgescli.exe that is included in IBM Director Agent for i5/OS.

Processing an event in IBM Director

Understanding how IBM Director processes an event can help you build and troubleshoot event-action plans.

IBM Director completes the following steps to determine which event actions to execute:

1. The managed object generates an event and forwards the event to all the management servers that have discovered the managed object (except for some events, such as those that are generated through meeting or exceeding a resource-monitor threshold, which are sent only to the management server where the thresholds are configured and applied).
2. IBM Director Server processes the event and determines which managed object generated the event and which group or groups the managed object belongs to.
3. IBM Director Server determines whether any event-action plans are applied to the managed object or to any of the groups of which the managed object is a member.
4. If an event-action plan has been applied, IBM Director Server determines whether any event filters match the event that was generated.
5. The management server performs any event actions for each matching event filter.

Event-action plans

When you create an event-action plan, you attach one or more event actions to a specified event filter. Then, you include one or more event filters in the event-action plan. Finally, you apply that event-action plan to a system or group of systems.

An event-action plan is composed of two types of components:

- Event filters, which specify event types and any related parameters
- Event actions, which occur in response to filtered events

You can apply an event-action plan to an individual managed object, several managed objects, or a group of managed objects.

By creating event-action plans and applying them to specific managed objects, you can be notified by e-mail or pager, for example, when a specified threshold is reached or a specified event occurs. Or you can configure an event-action plan to start a program on a managed object and change a managed-object variable when a specific event occurs. You can use process-monitor events and resource-monitor events to build an event-action plan.

Successful implementation of event-action plans requires planning and consideration of how you will implement them. In particular, developing and following strict naming conventions is important, so that you can easily identify what a specific plan does.

Modifying an existing event-action plan

You can modify an existing event-action plan, even one that is already applied to managed objects or groups, using the Event Action Plan Builder.

If you modify an event filter or an event action that is used in an existing event-action plan, the changes are applied automatically to any event-action plans that use those filters or actions. If you add or delete a filter or an action that is used in an existing event-action plan, the following warning is displayed.

Event filters

An *event filter* specifies an instance of one or more events that you want IBM Director to process. IBM Director ignores any event instances that do not meet the specifications of the event filter. Because the event filter can specify possible values for the extended attributes that are included in an event type's definition, an event instance can be customized for very specific problems and occurrences. To permit users to quickly event filters, the extended attributes include default values; however, users can customize the extended attribute settings.

You can use an event filter to capture a single event or multiple events. The following list includes some of the criteria that you can use to determine whether to include an event with other events:

- All managed objects that are targeted for the filter are able to generate all events that are included in the filter. If the managed object does not generate the event for which the filter is defined, the filter will not be effective on that managed object.
- The event actions that will be used to respond to the event are the same for all targeted objects.
- The other event filter options besides the event type are common for all targeted objects. These settings include the times the event filter is active, the severity of the event, and other attributes.

Event-action plans can include event filters with event types that are not generated by all managed objects. In such instances, you can apply the event-action plan to those managed objects, but it will have no effect. For example, if an event filter is based on a ServeRAID event and that event-action plan is applied to managed objects that do not have a ServeRAID adapter installed, the event filter has no events to filter, and therefore, no actions are performed. If you understand this concept, you can create more complex event-action plans, and you can reduce the number of event-action plans you have to build and maintain.

All currently available event types are displayed in the tree on the Event Type page in the Event Filter Builder window. The currently installed tasks and extensions publish their events in the Event Type tree when IBM Director Server or IBM Director Agent or IBM Director Core Services starts.

Note: Whether the events are published when IBM Director Server or IBM Director Agent or IBM Director Core Services starts depends on the tasks or extensions and how they are implemented.

If you add an extension to your IBM Director installation, the extension might publish its events either when it is added to the installation or when the extension sends its first event. If the extension publishes when it sends its first event, only that event is published.

Event-filter types:

IBM Director provides four types of event filters.

In the Event Action Plan Builder window, the Event Filters pane provides the following event filters.

Event filter	Description
Simple Event	Simple event filters are general-purpose filters; most event filters are this type. When you expand this tree, any customized simple event filters that you have created are displayed. Also, the following predefined, read-only event filters are displayed: • All Events • Critical Events • Environmental Sensor Events • Fatal Events • Hardware Predictive Failure Events • Harmless Events • Minor Events • Security Events • Storage Events • Unknown Events • Warning Events Some of these predefined filters use the severity of events to determine which events they will allow to pass through; other filters target a specific type of event. For example, the Critical Events filter processes only those events that have a Critical severity. The All Events filter processes any events that occur on any managed object, except for Windows-specific and i5/OS-specific events. Using one of these preconfigured event filters ensures that the correct event type or event severity is preselected. If you want to see what events are included in a predefined event filter, double-click that predefined event filter in the Event Filters pane. The “Simple Event Filter Builder” window opens, and the Event Filter Builder notebook is displayed. Select the applicable notebook page to view the selected event filters. For example, click the **Severity** tab to view the selections for the Critical Event filter. You cannot change predefined event filters; they are read-only. However, you can make changes and click **File → Save As** to save the modified event filter with another name.
Duplication Event	Duplication event filters ignore duplicate events, in addition to the options that are available in the simple event filters. To use this filter, you must specify the number of times (Count) that the same event is ignored during a specified time range (Interval). Then, this filter processes the first event that meets the criteria that are defined for this filter. Only the first event triggers the event actions that are associated with this event filter. For the associated event actions to be triggered again, one of the following conditions must be met: • The value that is specified in the **Count** field must be exceeded. • The time range that is specified in the **Interval** field must elapse. • The value that is specified in the **Count** field must be exceeded by 1 (Count+1) within the time range that is specified in the **Interval** field. For example, you can define a duplication event filter to filter on the occurrence of an offline event and define a corresponding event action to forward the event to IBM Director Server. Depending on the criteria that you define, only the first event announcing that the system is offline is processed, and all other instances in which an event meets the filtering criteria are discarded until the Count value is exceeded during the specified interval.

Event filter	Description
Exclusion Event	Exclusion event filters exclude certain event types, in addition to the simple event filter options. Using this filter, you define the criteria of the events to exclude.
Threshold Event	A threshold event filter processes an event after it has occurred a specified number of times within a specified interval, in addition to the simple event filter options. An event that meets the criteria that are defined in this filter triggers associated actions only after an event has met the criteria for the number of times that are specified in the **Count** field or only after the number of times specified in the **Count** field within the time range specified in the **Interval** field. For example, you can define a threshold event filter to monitor frequently occurring heartbeat events and forward the event to IBM Director Server only when the heartbeat event is received for the 100th time during a specified amount of time.

Event-filter criteria:

Depending on the event-filter type, you set specific values for these types of criteria.

Criteria	Description
Event Type	Use the Event Type page to specify the source or sources of the events that are to be processed. This tree is created dynamically; and entries are added by tasks and as new alerts are received. Entries in the tree can be expanded to display suboption events. Most event filters are created using only this page. It specifies the source or sources of the events that are to be processed by this filter. By default, the **Any** check box is selected, meaning that none of the events that are listed are filtered, except for Windows-specific and i5/OS-specific events. If you want to specify certain events on which to filter, clear the **Any** check box. You can highlight more than one event by pressing the Ctrl or Shift key. **Notes:** 1. When you select a root option in the Event Type tree, all suboption events are selected as well. For example, when you select **MPA** in the Simple Event Filter Builder window, all Component, Deployment, Environmental, and Platform suboption events are selected also. If additional event types are published after you create the event filter, the newly available event types are included in your event filter only if the new event types are suboption events of an event type that you selected. However, if you want to include a newly published event type that is not a suboption event, you must update the event filter by selecting the new event type. 2. The event types for BladeCenter events are displayed under **MPA**, except for BladeCenter Configuration Management events, which are displayed under **Configuration Management**.

Criteria	Description
Severity	Use the Severity page to indicate the urgency of the events that are filtered. If an event is received whose severity level is not included in the event filter, the filter will not process that event. By default, the **Any** check box is selected, indicating that all event severities are processed by the filter. When you select more than one severity, they are joined together using logical OR. The source of the event determines what severity the event is. Generally, the severity levels have the following meanings: **Fatal** The event caused a failure and must be resolved before the program or component is restarted. **Critical** The event might cause a failure and must be resolved immediately. **Minor** The event is not likely to cause immediate program failure but should be resolved. **Warning** The event is not necessarily problematic but might warrant investigation. **Harmless** The event is for information only. Most events of this severity do not indicate potential problems. However, offline events are categorized as harmless, and these events *can* indicate potential problems. **Unknown** The application that generated the event did not assign a severity level.
Day/Time	Use the Day/Time page to set the filter to accept and ignore events on certain days and at certain times of the day. By default, the **Any** check box is selected, indicating that events that occur at any time are processed by the event filter. The time zone that applies to the specified time is the time zone in which the management server is located. If your management console is not in the same time zone as the management server, the difference in time zones is displayed above the Selections pane as an aid to determining the correct time. By default, all events are passed through all filters. This includes events that were queued by IBM Director Agent because the link between the managed object and the management server was unavailable. However, you can prevent these queued events from being processed by a filter by selecting the **Block queued events** check box. This option can be useful if the timing of the event is important or if you want to avoid filtering on multiple queued events that are sent all at once when IBM Director Server becomes accessible. However, you can block queued events only if you filter events at a specified time. To block queued events, you must clear the **Any** check box.
Category	Use the Category page to specify an event filter according to the status of an event (alert or resolution of a problem). However, not all events have resolutions.

Criteria	Description
Sender Name	Use the Sender Name page to specify the managed object to which the event filter will apply. Events that are generated by all other managed objects will be ignored. By default, the **Any** check box is selected, indicating that events from all managed objects (including IBM Director Server) are processed by the event filter. Initially, only IBM Director Server is shown in the list. As other managed objects generate events, such as when a threshold is exceeded, this list is added to dynamically. If you anticipate that other managed objects will generate events, you also can type managed-object names into the field and click **Add** to add them.
Extended Attributes	Use the Extended Attributes page to specify additional event-filter criteria using additional keywords and keyword values that you can associate with some categories of events, such as SNMP. This page is available only when you clear the **Any** check box on the Event Type page and select certain entries from that page. If the Extended Attributes page is available for a specific event type but no keywords are listed, IBM Director Server is not aware of any keywords that can be used for filtering. To view the extended attributes of specific event types, expand the **Event Log** task in the IBM Director Console Tasks pane and select an event of that type from the list. The extended attributes of the event, if any, are displayed at the bottom of the Event Details pane, below the Sender Name category.
System Variables	Use the System Variables page to further qualify the filtering criteria by specifying a system variable. This page is available only if there are one or more system variables. A system variable consists of a user-defined pairing of a keyword and value that are known only to the local management server. **Note:** These user-defined system variables are not associated with the system variables of the Windows operating system.
Event Text	Use the Event Text page to specify event message text to associate with the event.

Event actions

The *event action* specifies the actions that you want IBM Director to take as a result of the occurrence of an event.

Event action types

IBM Director has several predefined event action types. With the exception of Add to Event Log, you must customize each event action type that you want to use.

Add/Remove “event” system to Static Group
Adds a managed object to or removes a managed object from a specified static group when the managed object logs a specific event.

Add/Remove source group members to target static group
Adds all specified managed objects in a source group to a target group or removes all specified managed objects from the target group.

Add a Message to the Console Ticker Tape
Displays a message in red type that scrolls from right to left at the bottom of IBM Director Console.

Add to the Event Log
Adds a description of the event to the IBM Director event log.

Define a Timed Alarm to Generate an Event
Generates an event only if IBM Director does not receive an associated event within the specified interval.

Define a Timed Alarm to Start a Program on the Server
Starts a program on the management server if IBM Director does not receive an associated event within the specified interval.

Log to Textual Log File
Generates a text log file for the event that triggers this action.

Post a News Group (NNTP)
Sends a message to a newsgroup using the Network News Transfer Protocol (NNTP).

Resend Modified Event
Creates or changes an event action that modifies and resends an original event.

Send an Alphanumeric Page (via TAP)
Windows only) Sends a message to a pager using the Telocator Alphanumeric Protocol (TAP).

Send an Event Message to a Console User
Displays a popup message on the management console of one or more specified users.

Send an Internet (SMTP) E-mail
Sends a Simple Mail Transfer Protocol (SMTP) e-mail message.

Send an SNMP Inform to an IP host
Sends an SNMP inform request to a specified IP host.

Send an SNMP Trap to a NetView Host
Generates an SNMP trap and sends it to a specified NetView host using a TCP/IP connection to the host. If delivery of the SNMP trap fails, a message is posted in the history log of the managed object.

Send an SNMP Trap to an IP Host
Generates an SNMPv1 or SNMPv2c trap and sends it to a specified IP address or host name.

Send a Numeric Page
(Windows only) Sends a numeric-only message to the specified pager.

Send a TEC Event to a TEC Server
Generates a Tivoli Enterprise Console event and sends it to a specified Tivoli Enterprise Console server.

Set an Event System Variable
Sets the managed system variable to a new value or resets the value of an existing system variable.

Start a Program on a System
Starts a program on any managed objects on which IBM Director Agent is installed.

Start a Program on the "event" System
Starts a program on the managed object that generated the event.

Start a Program on the Server
In response to an event, starts a program on the management server that received the event.

Start a Task on the "event" System
In response to an event, starts a noninteractive task on the managed object that generated the event.

Update the Status of the "event" System
When the selected resource status generates an event, causes a status indicator beside the icon of the managed object that is associated with the resource to be set or cleared according to your specification.

Event-data-substitution variables

For some event-action types, you can include event-specific information as part of the text message. Including event information is referred to as *event-data substitution*. You can use these event-data-substitution variables to customize event actions.

&date The date the event occurred.

&time The time the event occurred.

&text The event details, if they are supplied by the event.

&type The event-type criteria that are used to trigger the event. For example, the event that is generated when a managed object goes offline is of type Director > Topology > Offline. This corresponds to the entry on the Event Type page.

&severity
The severity level of the event.

&system
The name of the managed object for which the event was generated. The system name is either the name of IBM Director Agent or, in the case of an SNMP device, the TCP/IP address.

&sender
The name of the managed object from which the event was sent. This variable returns null if the name is unavailable.

&group
The group to which the target object belongs and is being monitored. This variable returns null if the group is unavailable.

&category
The category of the event, either Alert or Resolution. For example, if the managed object goes offline, the category is Alert. If the managed object goes online, the category is Resolution.

&pgmtype
A dotted representation of the event type using internal type strings.

&timestamp
The coordinated time of the event.

&rawsev
The nonlocalized string of event severity (Fatal, Critical, Minor, Warning, Harmless, Unknown).

&rawcat
The nonlocalized string of event category (Alert, Resolution).

&corr The correlator string of the event. Related events, such as those from the same monitor-threshold activation, will match this.

&snduid
The unique ID of the event sender.

&sysuid
The unique ID of the managed object that is associated with the event.

&prop:*file_name***#***property_name*
The value of the property string *property_name* from property file *file_name* (relative to the IBM\Director\classes directory).

Note: For i5/OS, the absolute directory path must be used.

&sysvar:*variable_name*
The event system variable *variable_name*. This variable returns null if a value is unavailable.

&slotid:*slot_id*
The value of the event detail slot with the nonlocalized ID *slot_id*.

&md5hash
The MD5 (message digest 5) hash code, or cyclic redundancy check (CRC), of the event data (an event-specific unique ID).

&hashtxt
Provides a full replacement for the field with an MD5 hash code (32-character hex code) of the event text.

&hashtxt16
Provides a full replacement for the field with a short MD5 hash code (16-character hex code) of the event text.

&otherstring
The value of the detail slot that has a localized label that matches *otherstring*. A detail slot is a record in an event detail. For example, an event has one event detail that has an ID of *key1* and a value of *value1*. You can use the substitution variable &soltid:*key1* to obtain the value *value1*. You also can use &key1 to obtain the value *value1*. In the description above, *otherstring* is a placeholder for the user-defined event detail ID. However, if the passed ID is not found, “Not applicable” is returned.

Message Browser

You can use the Message Browser to view events that are sent to IBM Director Console. The Message Browser is displayed automatically whenever an alert is sent to the management console.

You can chose to have events sent to the management console when an event occurs by configuring an event action plan with the Send an Event Message to a Console User event action.

The Message Browser displays all alerts, including management console ticker-tape alerts. However, the Message Browser does not display any ticker-tape messages. A ticker-tape message can display, for example, resource-monitor data.

Groups

IBM Director allows you to organize logical sets of managed objects into groups. For example, a group might contain all managed systems that have Linux installed.

When you start IBM Director Console for the first time, the default groups are displayed. This includes the All Systems and Devices group, which contains all discovered managed objects and devices.

There are two types of groups that you can create in IBM Director:

Static groups
Static groups contains a specified set of managed systems. IBM Director Server does not automatically update the contents of a static group. The members of a static group are fixed unless you change them thru the IBM Director Console or event action plan. You can also copy the members of any dynamic group to a static group.

Dynamic groups
Dynamic groups are based on specified inventory or task criteria. You can create a dynamic group by specifying criteria that the attributes and properties of the managed systems must match. IBM Director automatically adds or removes managed systems to or from the group when their attributes and properties change, affecting their match to the group criteria.

You also can create a dynamic group based on the types of tasks for which the group of managed systems is enabled. You can initiate a specific task on all members of the group with a single drag and drop operation without having to consider whether each managed object supports that task.

You can also create *group categories* to organize your groups.

When you select a group, the managed objects that are members of that group are displayed in the Group Contents pane.

You can perform a task on all managed objects in a specific group. To perform tasks simultaneously on multiple groups, create a new group and include managed systems that you want from the multiple groups, or combine several separate existing groups into one new group.

Tip: It is might be useful to consolidate all groups for which you have administrative authority into a single category. This enables you to focus on those managed systems that are your responsibility while removing other managed systems and devices from your immediate attention.

Group account

A *group account* is a collection of user accounts. On a management server that is running Windows, you can create a group account to manage the privileges of multiple accounts together. By making a user account a member of a group account, that user has the privileges and access to the tasks that are defined for that group.

Note: You cannot manage group accounts through IBM Director on management systems that are running Linux. The group management feature is currently supported only on Windows.

When IBM Director Server is installed, two groups of IBM Director users are created automatically at the operating system level: DirAdmin (diradmin in Linux) and DirSuper (dirsuper in Linux). To create a group account you must first create a group account for the operating system that is running on the management server. Once the account is created that group must be made a member of either the diradmin or dirsuper group. After a group account has been added to the applicable IBM Director group, you can log in to IBM Director Console as an administrator and configure that group's privileges to IBM Director tasks and groups. At the operating system level, you can add user accounts to the group accounts that you create. You can manage the privileges of all the user accounts in a group by configuring the group privileges. The changes that you make will affect all of the users who are members of the group.

Group membership

The group to which a user account belongs provides group membership. On a management server that is running Windows, you can create several user accounts and make them members of the same group. The privileges that are assigned to the group in IBM Director are also assigned to its members.

Groups that are used with scalable objects

IBM Director provides several default groups of scalable objects in the Groups pane for easier management of these objects.

The default groups that are relevant to scalable objects are shown in table below.

Table 3. IBM Director groups that are used with scalable objects

Group name	Managed objects
Logical Platforms	All logical-platform objects, which includes all scalable partitions.
Physical Platforms	All physical-platform objects, which includes all scalable nodes.
Platforms	All logical platforms and physical platforms.
Platforms and Platform Members	All logical and physical platforms and any managed systems that result from these platforms.
Scalable Partitions	Only scalable partitions.
Scalable Systems	Only scalable systems.
Scalable Systems and Members	All scalable systems and all members of those scalable systems. Members of a scalable system include its scalable partitions, its scalable nodes, and any remote I/O enclosures attached to its scalable nodes. This group also includes managed systems that result from its scalable partitions.

Groups used with storage managed objects

IBM Director provides several default groups of storage managed objects in the Groups pane for easier management of these objects.

This table lists the groups that support storage managed objects.

Group name	Storage managed objects
SMIS-Storage Devices	Only storage managed objects that comply with the SMI-S standard.
Storage Devices	All storage managed objects, regardless of compliance with SMI-S standards.

Managed objects

This topic describes the concepts of managed objects in IBM Director.

A managed object is a system or device that is managed by IBM Director. IBM Director manages these types of objects:

managed system

Any computer, such as a server, desktop, workstation, or mobile computer, that can be managed by IBM Director. In this release, managed systems are subcategorized as follows:

Level-0 (″agentless″) managed systems

Systems that are managed through the network services that are native to the operating system: Secure Shell (SSH) or Windows Management Instrumentation (WMI). No IBM Director software is installed.

Level-1 managed systems

Systems that are managed through installation of IBM Director Core Services, which provides a subset of IBM Director Agent functionality, including Remote Session, Power Control, Hardware Status, Event Log, hardware-only inventory data, and distribution of system level updates.

Level-2 managed systems

Systems that are managed through installation of IBM Director Agent, which provides added functionality for administering the system. The functionality of IBM Director Agent on the managed system will vary depending on the operating system and platform.

managed device

An SNMP device (such as a network device, printer, desktop computer, or server) that has an SNMP agent installed or embedded.

physical platform

A single physical chassis or server that has been discovered through the use of the Service Location Protocol (SLP). A physical platform also can be created when:

- A deployable system is discovered through an RDM scan.
- You right-click any blank space in the Group Contents pane to create the physical platform manually.
- IBM Director Server determines that a physical platform does not exist already for a blade server in a BladeCenter unit.
- IBM Director Server first discovers and gains access to a Level-1 or Level-2 managed system.
- IBM Director Server gains Internet Protocol (IP) access to a Remote Supervisor Adapter service processor. It will query the Remote Supervisor Adapter or Remote Supervisor Adapter II service processor for the topology of its associated ASM interconnect network, and for each ISMP system found, a physical platform is created.

A physical platform can identify some managed systems before an operating system or IBM Director Agent is installed.

Note: To delete a physical platform from IBM Director Console, you also must delete any associated managed system or systems.

Scalable objects

Systems with multinode configurations, including scalable systems, nodes, and partitions.

BladeCenter chassis

Windows cluster

Racks Racks that are created by Rack Manager.

Static partitions

Managed system

A managed system is a computer, such as a server, desktop, workstation, or mobile computer, that can be managed by IBM Director. There are three levels that are used to categorize managed systems.

Level 0: Agentless Systems

Level 0: Agentless Systems do not have any IBM Director software installed on them; however, although they can still be managed by using IBM Director. Level-0 managed systems can be IBM or non-IBM servers, desktop computers, workstations, and mobile computers.

A Level-0 managed system is a managed system that does not have IBM Director Agent or the Common Information Model (CIM) instrumentation installed, but does have the minimum set of protocols that are required for that system to be managed by IBM Director. To manage a system using IBM Director, that system, at a minimum, must support the Secure Shell (SSH) or Distributed Component Object Model (DCOM) protocol.

Level 0: Agentless Systems is supported on systems that are running the following operating systems. For a detailed list of supported operating-system versions, see *IBM Director Installation and Configuration Guide*.

- Linux
- Windows

IBM Director discovers Level-0 managed systems by verifying the IP addresses on your network and scanning the ports of those addresses using the SSH or DCOM protocols. The range of IP addresses that are verified is governed by the IBM Director discovery preferences that you configure in IBM Director Console. By default, IBM Director uses the range of addresses that are in the IP domain of the management server.

When a Level-0 managed system is discovered, it is locked by default. You can unlock the system by requesting access to it through IBM Director Console. If the object that is discovered supports SSH but is not a computer system (for example, a Remote Supervisor Adapter (RSA)), the object will be displayed in IBM Director Console but will not support any tasks. Systems and other network devices that support the SNMP protocol will display as SNMP Device Managed Objects in IBM Director Console.

The attributes that are returned for a Level-0 managed system include:

- Locked
  - System Name
  - TCP/IP Addresses
  - System State
  - MAC Address
  - System Presence Check Setting
- Unlocked
  - Computer Name
  - Architecture
  - OS Major Version

- – OS Minor Version
- – Access Denied
- – Operating System
- – Unique System ID
- – System UUID (xSeries)
- – Machine Type
- – Model Number
- – Serial Number

After you discover and unlock a Level-0 managed system, you can perform the following tasks on that system:

- Collect inventory that is available from the operating system.
- Install Level 1: Core Services or Level 2: IBM Director Agent by using software distribution.
- Restart the operating system (Linux only).
- Run command-line programs (only if SSH is present).

Level 1: IBM Director Core Services systems

Level 1: IBM Director Core Services provides managed systems with a subset of the Level 2: IBM Director Agent functionality that is used to communicate with and administer that system. Specifically, it provides hardware alerts and status information that can flow to Director or 3rd party management servers.

A *Level-1 managed system* is any system that has Level 1: IBM Director Core Services installed but does not have Level 2: IBM Director Agent installed. Level-1 managed systems can be IBM servers, desktop computers, workstations, and mobile computers.

You can perform these tasks on a Level-1 managed system:

- Collect inventory.
- Install Level 2: IBM Director Agent using software distribution.
- Manage events using event action plans, event subscription, and the event log.
- Monitor hardware status.
- Restart or shutdown the managed system.
- Run command-line programs.

Level 1: IBM Director Core Services is supported on systems that are running the following operating systems. For a detailed list of supported operating-system versions, refer to the *IBM Director Installation and Configuration Guide*see .

- Linux (xSeries, System p5 and pSeries, and System z9 and zSeries)
- Windows

Level 2: IBM Director Agent systems

Level 2: IBM Director Agent provides managed systems with the full complement of IBM Director Agent functionality that is used to communicate with and administer that system. The functionality of Level-2: IBM Director Agent on a managed system varies, depending on the operating system and platform.

A *Level-2 managed systems* is any system that has Level 2: IBM Director Agent installed. Level-2 managed systems can be IBM or non-IBM servers, desktop computers, workstations, and mobile computers.

Level 2: IBM Director Agent is supported on systems that are running the following operating systems. For a detailed list of supported operating-system versions, refer to the *IBM Director Installation and Configuration Guide*.
- AIX
- i5/OS
- Linux (xSeries, System p5 and pSeries, and System z9 and zSeries)
- NetWare
- Windows

Mass configuration profiles

Using mass-configuration profiles, you can to quickly configure a group of managed objects.

One of the advantages of IBM Director is its ability to make certain configuration changes on multiple managed systems at once. Even in a dynamic host control protocol (DHCP)-enabled environment, many critical servers tend to use static addresses. Using mass-configuration profiles, you can, for example, change the IP address that these managed systems use to locate their primary DNS server, without having to physically visit each system.

A *mass-configuration profile* identifies the configuration information that you want to distribute to a managed object or group. You can create a profile for these IBM Director tasks:
- Asset ID™
- Configure Alert Standard Format
- Configure SNMP Agent
- Network Configuration

The mass-configuration profiles are saved in the Tasks menu (and in the Tasks pane) under the task with which they are associated. After you create a mass-configuration profile, you then can apply that profile to a managed object or a group.

Restriction:

- You can configure only the community name and trap destination in a mass-configuration profile for the Configure SNMP Agent task. You cannot configure security information.
- Specifying a community name in the profile adds the community name as a trap destination only. For IBM Director to receive traps sent to that community, you must manually add the community name to the IBM Director Server.
- If you want to perform SNMP sets on managed systems, you must manually modify the Security pane in the SNMP service properties window, adding the correct Community Name as READ WRITE or READ CREATE enabled on each managed system.

Performance

IBM Director provides a robust tool for monitoring and managing the performance of your environment, called Capacity Manager.

Performance-analysis monitors

This concept describes the performance-analysis monitors.

Performance-analysis monitors are a subset of resource monitors that are considered critical and are used to make performance recommendations. The performance-analysis monitors are activated by default when you install Capacity Manager.

These are the types of performance-analysis monitors:

- Processor usage
- Memory usage
- Disk usage
- Network traffic

Note: You must turn on all four types of performance-analysis monitors for a report to display a performance-analysis recommendation.

Capacity Manager automatically discovers new disk and LAN resource monitors and removes monitors for devices that no longer exist. Performance-analysis monitors for Windows network adapters and physical disks are discovered when the Windows network adapters and physical disks are added to the managed system. If a checked network adapter or physical disk has been removed, Capacity Manager removes the corresponding performance-analysis monitor from the monitor list once every 24 hours or whenever the Capacity Manager Agent is restarted.

Note: Performance analysis is available for managed systems running a Windows or Linux operating system only.

Performance bottlenecks

This concept describes performance bottlenecks.

When you schedule Capacity Manager to check periodically for bottlenecks, or when you select to generate a report, the performance-analysis function looks for bottlenecks in managed-system hardware performance. When one or more performance-analysis monitors meet or exceed their preset threshold settings and you have selected the **Generate Bottleneck events** check box when you defined the report, a bottleneck event is generated. You can adjust the threshold settings on performance-analysis monitors, but you cannot change the default settings without impairing the performance-analysis function.

Corresponding to each types of performance-analysis monitors are the four types of bottlenecks:

- Processor
- Memory
- Disk
- LAN adapter

If a bottleneck is detected, two things happen:

- Each managed system with a bottleneck generates an event, and the event is displayed in the IBM Director event log.
- A report is generated and saved in the IBM\Director\reports directory (unless you specify another directory in the report definition).

When the performance-analysis function detects a bottleneck, it diagnoses the problem and determines a potential solution. The performance-analysis section of the report details the problem and recommendations.

Multiple bottlenecks can occur. For example, a disk bottleneck and a memory bottleneck can occur concurrently. In this case, the performance-analysis algorithm recognizes that insufficient memory can lead to disk thrashing, so the recommendation is to add more memory and leave the disk drives unchanged. Because systems and devices often interact in this way, each combination of bottlenecks (that is, microprocessor, memory, disk, and LAN adapter) constitutes a separate bottleneck with its own recommendation.

Often, when one bottleneck occurs, other bottlenecks are not evident because the first bottleneck slows the system. A *latent bottleneck* is one that is not evident even though the system has slowed down. Performance analysis reports a managed system or device as having a latent bottleneck if a performance monitor for that system or device exceeds the warning threshold at least 50% of the time that the performance monitor for another system or device is constrained.

Performance-analysis reports

This concept describes the performance-analysis reports.

When you create a report, you specify a report definition. A *report definition* identifies the details that you want to include in the report. You can create a customized report definition or use a predefine report definition. These predefined report definitions are included in Capacity Manager:

- Daily report to viewer
- Hourly bottleneck events to file
- Hourly report to viewer
- Monthly report to file
- Weekly report to file

You can generate a report for immediate viewing, or you can save the report to a file for later viewing.

Reports viewed from the Report Viewer window

The performance-analysis report consists of two sections:

Recommendations
Shows only the subset of details on which you have to act.

Details
Shows everything that was found and contains links so you can see a graph of the performance of the monitor in question.

The managed systems with the most severe bottlenecks appear at the top of the report list. A bottleneck that is reported in the Details section is displayed in the Recommendations section if it meets one of these criteria:

- It occurred on the last day of the report.
- It occurred more than 25% of the time, and it occurred more than any other bottleneck on that managed system.
- It has a high probability of occurring in the future. However, performance analysis must have enough data to make a reliable forecast.

Reports viewed from a saved HTML file

A report that is saved in HTML contains the following sections:

Table of Contents
Contains links to the other sections.

Report Table

Presents the same monitor and managed-system data that is also available in the Report Viewer in the Table view.

Report Information

Includes the file name, analysis start and end dates, days of the week and hours of coverage, name of the report definition, and a list of any managed systems that were requested but not included in the report.

Performance Analysis recommendations

Recommends remedies for the most serious bottlenecks.

Performance Analysis details

Includes information about the frequency and duration of both active and latent bottlenecks and their remedies.

Performance forecast

This concept describes how Capacity Manager predicts future performance.

Using the Forecast function of Capacity Manager, you can review a prediction of future performance of selected managed systems. Capacity Manager uses forecasting in these interfaces:

- In the performance-analysis section of a report. If there are no realized bottlenecks, Capacity Manager uses forecasting to predict, with a level of confidence, if and when it foresees a monitor performance bottleneck.
- In a managed-system monitor performance graph. On a graph of a selected monitor for one or more managed systems, you can click Forecast icon ( ) to see a forecast of the performance on the selected managed systems. The graph depicts both the observed data and the forecast.

To calculate future performance, Capacity Manager applies a wavelet transform to the monitor data before performing a least-squares linear regression . With this transformed data, it computes a forecast line with a 95% prediction interval. The forecast duration is equal to the duration of the observed data.

Tip: For the forecast to be valid, Capacity Manager must have a minimum of 24 days of previously collected data where the managed-system monitors have been running at least 50% of the time.

Performance forecast graph

The *forecast line* describes possible future data values that are consistent with the prediction that an actual future data value will fall within equal probability above or below the forecast line. This line is a dashed line with an arrow at the end.

The *forecast duration* is equal to your data-collection period. For example, if you have a month of collected data, the forecast will be for a month into the future.

The *prediction interval* is represented by the dotted lines above and below the forecast line. The prediction interval represents the range of data values that are located above and below the forecast line and are consistent with the prediction that an actual future data value will fall within the interval with a probability of 95%. The width of the interval depends on the variability of the observed monitor data: the greater the variability, the wider the prediction interval. The prediction interval is displayed when you request a forecast of a single managed system. Graphs of multiple managed-system forecasts do not show prediction intervals.

If you do not know how to interpret a wide prediction interval for a forecast, select a finer resolution of your data from the **Resolution** drop-down list located in the lower-right corner of the Graph pane. Your data points might have a broad variance that is hidden by averaging that occurs when data is displayed at a coarser resolution.

Notes:

1. The vertical bar at the beginning of the forecast data depicts the range.
2. The gap between the actual collected data and the beginning of the predicted data serves as a separator between these two data sets.

Scalable objects

Scalable objects are IBM Director managed objects that are used with multinode configurations of supported xSeries servers.

Scalable objects in IBM Director for xSeries 460 servers include:

- Scalable nodes
- Scalable partitions
- Scalable systems

IBM Director communicates out-of-band with service processors in xSeries 460 servers to manage hardware partitions. Each hardware partition can run a single image of the operating system and is defined as a *scalable partition* that consists of one or two xSeries 460 servers. The servers that are defined in a scalable partition have at least one SMP Expansion Module and are referred to as *scalable nodes*. A *scalable system* consists of scalable nodes and the scalable partitions that were created from those scalable nodes. These IBM Director managed objects are referred to as *scalable objects* throughout this documentation.

Note: IBM Director performs only discovery and power operations for scalable systems and scalable partitions that have been previously configured on xSeries 460 servers. It does not create or configure scalable systems or scalable partitions.

Scalable nodes

A scalable node is a server that has one or more SMP Expansion Modules. When IBM Director discovers such a server, it creates a physical-platform managed object. It also assigns attributes that record the number of SMP Expansion Modules, SMP Expansion Ports, and RXE Expansion Ports on the physical chassis.

In IBM Director Console, scalable nodes are identified with the same icon that is used for all physical platforms. To determine whether a physical platform has the additional attributes of a scalable node, in the Group Contents pane, double-click the icon for the physical platform. The Display System Attributes window opens and the attributes for SMP Expansion Modules and RXE Expansion Ports are in the list that is displayed.

The following requirements apply to multinode configurations:

- All servers in one scalable system must be of the same machine type and model, and must have the same number of SMP Expansion Modules.
- All servers in one scalable system must have the same type of service processor and the same firmware code level.
- The service processor of each server must be connected to an active network. This connection is necessary so that the service processors can communicate and

perform the necessary functions for the multiple servers to merge as one combined server or unmerge as separate servers. This connection also is required for out-of-band communication with IBM Director.

- All servers in one scalable system must be at the same basic input/output system (BIOS) code level.

Scalable partitions

A scalable partition contains one or more scalable nodes. Regardless of the number of scalable nodes it contains, a scalable partition can run a single image of an operating system.

Scalable partitions can:

- Be powered on and powered off
- Support an operating system
- Have a single, contiguous memory space and access to all associated adapters
- Identify the scalable nodes that are used by the scalable partition
- Be represented as managed systems after IBM Director Agent is installed on the scalable partition and the scalable partition is powered on.

You can view the state of a scalable partition from IBM Director Console. To do so, right-click the managed object for a scalable partition; then, click **Open** to display general attributes for that scalable partition. The scalable partition state is displayed under the general attribute State.

Furthermore, when you use the Status association in IBM Director Console, the Scalable Partition Power Status folder in the Group Contents pane includes several subcategories for scalable partition states.

Scalable systems

A scalable system is an IBM Director managed object that consists of scalable nodes and the scalable partitions that were created from the scalable nodes in the scalable system.

Attention: If you recable a multinode server into a different physical configuration after it has been used with IBM Director, you must notify IBM Director of the recabling changes by reestablishing out-of-band communication.

Groups that are used with scalable objects

IBM Director provides several default groups of scalable objects in the Groups pane for easier management of these objects.

The default groups that are relevant to scalable objects are shown in table below.

Table 4. IBM Director groups that are used with scalable objects

Group name	**Managed objects**
Logical Platforms	All logical-platform objects, which includes all scalable partitions.
Physical Platforms	All physical-platform objects, which includes all scalable nodes.
Platforms	All logical platforms and physical platforms.
Platforms and Platform Members	All logical and physical platforms and any managed systems that result from these platforms.

Table 4. IBM Director groups that are used with scalable objects (continued)

Group name	Managed objects
Scalable Partitions	Only scalable partitions.
Scalable Systems	Only scalable systems.
Scalable Systems and Members	All scalable systems and all members of those scalable systems. Members of a scalable system include its scalable partitions, its scalable nodes, and any remote I/O enclosures attached to its scalable nodes. This group also includes managed systems that result from its scalable partitions.

Power operations for scalable partitions

You can use IBM Director Console to power on and power off scalable partitions on xSeries 460 servers.

Power operation that are performed on managed objects that represent scalable partitions use out-of-band communication. Power operations that are performed on managed-system objects created from powered-on scalable partitions use in-band communication to power off the scalable partition.

Restriction: The out-of-band power operations in IBM Director 4.22 are only for use by xSeries 460 servers. Other supported servers (such as xSeries 455 and xSeries 445 servers) should install and use Scalable Systems Manager (SSM) 4.20 if needed.

IBM Director Console identifies all scalable partitions with the same scalable partition icon whether they are powered on or powered off. However, IBM Director Console uses additional icons with the scalable-partition icon to indicate the state of a scalable partition.

IBM Director Console uses the same icon to depict all physical platforms, including those that are not scalable nodes and those that are not in powered-on scalable partitions.

Discovering scalable objects

When a scalable node is unlocked, IBM Director performs additional discovery for the xSeries 460 server.

This discovery determines whether the NVRAM of the service processor contains a partition descriptor. If it does, IBM Director uses the partition-descriptor information to create scalable systems and scalable partitions. The partition descriptor in NVRAM was stored by the Web management interface for the xSeries 460 server.

IBM Director also creates the association between scalable systems and scalable nodes, and between scalable partitions and scalable nodes. The partition descriptor in NVRAM indicates how many scalable nodes are in a scalable system and how many scalable nodes are in a scalable partition.

The interrogation of NVRAM to locate a partition descriptor is performed in the background, in a manner similar to the discovery of physical platforms.

The following conventions are used to name the new scalable objects:

- The scalable system is named "Scalable System *xxxx*" where *xxxx* is the last four characters of the scalable system UUID that is read from NVRAM.
- The scalable partition is named "Scalable Partition *xxxx yyyy*" where *xxxx* is the last four characters of the scalable system UUID that is read from NVRAM and *yyyy* is the last four characters of the scalable partition UUID that is read from NVRAM.

Note: When more than one scalable system UUID ends with the same last four characters, this naming convention will result in duplicate names. For this reason, consider renaming automatically created scalable systems to avoid confusion.

When IBM Director Server discovers that IBM Director Agent is running on the newly started scalable partition, it creates a managed-system object to represent the active scalable partition. You can use IBM Director to manage this managed system as you would any other managed system. For example, by using Management Processor Assistant (MPA), system administrators can configure, monitor, and manage the service processors in xSeries servers. Further, IBM Director associates the managed-system object with its scalable partition object. Use the Scalable Partitions Membership association in IBM Director Console to view a tree structure of scalable partitions and their associations with any managed systems.

Security

IBM Director offers several security features, including authentication and user-administration options that enable system administrators to specify user privileges, support for Secure Sockets Layer (SSL), and optional encryption of interprocess communication.

Authentication

This topic provides conceptual information about authentication.

Integrated into IBM Director is a security mechanism by which a managed system can authenticate any management server attempting to access it. Authentication enables Level-1 and Level-2 managed system to accept commands from only an IBM Director Server that is trusted (that is, authorized to manage it). Authentication protects Level-1 and Level-2 managed system from access by unauthorized management servers or rogue managed-system applications.

The IBM Director authentication process is based on two interlocking concepts:

- Digital-signature certification
- Security state of the managed system

Digital-signature certification:

This topic provides conceptual information about digital-signature certification.

IBM Director authentication is based on the Digital Signature Algorithm (DSA). DSA is the public-key algorithm specified by the Digital Signature Standard of the National Institute of Standards and Technology. It enables the holder of a public key to verify the signature for a digital document that has been signed by a holder of the corresponding private key. In an IBM Director environment, it works in the following way:

1. IBM Director Server attempts to access IBM Director Agent. IBM Director Server bids the public keys that correspond to the private keys it holds.

2. IBM Director Agent checks these keys. If it considers the keys to be trusted, IBM Director Agent replies with a challenge that consists of one of the trusted public keys and a random data block.
3. IBM Director Server generates a digital signature of the random data block using the private key that corresponds to the public key included in the challenge. IBM Director Server sends the signature back to IBM Director Agent.
4. IBM Director Agent uses the public key to verify that the signature is a valid signature for the random data block. If the signature is valid, IBM Director Agent grants access to IBM Director Server.

This digital-signature scheme has the following benefits:

- The public keys stored on the managed systems can be used only for verifying access.
- Using a random data block for signing makes replay attacks unusable.
- Generating a private key corresponding to a given public key is cryptographically improbable, requiring 2^{128} or more operations to accomplish.

For Level-1 managed systems, the digital-signature certificate expires after 365 days. You can configure notification and polling settings when the certificate is about to expire. You can also create an event action plan to notify you of an expiring certificate.

Security states of managed systems:

This topic provides information about the security state of a managed system.

A managed system is in either an unsecured or secured state. A managed system is *unsecured* when any management server can access it and perform functions on it. A managed system is *secured* when only an authorized (trusted) management server can access it.

The initial security state of IBM Director Agent depends on the underlying operating system.

Table 5. Initial security state of IBM Director Agent

Operating system	Security state
AIX	Secured by default during installation of IBM Director Agent.
i5/OS	Secured by default during installation of IBM Director Agent.
Linux	Secured by default during installation of IBM Director Agent.
NetWare	Unsecured by default. Must be secured manually or during discovery. See Securing managed systems for more information.
Windows	Can be secured during installation of IBM Director Agent.

If IBM Director Agent is not secured during installation of IBM Director Agent, you can secure the managed system manually or during discovery.

Note: The IBM Director Agent running on a management server is secured automatically. It has a trust relationship with only the IBM Director Server installed on the same server.

On managed systems running Windows, the security state is determined by the secin.ini file. If the secin.ini file is initialized as unsecured, any management server can access the managed system and establish a trust relationship with IBM Director Agent. IBM Director Server establishes a trust relationship by giving IBM Director Agent a copy of its public key.

When the managed system has been secured by a management server, only that management server, any management servers that had previously established a trust relationship, and any future management servers that successfully request access are able to access the managed system.

Where the security information is stored:

This topic provides information about where security information is stored.

The information needed for authentication is stored in files on both the management server and the managed systems.

The public keys are stored in dsa*xxxxx*.pub files, where *xxxxx* is a unique identifier. The private keys held by IBM Director Server are stored in dsa*xxxxx*.pvt files. For example, the dsa23ef4.pub file contains the public key corresponding to the private key stored in the dsa23ef4.pvt file.

On systems running Windows, the secured or unsecured state data is stored in the secin.ini file, which is generated when you first start IBM Director Server or IBM Director Agent. On management servers, this file is initialized as secured; on managed systems, it is initialized as either secured or unsecured, depending on which options were selected during the installation of IBM Director Agent.

By default, the files are located in the following directories.

Operating system	Directory
AIX	/opt/ibm/director/data
i5/OS	/QIBM/UserData/Director/data
Linux operating systems for AMD64 and 32-bit systems	/opt/ibm/director/data
Linux operating systems for Intel Itanium and IBM iSeries and pSeries	/opt/ibm/director/data
NetWare	*d*:\IBM\Director
Windows	*d*:\Program Files\IBM\Director\Data

where *d* is the drive letter of the hard disk on which IBM Director is installed and IBM Director is installed in the default location.

How the keys and sec.ini files work together:

This topic provides information about how the keys and the secin.ini files work together.

When you first start IBM Director Server, it randomly generates a matching set of public and private key files (dsa*.pub and dsa*.pvt files). The secin.ini file is generated and initialized as secure.

The initial security state of a managed system depends on the following factors:

- Which operating system it is running

- Which features were selected during the installation of IBM Director Agent

Managed systems running NetWare are set to the unsecured state automatically. For all other managed systems, the initial security state depends on which features are selected when IBM Director Agent is installed. If either encryption or agent/server security is selected, the managed system is set automatically to the secured state.

While a managed system is in the unsecured state, it accepts a public key from *every* management server that attempts to access it. Through this process, the managed system establishes trust relationships with those management servers.

If a management server secures that unsecured managed system, it gives that managed system a copy of its public key *and* its secin.ini file, which is initialized as secure. After this has occurred, the managed system no longer accepts any new public keys from management servers. However, the managed system continues to grant access to any management server whose public key is stored on the managed system.

Key information and management:

This topic provides information about public and private keys and how to manage them.

The public and private key files are binary files, but they contain textual data that indicates their origin. If a dsa*.pub or dsa*.pvt file is printed using the type command at a command prompt, the following data is displayed in the first line:

```
DSAKeytypeString
```

where:

- *Keytype* indicates the type of the key. “P” denotes private, and “p” denotes public.
- *String* is the name of the management server that generated the key file.

For example, `DSAPtest1` indicates a private key file generated by a management server named test1, and `DSAptest1` indicates the public key file generated by the same management server.

It is *very important* to back up and protect the dsa*.pvt files. If they are lost, you cannot regenerate these files.

User accounts

A user account is an account that is set up for an individual that defines that user for IBM Director. The information that is saved to the account is full name, user ID, password, privileges, group access, task access, pager number, and e-mail address. From IBM Director Console, you can manage user accounts.

IBM Director user accounts are based upon the underlying operating-system accounts. When IBM Director Server is installed, two groups of IBM Director users are created automatically at the operating-system level: administrators and super users. The two user groups have different levels of access to IBM Director:

Administrator group

Members of the administrator group have general access to IBM Director, although the privileges available to the administrator group or an individual user can be restricted by a super user.

Super-user group

Members of the super-user group can define the privileges available to the administrator group. Also, they can create and edit individual user accounts. The privileges available to members of the super-user group cannot be restricted.

The following table lists the operating-system specific names of the IBM Director user groups.

Table 6. IBM Director user groups

Operating system	Administrator group	Super-user group
i5/OS	QIBM_QDIR_ADMINISTRATOR	QIBM_QDIR_SUPER_ADM_PRIVILEGES
Linux	diradmin	dirsuper
Windows	DirAdmin	DirSuper

To create a user that has access to IBM Director Server, you must first create a user account for the operating system that is running on the management server. Once the account is created, that user must be made a member of either the diradmin or dirsuper group. Root users (users belonging to the root group) or members of the Administrator group on Windows are also able to access IBM Director. After a user account has been added to the applicable IBM Director group, you can log in to IBM Director Console as an administrator and configure that user's privileges to IBM Director tasks and groups. On Windows, the IBM Director service account is automatically assigned to the super-user group (DirSuper). In addition, all operating-systems accounts with administrator privileges on the management server automatically can access IBM Director Console. Users with such operating-system accounts can access the same IBM Director privileges as members of the DirAdmin group.

On i5/OS, the groups are not automatically populated. A user with security administrator authority must assign users to the appropriate groups.

On Linux, the groups are not automatically populated. A user with root privileges must assign users to the appropriate groups.

On a management server that is running Windows, you also can create additional groups at the operating-system level and add these groups to the diradmin or dirsuper groups. The subgroups that you add to the diradmin or dirsuper groups can be managed in IBM Director.

Encryption

This topic provides conceptual information about encryption.

IBM Director contains a security feature that encrypts all data in interprocess communications, except transport-layer datagrams used during discovery. This encryption feature provides automatic key management and enables you to select an encryption algorithm from the provided libraries:

- IBM Java Cryptography Extension (JCE)
- OpenSSL

JCE provides ciphers for all Java-based platforms, including i5/OS and Linux; OpenSSL provides ciphers for 32-bit Windows operating systems.

Encryption is disabled by default. To encrypt data transmitted between Level-1 and Level-2 managed systems and IBM Director Server, you must enable encryption on both IBM Director Server and Level-1 and Level-2 managed systems.

When you install IBM Director Server, you can select one of the following encryption algorithms:

- Advanced Encryption Standard (AES)
- Data Encryption Standard (DES)
- Triple DES

IBM Director Server automatically generates a key, based on the encryption algorithm selected. IBM Director Server stores the key in memory and presents it to IBM Director Core Services or IBM Director Agent each time IBM Director Core Services or IBM Director Agent is started, using the Diffie-Hellman key exchange. This makes it unnecessary for a key to be stored on each managed system.

The following table outlines how data is transmitted between IBM Director Server and Level-1 and Level-2 managed systems, depending on whether encryption is enabled.

Table 7. Encryption state and data transmitted between IBM Director Server and IBM Director Agent

	IBM Director Core Services or IBM Director Agent (encryption enabled)	**IBM Director Core Services or IBM Director Agent (encryption disabled)**
IBM Director Server (encryption enabled)	Encrypted	Unencrypted
IBM Director Server (encryption disabled)	No data transmission possible	Unencrypted

Important: If two management servers have discovered each other (and they each appear in each other's IBM Director Console as managed nodes), and one management server (server A) has encryption enabled, and the other management server (server B) either has encryption disabled or has encryption enabled now but had it disabled when it was discovered and the communication has not ended since the discovery, then unencrypted transmissions sent by server B to server A will continue until the previous communication is ended. This occurs because server A (in its role as a management server) is already communicating with server B (in its role as managed system) in plain text. You can delete these managed objects from each other's console to end the unencrypted communication, and if you run multiple management servers that can discover each other, you can enable encryption on both management servers before they are started or before they can discover each other. You can also use the **dircli lsmo** command to check for previous communication.

Notes:

- Encryption is not supported on managed systems running NetWare or systems running 64-bit versions of Windows.

- Neither out-of-band communications nor communication used by Internet tools, such as Telnet or File Transfer Protocol (FTP), are encrypted.
- Enabling encryption imposes a performance penalty. Encrypting data packets and exchanging encryption keys has an effect on the speed with which IBM Director completes management operations. When either the management server or the managed systems are restarted, keys are regenerated and exchanged. Consequently, an unsecured managed system might appear to be unmanageable for a period of time.

Software distribution

This topic compares two methods for distributing software.

IBM Director supports the following methods of software distribution:
- Streaming from the management server
- Using redirected distribution

Streaming from the management server

This topic describes advantages and disadvantages of distributing software by streaming from the management server.

Software-distribution packages are copied directly from the management server to the managed system.

This method of software distribution is resource-intensive. It can have a negative effect on the management server performance. In addition, a package distributed by this method requires that the target managed system have empty disk space twice the size of the package.

For Level-2 managed systems, streaming from the management server has one advantage, however. If a network connection is broken during the transmission, IBM Director attempts to resume the connection from the point at which the transmission was interrupted. If the streaming operation can be resumed, retransmission time is saved.

Because of the ability to resume distribution, you might prefer to stream a software package from the management server if you have an unreliable or slow network link.

Using redirected distribution

This topic describes advantages and disadvantages of using redirected distribution to distribute software.

Many software packages are tens or hundreds of megabytes in size. Distributing software of this size across a large network can cause bottlenecks in network data transmission. To avoid this problem, you can set up a universal naming convention (UNC) or FTP share on a network server. IBM Director Server streams software packages to the network share, where they are cached. From the share, they are either streamed to the managed systems or, in the case of software that uses the Microsoft Windows Installer or InstallShield as the installation utility, installed directly from the file-distribution server.

Redirected distribution greatly reduces the software-distribution traffic in your network. It uses fewer system resources on the management-server. In addition, if you install InstallShield or Microsoft Windows Installer (MSI) packages directly from the file-distribution server, redirected distribution requires less disk space on the managed systems.

Redirected distribution has one limitation: if a redirected distribution of a software package is interrupted (for example, if the network connection is lost), the installation must begin all over.

File-distribution server considerations

Consider the following issues when setting up file-distribution shares:

- In a Windows environment, the file-distribution server must either be a member of the same domain as the management server or have a trust relationship with that domain.
- The management server must have full read/write access to the share. The user ID and password that were used to install IBM Director Server must also be present on the file-distribution server. Otherwise, software distribution uses streaming from the management server.
- The share must allow read access to all managed systems that you want to access the share.
- If the file-distribution server is configured as an FTP server, you can choose to use FTP when transferring packages from the management server to the share. For managed systems running Windows, the home directory for the FTP login must be the same directory as the file-distribution server. For example, if c:\stuff\swd_share is mapped to \\server\swd_share, then c:\stuff\swd_share must be the home directory for the FTP user ID login used on the FTP file-distribution server configuration panel.
- You can enable null credentials to access the share so that you do not have to specify a user ID and password for each managed system or group that needs to access the share. To enable null credentials, you must issue the **twgshare** command. This alters a registry setting on the file-distribution server, which enables managed systems to use null credentials to access the share. To issue the **twgshare** command, complete the following steps:
  1. Copy the twgshare.exe file to the file-distribution server. This file is in the \IBM\director\bin\ directory.
  2. From a command prompt, type the following command:

     ```
     twgshare -a sharename
     ```

     where *sharename* is the name of the share on the file-distribution server.
- If you do not want to use null credentials (which are a security risk), you must set up an operating-system account on the file-distribution server. This account must have read access to the share. Enter the user ID and password for this account when you configure distribution preferences for managed systems. In addition, the account must exist on each managed system so that IBM Director Agent can specify this user to mount the share.

SNMP devices

IBM Director discovers SNMP devices in your network according to discovery parameters that you can specify. The process that is used to discover SNMP devices in your network uses lists of initial IP addresses, SNMPv1 and SNMPv2c community names, subnet masks, and SNMPv3 profiles.

IBM Director works with SNMPv1, SNMPv2c, and SNMPv3 for all communications and recognizes Management Information Bases (MIBs) in System Management Information (SMI) version 1 and version 2 formats.

SNMPv1 and SNMPv2c devices and agents use community names to control their access. A community name can be any case-sensitive text string. By default, the community name of an SNMP device is set to `public`. If specific SNMP devices in your network have unique community names to restrict access, you can specify the correct name to gain access to a device. SNMPv3 devices and agents use profiles to control their access.

The subnet mask enables you to further refine the scope of the discovery process, limiting the search to certain subnets in the network. The default subnet mask is set to the subnet of each corresponding IP address.

Using your lists of IP addresses, community names, and subnet masks, a series of SNMP GET statements are performed against port 161 of the IP address to determine whether the address is associated with a valid SNMP device. A valid SNMP device for IBM Director has the following accessible values: sysName, sysObjectID, sysLocation, sysContact, sysDescr, and sysUpTime. If the object is determined to be a valid SNMP device, another series of SNMP GET statements are sent to obtain information in the ipNetToMediaNetAddress table, where additional IP addresses can be used to discover even more SNMP devices. The search continues until no new addresses are located. Newly discovered or created SNMP-device managed-object names default to the value of sysName. If sysName has no value, the host name of the device is used. If no host name is assigned, the IP address is used.

All SNMP traps that are configured with IBM Director Server as the destination are forwarded as events to the event log. Therefore, you can view an SNMP trap using the event log on the SNMP managed device that originated the trap. If a trap is received that corresponds to an SNMP device that has not been discovered, IBM Director creates the device automatically, if you selected the **Auto-add unknown agents which contact server** check box on the SNMP Discovery page in the Discovery Preferences window.

The MIB file is used to translate raw SNMP dotted decimal notation into human-readable text. This is especially useful for SNMP devices for Level-0 managed devices, which do not have IBM Director Core Services or IBM Director Agent installed (such as network hubs, switches, printers, and USPs). MIBs that are placed in the data\snmp directory on the management server are compiled automatically. You can also compile MIBs manually from the SNMP Browser window.

Storage managed objects

IBM Director recognizes certain storage devices that comply with the Storage Management Initiative Specification (SMI-S). This is an industry standard developed by the Storage Networking Industry Association (SNIA). IBM Director supports SMI-S versions 1.1 and 1.0.2.

IBM Director provides support these storage devices:
- IBM System Storage DS300
- IBM System Storage DS400
- IBM System Storage DS4000

IBM Director communicates with the storage devices through their respective SMI-S providers. Their Service Location Protocol (SLPv2) component enables the devices to be discovered by IBM Director, which looks for SNIA-defined SLP service types.

IBM Director obtains information about storage devices through the SMI-S provider's Common Information Model Object Manager (CIMOM) component. Communication occurs using the Distributed Management Task Force (DMTF) standard for Web Based Enterprise Management (WBEM) as required by SNIA. The information is organized according to the DMTF standard for the Common Information Model (CIM) using the profiles defined by SNIA.

Note: For DS4000, the SANtricity SMI Provider from Engenio, which complies with SMI-S version 1.1, is required. You can obtain this provider from the Engenio web site at www.engenio.com. Note that SMI providers version 1.1 are embedded in DS300 and DS400. The providers are from Adaptec.

For information on standards, see these Web sites:

SMI-S www.snia.org/smi/about

SNIA www.snia.org

DMTF
www.dmtf.org

WBEM
www.dmtf.org/standards/wbem

CIM www.dmtf.org/standards/cim

SMI-S Attributes

SMI-S attributes can include four sets of data values: Default, CIMOM, Base and Extended.

Default attributes

SMI-S storage managed objects appear on the console immediately after being discovered through SLPv2 as a locked object labeled by the host name (or IP address if DNS lookup fails). When the storage managed object is locked, only the public data surfaced through SLPv2 is available (such as a locked object that represents the CIMOM required by the SMI-S standard). Default attributes that are common to IBM Director managed objects are displayed. These attributes include:

SMI-S Storage Device Name
Starts as the primary host name or IP address of the CIMOM. Once unlocked, the generated name for the device.

System Factory ID
The managed object factory that creates the objects—SMI-S Storage Devices.

System State
State of the object. Starts as Unknown because access is required to determine the actual online, error or offline status.

System Presence Check Setting
Polling interval in minutes for checking the status.

Secure/Unsecure supported
Specifies whether securing of the target is supported.

Access Denied
True for a locked object. This attribute goes away once the object is unlocked.

Encryption Enabled
Specifies whether encryption to the target is enabled.

CIMOM attributes

Until the storage managed object is unlocked, the actual device and its associated data are unavailable. The storage managed object is unlocked by obtaining access to represent the CIMOM and the CIM schemas minded for data. The SLP provided CIMOM attributes includes:

SMI-S Provider Service ID
The unique service id of the CIMOM. As specified in SMI-S, responses from different IP addresses with the same Service ID value are assumed to be the same CIMOM and are collapsed together.

Interop Namespace
The entry namespace in CIM containing the Server profile that has the CIMOM data and list of registered profiles.

Registered Profiles Supported
The list of profiles this CIMOM claims to support.

SLP Service Names
The full service names for this CIMOM

SMI-S Provider IP Addresses
The list of one or more IP addresses for the CIMOM.

SMI-S Provider IP Port Numbers
The IP ports on this CIMOM.

Base attributes

After access is granted to the storage managed object, the full CIM data can be obtained. This might result in the device representation going away successfully because the entity already exists, or the object access failing or a security problem. If access is successful, the object becomes a representation of an actual device. The name changes based on the rules specified in the Discovery Preferences panel for SMI-S Storage Devices. The following base attributes then appear:

Registered Profile
The profile this device represents along with the version (for example, SNIA:Array 1.0.2).

Namespace
The namespace with the schema details about this particular profile.

Storage Server Manufacturer ID
The manufacturer of this device.

Storage Server Type and Model
Type and model string representing this device.

Storage Server Serial Number
Serial number for the device.

Endpoint Unique Name
Unique identifier (primary FC port in fibre channel).

Storage Device Type
Specifies whether this is a fibre channel or iSCSI device.

Extended attributes

Any extended attributes that are specified by the attribute extension files are created. These may vary based on the storage managed object. For example, for an IBM DS4000 disk array, the following attributes are displayed:

Subscriptions
Indicates whether a managed object is registered to receive state or alert indications.

Nickname
A name for the device

Controller IP Addresses
IP addresses for the controllers on the target.

External applications for storage managed objects

You can use IBM Director to start external applications for targeted storage managed objects, including Storage Manager Client for IBM System Storage DS4000 storage systems series.

Important: You must install Storage Manager Client on each management console from which you intend to use it.

Before you can start Storage Manager Client from IBM Director Console on systems running a Windows operating system, you must set the following environment variables:

JAVA_FAStT
Defines the Java Runtime Environment (JRE) directory that is associated with Storage Manager Client. The default directory is c:\Program Files\Common Files\IBM_FAStT\jre\1.4.

STORAGE_MANAGER
Defines the working directory of Storage Manager Client. The default directory is c:\Program Files\IBM_FAStT\client.

Groups used with storage managed objects

IBM Director provides several default groups of storage managed objects in the Groups pane for easier management of these objects.

This table lists the groups that support storage managed objects.

Group name	Storage managed objects
SMIS-Storage Devices	Only storage managed objects that comply with the SMI-S standard.
Storage Devices	All storage managed objects, regardless of compliance with SMI-S standards.

Service processors

Hardware-based service processors, also known as management processors, work with hardware instrumentation and systems management software and are key to problem notification and resolution and allow you to remotely manage your system. In an IBM Director environment, service processors send alerts IBM

Director Server when error conditions occur in a specific managed system and are key in helping you effectively managed your environment.

Communication with IBM Director Server

This topic provides information about the pathways on which data is transmitted between service processors and IBM Director Server.

There are several pathways along which communication between IBM Director Server and the service processors present in IBM Netfinity® or xSeries servers takes place:

In-band communication
IBM Director Server communicates with IBM Director Agent; IBM Director Agent uses a device driver to pass data to and from the service processor. This also is called interprocess communication (IPC).

Over the local area network (LAN)
Data is transmitted between the service processor and IBM Director Server over the LAN. This is possible if the service processor has an integrated network interface card (NIC) or access to a NIC shared with the server.

Over the ASM interconnect
Data is passed from the service processor over an ASM interconnect network to a second service processor. The second service processor serves as a gateway between IBM Director Server and the first service processor.

Both of the latter types of communication are known as *out-of-band* communication, because they take place independent of an operating system.

An *ASM interconnect network* is a group of service processors that are networked together using the ASM interconnect feature. Connected through the RS-485 ports, the service processors can communicate with and send alerts out-of-band to IBM Director Server through a *gateway service processor* (sometimes called an ASM interconnect gateway). An ASM interconnect network eliminates the need for multiple modems, telephones, and LAN ports; it also permits service processors without network interface cards to communicate out-of-band with IBM Director Server.

Notes:

1. For IBM Director and Scalable Systems Manager (SSM) to communicate out-of-band, the following conditions must be met:
   - Service processors must maintain consistent IP addresses. You either must assign static IP addresses or configure Dynamic Host Configuration Protocol (DHCP) to maintain consistent IP addresses for the service processors.
   - The service processor IP addresses cannot change after IBM Director has discovered the server.
2. Only one of the following systems management applications can communicate with a service processor at any given time:
   - IBM Director Server
   - IBM Management Processor Command-Line Interface (MPCLI)

In-band communication and alerts

This topic provides information about when service processors can communicate in-band with and send alerts to IBM Director Server .

Whether a service processor can communicate in-band with IBM Director Server depends on both the type of service processor and the operating system running on the managed system.

Table 8. In-band communication between service processors and IBM Director Server

	Operating system		
Primary service processor	**Linux**	**NetWare**	**Windows**
Advanced System Management PCI Adapter (ASM PCI Adapter)	Yes	Yes	Yes
Advanced System Management processor (ASM processor)	Yes	Yes	Yes
Integrated system management processor (ISMP)	Yes	No	Yes
IPMI baseboard management processor	Yes	No	Yes
Remote Supervisor Adapter	Yes	Yes	Yes
Remote Supervisor Adapter II	Yes	Yes[1]	Yes
[1] Novell NetWare 6.5 only			

In addition, to enable in-band communication between IBM Director Server and a managed system that contains a service processor, both the service processor device driver and MPA Agent must be installed on the managed system.

When in-band communication is possible, alerts are handled either by MPA Agent or System Health Monitoring. Unless the server supports System Health Monitoring, ISMPs in servers running Linux cannot send alerts in-band, although in-band communication between the service processor and IBM Director Server is possible.

The following table specifies which IBM Director Agent feature handles in-band alerting.

Table 9. IBM Director Agent features that handle in-band alerts

	Operating system running on managed system		
Type of service processor	**Linux**	**NetWare**	**Windows**
ASM PCI Adapter	MPA Agent	MPA Agent	System Health Monitoring
ASM processor	MPA Agent	MPA Agent	System Health Monitoring
ISMP	None or System Health Monitoring[1]	Not applicable	System Health Monitoring
IPMI baseboard management processor	System Health Monitoring	Not applicable	System Health Monitoring
Remote Supervisor Adapter	MPA Agent or System Health Monitoring[2]	MPA Agent	System Health Monitoring
Remote Supervisor Adapter II	MPA Agent or System Health Monitoring[2]	MPA Agent	System Health Monitoring
[1] If System Health Monitoring is supported on the server. [2] MPA Agent handles the alerts, unless System Health Monitoring is supported on the server.			

See the *IBM Director Hardware and Software Compatibility* document for a list of servers on which System Health Monitoring is supported when the server is running Linux. This PDF file is updated every six to eight weeks. You can download it from the IBM Director Web site at www.ibm.com/servers/eserver/xseries/systems_management/ibm_director/.

Out-of-band communication and alerts

This topic provides information about when service processors can communicate out-of-band with IBM Director Server. It also contains information about the pathways on which service processors can provide out-of-band alerts to IBM Director Server.

The type of service processor present in a server determines which paths out-of-band communication can take. Servers that contain ISMPs can communicate out-of-band with IBM Director Server only through a gateway service processor.

The following service processors all can serve as gateway service processors:
- ASM PCI Adapter
- ASM processor
- Remote Supervisor Adapter
- Remote Supervisor Adapter II

However, some of these service processors cannot communicate with certain other service processors. In addition, an ASM processor can communicate with IBM Director Server only through interprocess communication.

The following table details the possible gateway service processors and the types of service processors located on an ASM interconnect network with which they can communicate.

Table 10. Gateway service processors and communication with service processors on an ASM interconnect network

	Service processor on an ASM interconnect					
Gateway service processor	**ASM processor**	**ASM PCI adapter**	**ISMP**	**IPMI baseboard management controller**	**Remote Supervisor Adapter**	**Remote Supervisor Adapter II**
ASM PCI adapter	Yes	Yes	No	Not applicable	No	No
ASM processor	Yes	Yes	No	Not applicable	No	No
Remote Supervisor Adapter	Yes	Yes	Yes	Not applicable	Yes	Yes
Remote Supervisor Adapter II	Yes	Yes	Yes	Not applicable	Yes	Yes

To maximize the possibility of IBM Director Server receiving alerts from service processors located on an ASM interconnect network, consider using a Remote Supervisor Adapter or a Remote Supervisor Adapter II as a gateway service processor.

Note: If you have one of the following servers attached to an RXE-100 Remote Expansion Enclosure, you cannot use the on-board Remote Supervisor Adapter as a gateway service processor:

- xSeries 360
- xSeries 365
- xSeries 440
- xSeries 445
- xSeries 455

The Remote Supervisor Adapter is dedicated to managing the RXE-100 Remote Expansion Enclosure.

The following table contains information about the pathways available for out-of-band alerting.

Table 11. Out-of-band alerting pathways

Type of service processor	Pathways for out-of-band alerting	Possible gateway service processors
ASM PCI adapter	• LAN • Over an ASM interconnect	• ASM PCI adapter • Remote Supervisor Adapter • Remote Supervisor Adapter II
ASM processor	• Over an ASM interconnect	• ASM PCI adapter • Remote Supervisor Adapter • Remote Supervisor Adapter II
ISMP	• Over an ASM interconnect	• Remote Supervisor Adapter • Remote Supervisor Adapter II
IPMI baseboard management processor	• LAN	• Not applicable
Remote Supervisor Adapter	• LAN • Over an ASM interconnect	• Remote Supervisor Adapter • Remote Supervisor Adapter II
Remote Supervisor Adapter II	• LAN • Over an ASM interconnect	• Remote Supervisor Adapter • Remote Supervisor Adapter II

See the documentation that came with the server for information about how to configure your service processor and ASM interconnect to ensure that IBM Director Server receives alerts. The IBM Redbooks publication *Implementing Systems Management Solutions Using IBM Director* (SG24-6188-01) also contains information that might be helpful.

Out-of-band alert-forwarding strategies

This topic provides information about the out-of-band alert-forwarding strategies that are supported by xSeries service processors.

The type of service processor also determines what type of alert-forwarding strategy is possible. The following table contains information about possible alert-forwarding strategies.

Table 12. Out-of-band alert-forwarding strategies

Type of service processor	Possible alert-forwarding strategies
ASM PCI adapter	IBM Director over LAN
ASM processor	IBM Director over LAN

Table 12. Out-of-band alert-forwarding strategies (continued)

Type of service processor	Possible alert-forwarding strategies
ISMP	Not applicable
IPMI baseboard management processor	IBM Director comprehensive
Remote Supervisor Adapter	IBM Director over LAN IBM Director comprehensive
Remote Supervisor Adapter II	IBM Director comprehensive

Some service processors also support SNMP as an alert-forwarding strategy.

Upward integration

Upward integration modules (UIMs) enable third-party workgroup and enterprise systems-management products to interpret and display data that is provided by Level-1 and Level-2 managed systems. The UIMs provide enhancements to the systems-management products that you can use to start IBM Director Agent from within the systems-management platform, collect inventory data, view IBM Director event notifications, and for some UIMs, distribute IBM Director managed system software packages.

With the UIMs, you can use your systems-management software to manage systems installed with IBM Director Core Services or IBM Director Agent software. You can use IBM Director Agent software to:

- Gather detailed inventory information about your systems, including operating system, memory, network adapters, and hardware.
- Track your systems with features such as power management, event log, and system monitor capabilities.

IBM Director Agent uses some of the latest systems-management standards, including Common Information Model (CIM), Web-Based Enterprise Management (WEBM) and Extensible Markup Language (XML), to provide compatibility with your existing enterprise-management software.

IBM Director enables you to make the most of your existing enterprise management structure by upwardly integrating with Tivoli Management Framework, Tivoli NetView, HP OpenView, and Microsoft Systems Management Server (SMS), and Microsoft Operations Manager (MOM).

IBM Director UIM for HP OpenView

With the IBM Director UIM for HP OpenView, you can use your systems-management software to manage systems installed with IBM Director Core Services or IBM Director Agent software.

When you install IBM Director UIM for HP OpenView, the following functions are added to the HP OpenView environment:

- **Event notification**: Provides notification of events that occur on managed systems on which IBM Director Agent is installed. These notifications are delivered using SNMP traps.
- **Inventory**: Scans inventory using an inventory plug-in that starts a Java application that collects the inventory from IBM Director Agent, including Asset ID data, BIOS details, and lease information.

- **Web browser launch**: Provides Web browser capability from within the HP OpenView environment so that you can display and manage real-time asset and health information about managed systems on which IBM Director Agent is installed.
- **Discovery**: Provides SNMP-based discovery of managed systems on which IBM Director Agent is installed.

  Note: You must configure the SNMP community name of the managed system.

IBM Director UIM for Microsoft Operations Manager

With the IBM Director UIM for Microsoft Operations Manager, you can use your systems-management software to manage systems installed with Level-1: IBM Director Core Services or Level-2: IBM Director Agent software.

When you install IBM Director UIM for Microsoft Operations Manager (MOM), the following functions are added to the Microsoft Operations Manager environment:

- **Discovery**: Provides discovery of Level-1 and Level-2 managed systems.
- **Events**: Captures events that occur on Level-1 and Level-2 managed systems.
- **Alerts**: Sends a notification when certain events occur on Level-1 and Level-2 managed systems.
- **State**: Changes the state of Level-1 and Level-2 managed systems based on event criteria.

IBM Director UIM for Microsoft Systems Management Server

With the IBM Director UIM for Microsoft Systems Management Server (SMS), you can use your systems-management software to manage systems installed with IBM Director Core Services or IBM Director Agent software.

When you install IBM Director UIM for SMS, the following functions are added to the SMS environment:

- **Event notification**: Provides notification of events that occur on managed systems on which IBM Director Agent is installed. These notifications are translated into SMS status messages.
- **Collections**: Adds an SMS Collection to easily identify all managed systems on which IBM Director Agent is installed.
- **Inventory**: Scans inventory directly from IBM Director Agent, including Asset ID data, BIOS details, field-replaceable unit (FRU) numbers, lease information, and network details.

  Tip:

  - The inventory feature is compatible only with IBM Director Agent 4.20 or later.
- **Queries**: Adds an SMS Query to identify all managed systems on which IBM Director Agent is installed.
- **Software distribution**: Distributes an IBM Director Agent software package and performs an unattended installation on any system in the Microsoft SMS environment.
- **Wake on LAN®**: Remotely turns on managed systems on which IBM Director Agent is installed, and are Wake-on-LAN-capable.

IBM Director UIM for Tivoli Management Framework

With the IBM Director UIM for Tivoli Management Framework, you can use your systems-management software to manage systems installed with IBM Director Core Services or IBM Director Agent software.

When you install IBM Director UIM for Tivoli Management Framework , the following functions are added to the Tivoli Management Framework environment:

- **Event notification**: Provides notification of events (such as failing components) occurring on IBM Director Agent systems and IBM Management Processors, allowing IT personnel to take immediate corrective action. These notifications can be sent as native Tivoli Enterprise Console events, SNMP traps, and Windows event log events.
- **Inventory**: Collects inventory data directly from IBM Director Agent is installed using custom MIF files, SQL scripts, and inventory queries.
- **Monitors**: Provides hardware status monitors for managed systems on which IBM Director Agent is installed. This feature enhances the Tivoli Console interface by providing a richer set of features and more comprehensive hardware monitoring capabilities. You can monitor hardware status and various thresholds.
- **Software distribution**: Enables you to build and distribute update packages for IBM Director Agent software and perform an unattended installation of these packages on any Tivoli endpoint running Microsoft Windows.
- **Tasks**: Allows you to view additional information and restart or shut down managed systems on which IBM Director Agent is installed remotely using Wake on LAN.

IBM Director UIM for Tivoli NetView

With the IBM Director UIM for Tivoli NetView, you can use your systems-management software to manage systems installed with IBM Director Core Services or IBM Director Agent software.

When you install IBM Director UIM for Tivoli NetView, the following functions are added to the Tivoli NetView environment:

- **Event notification**: Provides notification of events (such as failing components) occurring on IBM Director Agent systems and IBM Management Processors, allowing IT personnel to take immediate corrective action. Notifications are delivered through SNMP traps.
- **Inventory**: Collects the inventory data from IBM Director Agent, including Asset ID data, BIOS details, FRU service numbers, lease information, and network details.
- **Web browser launch**: Provides Web-browser capability from within the NetView environment that allows you to view and manage real-time asset and health information about managed systems on which IBM Director Agent is installed.
- **Discovery**: Automatically finds systems with the IBM Director agent installed, using SNMP. From NetView, you can identify Director agent systems at a glance.

  Note: You must configure the SNMP community name of the managed system.

User interfaces

There are three methods for managing an IBM Director environment: a graphical user interface, called the IBM Director console, a command-line interface (dircli), and a Web-based interface.

IBM Director Console

The IBM Director Console allows you to control and monitor managed systems and devices from an application-based graphical user interface. You can install this console on a desktop computer, workstation, or mobile compute that exists on the same network as the management server.

IBM Director command-line interfaces

You can use the IBM Director the command-line interfaces to manage and monitor the managed objects and devices.

You can use the administrative command-line interface interactively using the **dircli** or **dircmd** utilities. This administrative command-line interface is an important primary interface into IBM Director and may be used either as an efficient way to accomplish simple tasks directly or as an embeddable and scriptable framework for achieving higher level goals. For security reasons, administrative command-line interface runs only on the management server.

Note: The IBM Director **dircli** supports a subset of the commands that were available previously through the deprecated **dircmd** utility.

To access **dircli** or **dircmd** you must log in to an management server as an IBM Director super user. Access to **dircli** and interfaces is limited to IBM Director super-users (members of the DirSuper group). By default, the connection between the CLI client and the management server is a nonsecure TCP/IP data link. You can use Secure Sockets Layer (SSL) to secure the data transmission.

Web-based Access

Web-based Access allows you to view managed system information, change alert standard format (ASF) alerts, change system settings and configurations from a Web-based graphical user interface. When you install Web-based Access on a managed system, you can access IBM Director Agent and view real-time asset and health information about the managed system.

Note: This feature is supported only on Windows 32-bit operating systems.

z/VM Center concepts

This topic provides information about the z/VM Center task and the z/VM environment that z/VM Center manages.

z/VM Center uses the IBM System z9 and @server zSeries virtualization technologies, in particular z/VM, to provision System z9 and zSeries resources in form of z/VM virtual servers. To z/VM, a z/VM virtual server is a guest virtual machine.

A z/VM guest virtual machine consumes a portion of the processor cycles, memory, and I/O bandwidth of the System z9 or zSeries hardware. All operating systems that can run natively on System z9 or zSeries can also run on a z/VM virtual server.

The z/VM Center task comprises two subtasks:

- Virtual Server Deployment
- Server Complexes

What you can do with the z/VM Center tasks

This topic explains what you can do with Virtual Server Deployment and Server Complexes and how they relate.

You can use the z/VM Center subtasks to virtualize your System z9 or zSeries hardware resources into z/VM virtual servers (guest virtual machines) and to deploy Linux instances on them. You interact with z/VM through a graphical user interface that runs on the IBM Director Console. You only have to directly work with z/VM for setting up z/VM Center and to install and set up master Linux instances that serve as the source for the Linux instances you deploy on your z/VM virtual servers.

The Virtual Server Deployment and Server Complexes task complement one another.

Virtual Server Deployment

With the Virtual Server Deployment task, you can define configurations of guest virtual machines and save them as virtual server templates. From a virtual server template you can then create numerous z/VM virtual servers all with the characteristics defined in the template. You can use specially prepared master Linux instances as sources for deploying Linux instances on z/VM virtual servers.

With Virtual Server Deployment you can manage the characteristics of your z/VM virtual servers.

Use the Virtual Server Deployment task to manage individual z/VM virtual servers and operating system instances and to set up templates and Linux guest systems to be used by Server Complexes.

Server Complexes

With the Server Complexes task you can manage configurations of *Linux guest systems*. A Linux guest system is a combination of a Linux instance and the z/VM virtual server on which the Linux instance is installed. A server complex is a configuration profile for Linux guest systems and includes both Linux and z/VM aspects. A server complex can define network settings, Linux configuration scripts, disk access, and VM Resource Manager (VMRM) performance goals.

You can automatically configure a Linux guest system by assigning it to a server complex. You can also create a new Linux guest system within a server complex. When creating a new Linux instance in a server complex, you automatically create a z/VM virtual server with a Linux instance that is configured according to the server complex. For creating a Linux guest system, you require a virtual server template and an operating system template that have been created by the Virtual Server Deployment task.

You can make changes to a server complex and then apply the configuration changes to all Linux instances in the server complex.

Use Server Complexes to manage numerous Linux instances with similar configurations.

System z9 and zSeries virtualization

This topic provides information about the zSeries virtualization technologies.

System z9 and zSeries provide two layers of virtualization:

zSeries LPAR hypervisor

is a virtualization technology built into the System z9 and zSeries hardware. With the LPAR hypervisor, you can divide a System z9 or zSeries mainframe into logical partitions (LPARs). Each LPAR has a dedicated portion of the available physical memory (*central storage*, in System z9 and zSeries terminology). Storage devices, I/O channels, and processors can be shared across LPARs or dedicated to a particular LPAR.

You can use the Integrated Facility for Linux (IFL) feature of the System z9 and zSeries hardware to set up LPARs that are restricted to Linux workloads. Such LPARs have processors that cannot run operating systems other than Linux and z/VM.

z/VM is a System z9 and zSeries operating system that acts as virtualization software. z/VM can run in an LPAR. z/VM can virtualize all system resources, including processors, memory, storage devices, and communication devices.

With z/VM, you can run hundreds of operating system instances concurrently, all on the same System z9 or zSeries hardware.

You can use a number of LPARs to concurrently run multiple instances of z/VM while other LPARs run other mainframe operating systems. Each z/VM can run a multitude of mainframe operating systems, including instances of z/VM itself.

z/VM basics:

This topic gives a brief introduction to z/VM and provides a reference to more information on z/VM.

For more detailed information on z/VM, visit the z/VM 5.1 library at ibm.com/servers/eserver/zseries/zos/bkserv/zvmpdf/zvm51.html or the corresponding z/VM 5.2 library.

The z/VM control program

At the core of the z/VM operating system is the control program (CP). CP is a virtualization layer between the System z9 or zSeries hardware and the z/VM guest virtual machines. CP runs on the System z9 or zSeries machine architecture.

As illustrated in Figure 4, a guest virtual machine is a virtualized System z9 or zSeries machine with a fraction of the actual hardware resources but with the same machine architecture.

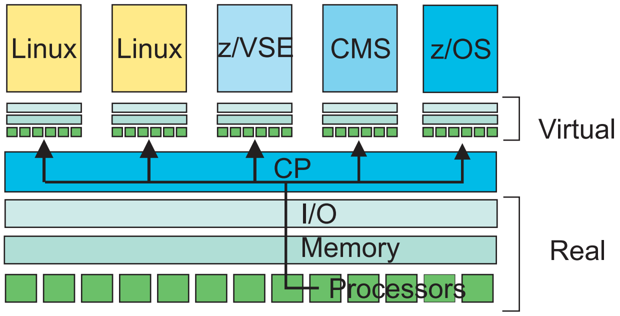

Figure 4. z/VM guest virtual machines

Any operating system or standalone program that can run natively on a System z9 or zSeries machine can also run in a guest virtual machine. To a System z9 or zSeries operating system a guest virtual machine looks like real System z9 or zSeries hardware. An operating system that runs in a z/VM guest virtual machine is called a z/VM guest operating system.

A guest virtual machine runs in the context of a z/VM user ID. With this z/VM user ID you can log on to the virtual hardware. To log on to a Linux guest operating system you do not need the z/VM user ID. You can use your Linux user ID and password, as usual.

The z/VM directory

z/VM uses the z/VM directory to keep track of its guest virtual machines. For each guest virtual machine, there is a directory entry with a number of statements that define its characteristics.

For example, the directory entry defines the processing power, memory size (virtual storage, in z/VM terminology), disk access permissions and other privileges.

The directory is well-protected from general access. There are predefined z/VM users that are privileged to perform administrative functions. Maintaining the directory is among the tasks that require the highest privilege level in z/VM. Many installations use a security manager in addition to this built-in security (see “Security manager” on page 73).

Service machines

z/VM includes a number of service machines. Service machines are guest virtual machines that provide specific services to other guest virtual machines. For example, there are service machines that run programs required for communications or printing. Like all guest virtual machines, service machines are associated with user IDs.

Table 13 shows examples of service machines that are directly relevant to z/VM Center:

Table 13. Examples of service machines

User ID	Purpose
VSMSERVE	This service machine implements the z/VM systems management API. z/VM Center uses it to interact with z/VM and its guest virtual machines.
TCPIP	This service machine runs a TCP/IP stack and defines an IP address through which z/VM can be addressed. z/VM Center accesses VSMSERVE through the IP address provided by TCPIP.
DIRMAINT	This service machine provides an interface for maintaining the z/VM directory. Instead of using DIRMAINT, you can use any directory management program that provides equivalent functions and the required interface.
DATAMOVE	This service machine has privileges to perform disk copy operations. z/VM Center uses DATAMOVE to make copies of disks it requires for setting up new operating system instances and templates.

Note: z/VM is a highly customizable operating system. The user IDs of Table 13 are the standard user IDs for the respective service machines. These IDs can be renamed and might be different on your installation. Your installation might also have multiple instances of a particular service machine, each with a different user ID.

Virtual networking

z/VM allows for a multitude of methods for communication between a guest operating system and another guest operating system on the same z/VM, an operating system instance elsewhere on the same System z9 or zSeries mainframe, or a networked operating system that runs on a separate hardware. For a comprehensive description of z/VM communications refer to *z/VM Connectivity*, SC24-6080.

This section briefly introduces three methods that are particularly relevant to Linux as a guest operating system:

- Direct connections from the z/VM guest virtual machines to an Open Systems Adapter (OSA) card
- Guest LAN
- Virtual switch

You can set up a virtual connection from each guest virtual machine to an OSA card. The OSA card provides a connection to a LAN outside the System z9 or zSeries mainframe. All guest virtual machines that are connected to the same OSA card can also communicate with one another (see Figure 5). *Connecting* in this context does not involve physical cables but means issuing commands that define virtual connections.

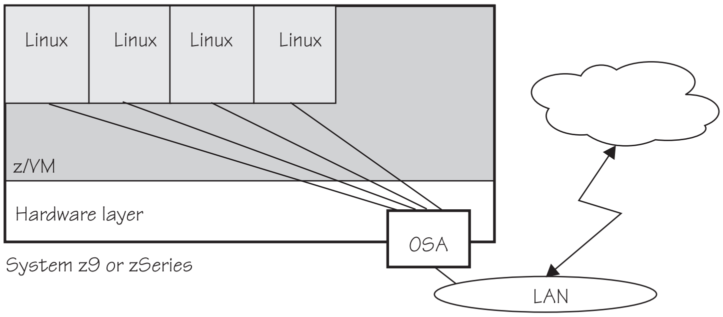

Figure 5. z/VM guest virtual machines directly connected to an OSA card

You can also define a guest LAN. A guest LAN is a virtual LAN, emulated by z/VM. Because a guest LAN does not use physical cables and is contained entirely within the mainframe, it is fast and, if configured correctly, highly secure.

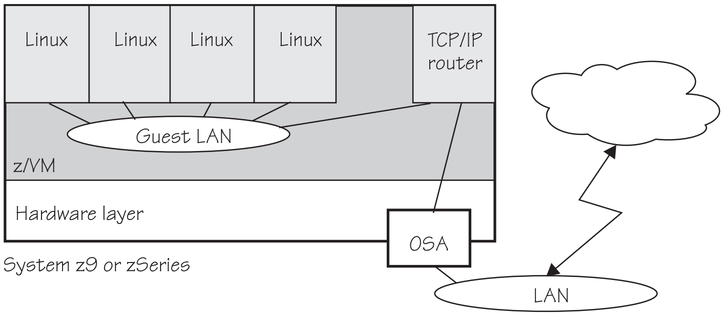

Figure 6. z/VM guest virtual machines connected to a guest LAN

If you want to provide your guest virtual machines with a connection outside the z/VM, you can include a TCP/IP router in your guest LAN (see Figure 6). The router can be a guest virtual machine with a Linux instance as the guest operating system or it can be a TCPIP service machine.

You can also use a virtual switch to connect your guest virtual machines. Like a guest LAN, a virtual switch does not use physical cables and can provide a fast and highly secure connection.

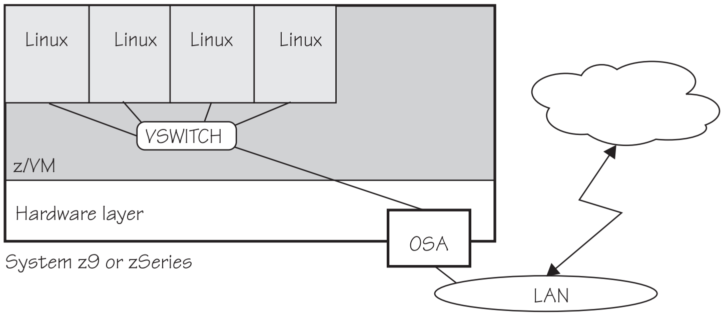

Figure 7. z/VM guest virtual machines connected to a virtual switch

To provide your guest virtual machines with a connection outside the z/VM, you can directly connect the virtual switch to an OSA card (see Figure 7). No router is needed in conjunction with a virtual switch.

Cloning:

This topic explains how copy and personalization techniques can be used to create new instances of guest operating systems.

When z/VM Center creates a new instance of an operating system, it creates a copy of an existing instance. The copy has the same characteristics as the original but instance-specific data is personalized to make the copy a unique instance, rather than a backup of the original. This copying in conjunction with a personalization is called *cloning*.

For example, a cloned Linux instance has the same network interfaces, directory structure, and installed programs as the original, but it has its own unique host name and IP addresses.

Cloning in the Virtual Server Deployment task

With the Virtual Server Deployment task, you do not directly clone operating system instances. You first use cloning to create an operating system template. You then use cloning again to create one or more operating system instances from the template. To create a clone of an operating system instance, you need an existing z/VM virtual server.

Cloning in the Server Complexes task

With the Server Complexes task, you work with z/VM virtual servers that have Linux installed on them. In single cloning operation, you directly create a z/VM virtual server, apply an operating system image to it, and configure it according to the target server complex properties.

For a direct cloning operation you need:

- A server complex that has been created with the Server Complexes task
- A virtual server template that has been created with the Virtual Server Deployment task
- An operating system template that has been created with the Virtual Server Deployment task
- A disk pool to provide the disk resources for the clone

Disk sharing:

Disk sharing can save disk resources.

To z/VM, the disks that a Linux guest operating system can access are minidisks. A minidisk is a logical representation of a part or all of a Direct Access Storage Device (DASD). If multiple operating systems that run on a z/VM require access to the same data, you can share the minidisk where this data resides. When a minidisk is shared, multiple operating system instances can access it.

z/VM Center typically provisions multiple operating system instances that are based on the same master system. Because of the common base, these systems have similar data on the respective system disks. z/VM Center allows the operating system instances to share system disks that:

- Are identical for all operating system instances that need them
- Do not need to be written to

All operating system instances use the same physical copy of a shared minidisks which saves disk space. z/VM ensures data integrity on shared minidisks by restricting access by the operating system instances to read-only.

You can use minidisk sharing most effectively if you set up your master operating systems such as to have the following data on separate minidisks:

- Read-only data
- Read-write and instance specific data

For example, when installing a master Linux you might want to design the file system such that the /etc, /var, and /home directories and any other directories that contain instance-specific data or data that needs to be written to are mounted from a separate minidisk, while the /usr and other directories with fixed data are on a minidisk that can then be shared.

Security manager:

This topic describes considerations if your z/VM installation uses a security manager in addition to the built-in security.

z/VM resources are protected through the access permissions and privileges defined in the z/VM directory. For example, a guest virtual machine cannot access a disk, unless there is a specific or generic directory statement that permits this access.

The directory source file and the facilities to make directory changes take effect are sensitive resources for z/VM security. Accordingly, these resources are well protected in a properly set up z/VM.

In addition to the built-in security, z/VM offers an API for security managers. Many installations use Resource Access Control Facility (RACF®) as their security manager.

If you are using a resource manager in conjunction with z/VM Center, you might need to use the security manager to grant access permissions and privileges for any objects you create with z/VM Center. z/VM Center handles the z/VM directory for you. If your security manager requires additional permissions, you need to define them using your security manager interfaces, outside z/VM Center. Table 14 provides an overview of the objects that you might have to define to your security manager.

Table 14. z/VM Center objects and required security manager definitions

z/VM Center object	**z/VM view of the object**	**Required security manager definitions**
z/VM virtual server	A guest virtual machine	• User ID • Optional: Password
Operating system template	A guest virtual machine that cannot be logged on to but owns a number of disks with an install image, that is an installed operating system that is ready to be booted for the first time	• User ID • Optional: Disk access
Operating system instance	A number of disks where an operating system instance resides; if there are shared disks they are owned by the guest virtual machine that corresponds to an operating system template; the remaining disks are owned by the guest virtual machine on which the operating system has been installed	• Access to network interfaces • Optional: Disk access

When you set up your z/VM for z/VM Center, you also create some z/VM guest virtual machines and need to ensure that the security manager grants the required privileges and permits access to the required disks.

Refer to your security manager documentation to find out how to grant a particular access or privilege.

The z/VM manageability access point

The z/VM manageability access point provides a CIM based remote interface for managing z/VM.

The Common Information Model (CIM) is a standard defined by the Distributed Management Task Force (DMTF). The management functions provided at that interface satisfy the CIM profile for z/VM management. Visit dmtf.org for more information on CIM and DMTF.

In terms of z/VM, the z/VM manageability access point is a guest virtual machine that is privileged to use the z/VM systems management API and the z/VM directory manager interface.

In terms of IBM Director, the z/VM manageability access point is a managed object that can be discovered by IBM Director. It is a Linux system that runs in a z/VM guest virtual machine and has IBM Director Agent and implements the CIM profile for z/VM management (z/VM management profile). Figure 8 shows the manageability access point in relation to z/VM.

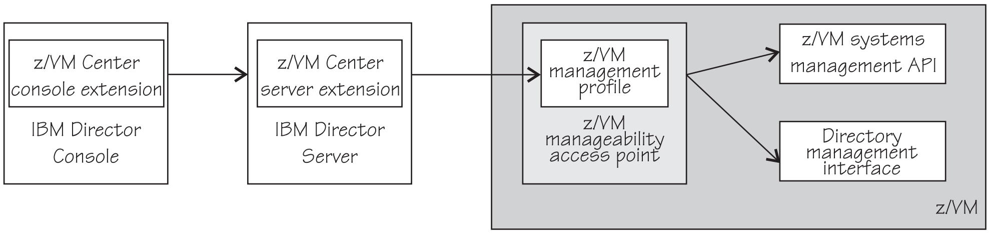
z/VM Center console extension
IBM Director Console
z/VM Center server extension
IBM Director Server
z/VM management profile
z/VM manageability access point
z/VM systems management API
Directory management interface
z/VM

Figure 8. z/VM manageability access point

z/VM Center uses the z/VM management profile to provision systems under z/VM.

To interact with the z/VM manageability access point, you need z/VM Center extension code installed on the IBM Director Console on which you are working and on your IBM Director Server. You can then work with the z/VM Center task on the IBM Director Console. On z/VM, the z/VM manageability access point uses the z/VM systems API and the directory management interface to perform the operations according to the user actions on the IBM Director Console.

The z/VM management profile is not for exclusive use by z/VM Center but is available to any z/VM systems management application.

Virtual Server Deployment concepts

This topic introduces some of the basic concepts of the Virtual Server Deployment task.

z/VM virtual servers and virtual server templates:

All z/VM virtual servers are guest virtual machines but not all guest virtual machines are z/VM virtual servers.

A guest virtual machine is not a z/VM virtual server, if its z/VM directory entry includes a NOLOG statement, that is, if it cannot be logged on.

You use a virtual server template to create a z/VM virtual server. A virtual server template contains configuration data for creating a directory entry for the z/VM virtual server on z/VM. The template relieves you from having to enter all details each time you want to create a z/VM virtual server. You can maintain multiple templates for z/VM virtual servers with different characteristics. Virtual server templates are stored in the CIMOM data repository and are not defined in the z/VM directory.

When creating a z/VM virtual server, the Virtual Server Deployment task uses:

- z/VM defaults. Advanced z/VM users can create their own set of defaults by defining a prototype in the z/VM directory.
- Data from the virtual server template. If you are using a prototype, the data from the virtual server template complements or overrides data from the prototype.
- Data you provide when creating the z/VM virtual server. Data you provide when creating the z/VM virtual server complements or overrides data from the virtual server template.

Once a z/VM virtual server is in place, you can create an operating system instance on it.

Master systems, operating system templates, and personalization:

Master systems are the sources on which Virtual Server Deployment models the operating system instances it creates.

Virtual Server Deployment needs the following to create an operating system instance on a z/VM virtual server:

- An operating system template with a fully configured operating system instance.
- Instance-specific data for the new operating system instance.

Virtual Server Deployment uses cloning techniques to create new operating system instances. You first configure an operating system instance as a model for new operating system instances. To make this operating system instance into a *master operating system instance*, you install a personalization script on it.

When you create a new operating system instance with Virtual Server Deployment, you must specify some instance-specific data. When the newly created operating system instance is started for the first time, the personalization process reads the provided instance-specific data and personalizes the operating system instance accordingly.

Virtual Server Deployment does not clone new operating system instances directly from a master operating system instance. Instead, you create an operating system template. To z/VM, an operating system template is a guest virtual machine that owns disks but cannot be activated. In more general terms, an operating system template can be considered an *install image*.

You can activate the master operating system instance, for example, to make software updates. You can create more than one operating system template from the same master operating system instance.

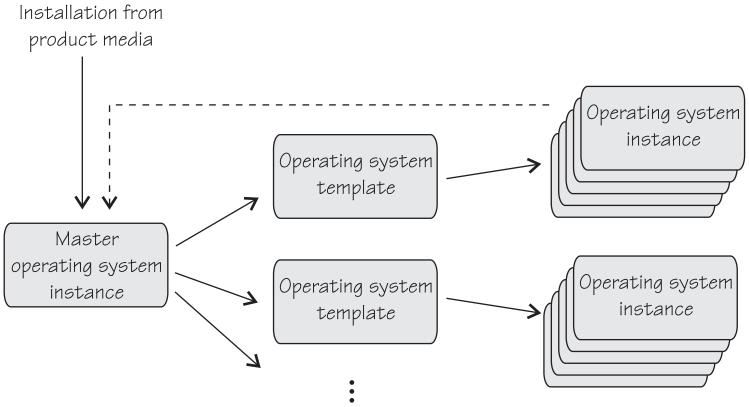
Installation from product media
Master operating system instance
Operating system template
Operating system template
Operating system instance
Operating system instance

Figure 9. Operating system creation flow

As illustrated in Figure 9, you can create one or more operating system templates from a master operating system instance. From each operating system template, you can create one or more operating system instances. The master operating system instance can be from an installation that takes place outside z/VM Center or it can be one of the operating system instances that have been created from an operating system template, within z/VM Center.

At least an initial master operating system instance must have been installed outside z/VM Center.

Registration of operating system instances:

When you register an operating system instance, you provide configuration data about the operating system instance to Virtual Server Deployment.

z/VM Center cannot detect data on the internal configuration of z/VM guest operating system instances. If you have installed an operating system instance outside z/VM Center, you need to provide configuration data about the operating system instance to z/VM Center. This process is called registration. z/VM Center cannot work with operating systems that have been installed outside z/VM Center, unless you register them.

You do not need to register operating system instances that have been created from an operating system template. z/VM Center takes the required configuration data from the template from which they were created.

Regardless of how configuration data on a particular operating system instance has been provided to Virtual Server Deployment—by registration or through a template—z/VM Center cannot detect any changes you might make to that configuration data. If you want to use an instance as a master operating system instance, the actual operating system configuration data must match the data that was registered or taken from an operating system template.

You re-register an operating system instance to make updates to the information z/VM Center holds on an operating system instance.

z/VM Center supports only a single operating system instance on each z/VM virtual server. Before you can create or register a new operating system instance, you must first delete an existing one. De-registering an operating system instance in Virtual Server Deployment only deletes the data Virtual Server Deployment holds on the instance. It does not delete any data on the disks where the instance resides. To delete an operating system instance itself and free the disks where it resides, you can delete the z/VM virtual server on which the instance is installed.

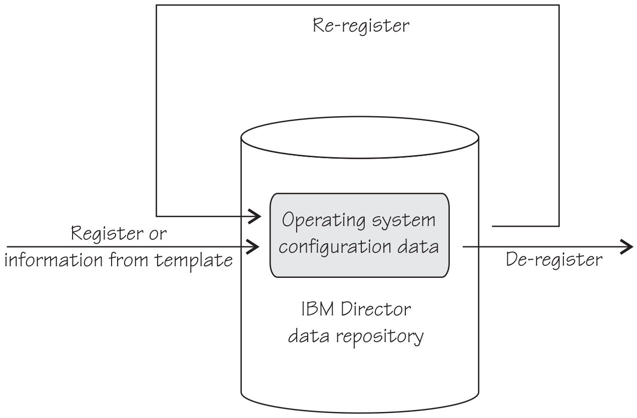
Re-register
Register or information from template
Operating system configuration data
De-register
IBM Director data repository

Figure 10. Register, Re-register, and De-register

Figure 10 summarizes the relationship between register, re-register, and de-register.

Resource names and descriptions:

The z/VM Center graphical user interface provides fields for names and descriptions that you can assign to objects, such as, templates, z/VM virtual servers, disks, and ports.

z/VM Center reads most of the data on the objects from the z/VM directory. Any descriptions and names that you assign to resources are stored in the CIMOM data repository. As an exception, the virtual server template resides entirely in the CIMOM data repository.

z/VM Center uses definitions from z/VM to establish the association between the data in the CIMOM data repository and the z/VM directory. Be aware that names and descriptions can become dissociated from the actual objects when changes are made directly on z/VM.

Server complexes

This topic introduces some of the basic concepts related to the Server Complexes task.

The Server Complexes task lets you control the configuration of your Linux guest systems in an automatic fashion. Within the context of a managed z/VM, you can create as many server complexes as you need.

You set the properties of a server complex to control various configuration aspects of Linux guest systems. Then, when you add a Linux guest system to a server complex, it is automatically configured according to the properties, taking care of the required configuration in both sides—the underlying z/VM and the Linux operating system itself.

You can also directly create a Linux guest system in a server complex. This way, you get a new Linux guest system configured according to the server complex properties in one action. You can also create and configure multiple Linux guest systems this way in one action.

Before beginning to work with server complexes, you should be familiar with the concepts described here:

Server complex:

A server complex represents a (possibly) multi-tier grouping of Linux guest systems. It governs the creation and configuration of included Linux guest systems by persistent properties.

The supported configuration domains are virtual networking, z/VM minidisk attachments, VM Resource Manager performance goals, and configuration scripts. In each configuration domain, you can define the properties separately for each tier, or for the whole server complex.

Linux guest system:

A Linux guest system is an IBM Director managed object. This object represents a Linux system running as a guest operating system in a z/VM virtual server.

From the point of view of IBM Director, a Linux guest system is simply a managed object that is a Linux system.

Server complex properties:

You configure the properties of a server complex, with the intention that these properties will govern the configuration of Linux guest systems contained in that server complex.

You can configure the properties of a server complex within four domains:

- **Virtual Machine Resource Manager (VMRM) velocity goals** – define the CPU and DASD/IO velocity goals to be monitored and adjusted by the VMRM
- **Virtual networking** – define the virtual networking of the Linux guest systems, based on VM guest LANs, direct OSA attachments, or VSWITCHes
- **Post configuration scripts** – add scripts to be run on Linux guest systems when you add them to, or remove them from, the server complex (or tier)
- **Minidisk attachments** – attach minidisks up to the Linux mount point

You can set the properties for all four configuration domains or for a single one. For example, if you are only interested in VMRM configuration, you can create server complexes and specify their properties only for the VMRM domain. Similarly, you can use them for networking or minidisk management.

For each configuration domain, you can specify the properties for either the whole server complex or for each tier separately.

Inconsistencies in server complexes:

Server complex properties govern the configuration state of the Linux guest systems contained within them. However, there are situations in which the configuration of Linux guest systems is not consistent with the configuration implied by the properties of the server complex containing them.

There are two kinds of situations that may cause a state of inconsistency.

Cause 1: If a user changes configuration properties directly on the managed systems and not through the Server Complexes task functionality.

For example, a user sets the CPU velocity goal property value of a tier to 10 and moves Linux guest systems lnx001, lnx002 (or directly creates them) into the tier. This triggers a VMRM configuration of a workload for that tier and a corresponding CPU velocity goal of 10. Then, the user manually edits the VMRM configuration and changes the velocity goal to 20. This results in an inconsistency.

Cause 2: Configuration failure.

For example, a user changes a property value and applies the reconfiguration on all the Linux guest systems in the server complex (or the tier). However, one Linux guest system is down, so the reconfiguration can not be performed on it. The configuration failure for that Linux guest system is 'remembered' according to each configuration domain.

Chapter 2. Managing IBM Director

You can use the administration features of IBM Director to manage the product effectively.

Accessing online help in IBM Director

The IBM Director interface allows you the flexibility to access online help using a variety of methods.

To access online help for IBM Director, choose from one of the following methods:

Using the Console menu

To access online help for an IBM Director task from the **Console** menu, complete the following steps:

1. In the Tasks pane of IBM Director Console, click the task about which you want to find help.
2. If the Tasks pane does not appear, click the triangle along the right side of IBM Director Console, or from the menu bar click **View** → **Tasks Pane**.
3. In the IBM Director Console window, click **Console** → *IBM Director task* → **Help**.
4. The IBM Director Help window displays an overview of the IBM Director task.

Using the Tasks menu

To access online help for an IBM Director task from the **Tasks** menu, complete the following steps:

1. In the IBM Director Console window, click **Tasks** → *IBM Director task* → **Help**.
2. The IBM Director Help window displays an overview of the IBM Director task.

Using the Help menu

To access online help for an IBM Director task from the **Help** menu, complete the following steps:

1. In the IBM Director Console window, click **Help** → **IBM Director Help**.
2. The IBM Director Help window displays the table of contents for the online help.
3. If a specific task window or option window displays, you can click **Help** → *IBM Director task*.
4. The IBM Director Help window displays field-level help for the task window or the option window.

Using the F1 key

To access online help for an IBM Director task using the F1 key, complete the following steps:

1. When the IBM Director Console window is displayed, press the F1 key.
2. The IBM Director Help window displays the table of contents for the online help.

3. If you want to find help for a specific task window or option window, press F1 when that window displays.
4. The IBM Director Help window displays field-level help for the window.

Configuring IBM Director Console settings

Console Preferences enables you to customize your IBM Director Console for color and font, background, toolbar options, confirmation prompting, startup options, and monitor attributes.

Configuring IBM Director Console appearance

The Appearance Preferences page enables you to customize the look of your IBM Director Console. You can set the color for text, backgrounds, and links. You can choose a background image and decide whether to show a shadow.

Complete the following steps to customize the background for IBM Director Console:

1. On the IBM Director Console window, select **Options** → **Console Preferences**.
2. On the Console Preferences window, click the **Appearance** tab.
3. Select the text and background colors for IBM Director Console GUI components. You can also select shadow settings and a background image.
4. Click **OK**.

Configuring IBM Director Console colors and fonts

The Accessibility Preferences page allows you to customize your IBM Director Console colors and fonts.

Note:

- If you change the **Accessibility Preferences** while other windows are visible, the window might not be displayed correctly after the change. If this occurs, close and reopen the particular window to fix the problem.
- Operating System changes in font, color and size can be reflected in both the title bar and the client area of the application. IBM Director Console settings for font size and color affect only the client area of the application.

To customize your IBM Director Console colors and fonts, complete the following steps:

1. On the IBM Director Console window, select **Options** → **Console Preferences**.
2. On the Console Preferences window, click **Accessibility Preferences**.
3. Select the color and font settings that you want.
4. Click **OK**.

Configuring IBM Director Console column details

The Details View Preferences page allows you to customize the attributes that you want to display in the columns for the details view.

To customize the IBM Director Console columns for the details view, complete the following steps:

1. On the IBM Director Console window, select **Options** → **Console Preferences**.
2. On the Console Preferences window, click the **Details View Preferences** tab.

3. Select the attributes that you want to display in the columns for the details view.
4. Click **OK**.

Configuring IBM Director Console confirmation prompting

The Prompting Preferences page allows you to customize when you are prompted to confirm specific tasks or procedures.

To specify the situations in which IBM Director Console prompts you for a safety confirmation, complete the following steps:

1. On the IBM Director Console window, select **Options** → **Console Preferences**.
2. On the Console Preferences window, click the **Prompting Preferences** tab.
3. Select the situations in which you want to be prompted for a safety confirmation.
4. Click **OK**.

Configuring IBM Director Console monitor attributes

The Monitor Preferences page enables you to specify whether warnings are displayed for empty attributes, as well as the number of levels to show for the attribute path.

To customize the IBM Director Console monitor attributes, complete the following steps:

1. On the IBM Director Console window, select **Options** → **Console Preferences**.
2. On the Console Preferences window, click the **Monitor Preferences** tab.
3. Select your preference for the display of warnings for empty attributes, and indicate the number of levels you want to show for the attribute path.
4. Click **OK**.

Configuring IBM Director Console startup options

The General Preferences page enables you to customize the IBM Director Console startup options.

To customize the IBM Director Console startup options, complete the following steps:

1. On the IBM Director Console window, select **Options** → **Console Preferences**.
2. On the Console Preferences window, click the **General Preferences** tab.
3. Select your preferences for maintaining the current task consoles or ticker tape monitors when IBM Director Console is restarted.
4. Click **OK**.

Configuring the IBM Director Console toolbar

The Toolbar Preferences window enables you to customize your IBM Director Console toolbar options.

To customize the IBM Director Console toolbar attributes, complete the following steps:

1. On the IBM Director Console window, select **Options** → **Console Preferences**.
2. On the Console Preferences window, click the **Toolbar Preferences** tab.

3. Select the tasks and options you want to display on the toolbar for quick access.
4. Click **OK**.

Configuring SSL settings for IBM Director Console and IBM Director Server

If necessary, you can modify the ports that are used for Secure Sockets Layer (SSL) and restrict IBM Director Console sessions to particular ports and session keys.

Modifying SSL ports

SSL is enabled by default for connections to IBM Director Console. If you need to modify the SSL ports that are used, you can edit the TWGServer.prop and TWGConsole.prop files.

The following table describes the default ports for TCP and SSL communication:

Communication type	Default port number
TCP communication for IBM Director Console	2033
TCP communication for DIRCLI	2044
SSL communication for IBM Director Console and DIRCLI	4066

To modify the ports used for communication between a management server and a management console, you must modify the TWGServer.prop and TWGConsole.prop files located on each system. If you installed IBM Director Server and IBM Director Console in the default location, these files are located in the following directories:

For i5/OS	/QIBM/UserData/Director/data/
For Linux and AIX	/opt/ibm/director/data/
For Windows	*c*:\Program Files\IBM\Director\data

c is the drive letter of the hard disk on which IBM Director is installed.

Complete the following steps to modify the SSL ports used for communication between a management server and management console:

1. Using an ASCII text editor or the i5/OS Edit File (EDTF) command, open the TWGServer.prop file located on the management server.
2. To modify the ports used for IBM Director Console and DIRCLI connections, edit the following statements:

To modify the port for DIRCLI TCP connections	`twg.cli.gateway.link.1.initparam=`*PortNumber*
To modify the port for IBM Director Console TCP connections	`twg.gateway.link.1.initparam=`*PortNumber*

To modify the port for DIRCLI SSL connections	`twg.cli.gateway.link.2.initparam=`*`PortNumber`*
To modify the port for IBM Director Console SSL connections	`twg.gateway.link.2.initparam=`*`PortNumber`*

PortNumber is the port.

3. Save and close the TWGServer.prop file.
4. Stop and restart IBM Director Server.
5. Using an ASCII text editor or the i5/OS Edit File (EDTF) command, open the TWGConsole.prop file located on the management console.
6. In the TWGConsole.prop file, modify the same statements that you modified in the TWGServer.prop file in step 2 on page 84
7. Save and close the TWGConsole.prop file.
8. Stop and restart IBM Director Console.

Restricting IBM Director Console sessions

This topic describes how to modify the TWGServer.prop and TWGConsole.prop files to specify the port used by IBM Director Console sessions and the specific cipher suites to be used.

If you installed IBM Director Server and IBM Director Console in the default location, these files are located in the following directories on the management server and management console:

For i5/OS	/QIBM/UserData/Director/data/
For Linux	/opt/ibm/director/data/
For Windows	*c*:\Program Files\IBM\Director\data

c is the drive letter of the hard disk on which IBM Director is installed.

Complete the following steps:

1. Using an ASCII text editor or the i5/OS Edit File (EDTF) command, open the TWGServer.prop file.
2. Modify the file so that it contains the following properties:
   ```
   twg.gateway.link.1=com.tivoli.twg.libs.TWGSSLLink
   twg.gateway.link.1.initparm=port_number -cipherSuites cipher_suite
   ```
   port_number is the port and *cipher_suite* is the cipher suite.

   Note: Separate multiple cipher suites with a comma; do not add a space after the comma.
3. Save and close the TWGServer.prop file.
4. Stop and restart IBM Director Server.
5. Using an ASCII text editor or the i5/OS Edit File (EDTF) command, open TWGConsole.prop file.
6. In the TWGConsole.prop file, modify the same properties that you modified in the TWGServer.prop file in step 2.
7. Save and close the TWGConsole.prop file.
8. Stop and restart IBM Director Console.

9. (Management servers running Linux or Windows) Copy a cacerts file to the following directory on the management server and name it `cacerts.ssl`:

For Linux	/opt/ibm/director/data
For Windows	*c*:\Program Files\IBM\director\data

10. You can find an existing cacerts file in the following directory:

For Linux	/opt/ibm/director/jre/lib/security
For Windows	*c*:\Program Files\IBM\director\jre\lib\security

11. Import the applicable Rivest-Shamir-Adleman (RSA) or Secure Hash Algorithm (SHA) certificate into the cacerts.ssl file. You can use the keytool program located in one of the following directories:

For Linux	/opt/ibm/director/jre/bin
For Windows	*c*:\Program Files\IBM\Director\jre\bin

To establish an SSL session without importing an RSA or SHA certificate, use an anonymous cipher suite.

Configuring IBM Director Server preferences

Server Preferences allows you to select the types of console connections to allow, view and configure database settings, define the number of event logs and action histories to store, define file distribution servers, set inventory collection preferences, set remote control preferences, select the SNMP version to use for communication, and define software distribution preferences.

Configuring event log and history settings

The Server Preferences window: Event Management page allows you to change the maximum number of event log entries and action history entries that are stored.

Complete the following steps to change event management preferences:

1. On the IBM Director Console window, select **Options** → **Server Preferences**.
2. On the Server Preferences window, click **Event Management**.
3. Edit the **Event Management** settings to fit your needs.
4. Click **OK**.

Viewing and configuring database settings

The Server Preferences window: Database page allows you to view information about the database that is configured for IBM Director, and in some cases you can edit database login settings.

To change software distribution preferences, complete the following steps:

1. On the IBM Director Console window, select **Options** → **Server Preferences**.
2. On the Server Preferences window, click **Database**.
3. If applicable, edit the **Database** login settings to fit your needs.
4. Click **OK**.

Configuring inventory collection preferences

The Server Preferences window: Inventory Collection page allows you to set inventory collection parameters: how often inventory is refreshed, timeout values, the number of managed objects on which inventory is collected simultaneously, and the type of inventory data you want to collect.

Complete the following steps to change inventory collection preferences:

1. On the IBM Director Console window, select **Options** → **Server Preferences**.
2. On the Server Preferences window, click **Inventory Collection**.
3. Edit the **Inventory Collection** settings to fit your needs.
4. Click **OK**.

Configuring remote control preferences

The Server Preferences window: Remote Control page allows you to define timeout values for an inactive remote control session and edit the starting state for a remote control session.

Complete the following steps to change remote control preferences:

1. On the IBM Director Console window, select **Options** → **Server Preferences**.
2. On the Server Preferences window, click **Remote Control**.
3. Edit the **Automatic Timeout** and **Starting Session State** settings to fit your needs.
4. Click **OK**.

Configuring security for console sessions

The Server Preferences window: Connections page allows you to select which type of console connections to allow. By default, both nonsecure TCP connections and secure SSL connections are enabled.

Complete the following steps to change security settings for IBM Director Console communication:

1. On the IBM Director Console window, select **Options** → **Server Preferences**.
2. On the Server Preferences window, click **Connections**.
3. Edit the **Connections** settings to fit your needs.
4. Click **OK**.

Configuring SNMP version for communication

The Server Preferences window: SNMP page allows you to define the SNMP version and profiles to use for communication.

Complete the following steps to change SNMP preferences:

1. On the IBM Director Console window, select **Options** → **Server Preferences**.
2. On the Server Preferences window, click **SNMP**.
3. Edit the **Protocol Settings** to select the version of SNMP to use for communication.
4. You can add, edit or remove **SNMPv3 Profile Settings** to fit your needs.

Setting up software-distribution preferences

The Server Preferences window: Software Distribution page allows you to configure software distribution preferences such as the maximum number of managed systems to which you want to stream software packages concurrently and the bandwidth you want to assign to streaming software packages. You can also specify not to stream a package if a redirected distribution fails as well as restrict the server access check.

Complete the following steps to configure software-distribution preferences:

1. If necessary, start IBM Director Console.
2. Click **Options** → **Server Preferences**.
3. In the Server Preferences window, click the **Software Distribution** tab.
4. On the Software Distribution page, locate the **Maximum Managed Systems** field and type the maximum number of managed systems to which IBM Director Server can concurrently stream software packages. The default value is three.
5. To limit the bandwidth used to stream packages, select the **Enter streaming bandwidth (kbps) for managed systems** check box. In the entry field, type the bandwidth, in kilobytes per second (KBps), that you want to use to stream packages from either IBM Director Server or a file-distribution server to the managed system.

   Note: To specify values less than 1 KBps, type a decimal. The minimum acceptable value is 0.25 (256 bytes per second).
6. To avoid streaming a package in the event that a redirected distribution fails, select the **Do not stream distribution if redirected distribution fails** check box.
7. To prevent IBM Director Server from performing an access check of *all* of the file-distribution shares, select the **Restrict server access check** check box.

   This selection restricts the access check to *only* the file-distribution shares you configure for a specific managed system or group.
8. Click **OK**.

Defining file distribution servers

The Server Preferences window: File Distribution Servers page allows you to define the file distribution servers that you want to use for distributing software.

Complete the following steps to change file distribution preferences:

1. On the IBM Director Console window, select **Options** → **Server Preferences**.
2. On the Server Preferences window, click **File Distribution Servers**.
3. Click **Add**, **Edit**, or **Remove** to define file distribution servers for the shares you have created.
4. Click **OK**.

Setting up software-distribution preferences

The Server Preferences window: Software Distribution page allows you to configure software distribution preferences such as the maximum number of managed systems to which you want to stream software packages concurrently and the bandwidth you want to assign to streaming software packages. You can also specify not to stream a package if a redirected distribution fails as well as restrict the server access check.

Complete the following steps to configure software-distribution preferences:

1. If necessary, start IBM Director Console.
2. Click **Options** → **Server Preferences**.
3. In the Server Preferences window, click the **Software Distribution** tab.
4. On the Software Distribution page, locate the **Maximum Managed Systems** field and type the maximum number of managed systems to which IBM Director Server can concurrently stream software packages. The default value is three.
5. To limit the bandwidth used to stream packages, select the **Enter streaming bandwidth (kbps) for managed systems** check box. In the entry field, type the bandwidth, in kilobytes per second (KBps), that you want to use to stream packages from either IBM Director Server or a file-distribution server to the managed system.

   Note: To specify values less than 1 KBps, type a decimal. The minimum acceptable value is 0.25 (256 bytes per second).
6. To avoid streaming a package in the event that a redirected distribution fails, select the **Do not stream distribution if redirected distribution fails** check box.
7. To prevent IBM Director Server from performing an access check of *all* of the file-distribution shares, select the **Restrict server access check** check box.

   This selection restricts the access check to *only* the file-distribution shares you configure for a specific managed system or group.
8. Click **OK**.

Setting up file-distribution servers

This topic describes how to set up file-distribution servers.

IBM Director supports UNC-based and FTP-based file distribution. See your server documentation for information about setting up a shared subdirectory.

Note: You do not need to install IBM Director on the file-distribution server.

See "File-distribution server considerations" on page 55 for more information.

Configuring IBM Director to use a file-distribution server

This topic describes how to configure IBM Director to use a file-distribution server to distribute software to Level-2 managed systems.

Complete the following steps to configure IBM Director Server to use a file-distribution server:

1. Start IBM Director Console.
2. Click **Options** → **Server Preferences**.
3. In the Server Preferences window, click the **File Distribution Server** tab. A list is displayed of all configured file-distribution servers.
4. Click **Add**.
5. In the Add Share Name window, type the name of the file-distribution server (using UNC notation) in the **Share Name** field. To specify FTP as the transport protocol, begin the share-name entry with `ftp:`, for example `ftp:\\ServerName\AccountName`.

6. In the **Maximum Disk Space** field, type the maximum amount of disk space (in MB) that can be allocated on the file-distribution server for software distribution.
7. In the **Maximum Managed Systems** field, type the maximum number of managed systems that can receive a software package at the same time.
8. To limit the bandwidth that can be used to send packages between IBM Director Server and the file-distribution server, select the **Limit bandwidth between server and share (kbps)** check box. In the entry field, type the maximum bandwidth, in kilobytes per second (KBps), that can be used to send packages between IBM Director and the file-distribution server.

   Note: You might want to limit the bandwidth when a dedicated connection, such as integrated services digital network (ISDN), is used for copying the files from IBM Director Server to the share.
9. If you specified an FTP-based server in step 5 on page 89, you must provide information about the FTP server:
   a. In the **User ID on FTP** server field, type a user ID authorized to access the FTP server installed on the share.
   b. In the **Password** field, type the password associated with the user ID.
   c. In the **Confirm password** field, retype the password associated with the user ID.
   d. In the **Home Directory** field you can specify the directory where you want to cache software packages. If you do not specify a directory, the packages will be cached in the default home directory defined for the FTP user.

      Note: For i5/OS, you must specify a directory, or configure the FTP server to operate in regular mode. By specifying a directory, IBM Director automatically changes the FTP server to operate in regular mode.
10. Click **OK**.

If you have multiple file-distribution servers, repeat this procedure for each server.

Configuring distribution preferences for managed systems

This topic describes how to configure distribution preferences for managed systems.

After you configure IBM Director to use a file-distribution server, you can assign unique policies to a managed system and groups. By default, a managed system attempts to access all shares that have been defined to the management server. You can configure the following software-distribution preferences for a managed system or group:

- Restrict access to specific shares
- Specify whether software distribution occurs through streaming or redirected distribution
- Limit the bandwidth used for software distribution

Complete the following steps to define distribution preferences:

1. If necessary, start IBM Director Console.
2. In the Group Contents pane, right-click the managed system or group.
3. Click **Distribution Preferences**.
4. In the Distribution Preferences window, select the method of software distribution:

- If you want to copy packages directly from IBM Director Server to the managed system or group, click **Always stream to the Managed System(s)**.
- If you want to copy packages from a share to the managed system or group, click **Use File Distribution Server Shares**.

Note: If a file-distribution server share cannot be located during a software distribution when you have selected **Use File Distribution Server Shares**, the default action is to stream the package from IBM Director Server. You can prevent streaming from IBM Director Server by selecting **Do not stream if redirection fails** in the Software distribution server preferences.

5. To add a share, click **Add**.
6. In the Add Share Name window, in the **Share Name** field, select the share. If necessary, specify a user ID and password for an account that can access the share.
7. Click **OK**.
8. Repeat steps 5 through 7 until you have added all of the shares that you want the managed system or group to access.
9. If you want to limit the shares that the managed system or group can access to only those displayed, select the **Restrict share selection to list** check box.

   Note: Windows only: If you select to use a file-distribution server share and specify a user id and password with which to distribute the package, rather than using null credentials, the **Stream from File Distribution Server** check box must be selected for a distribution to complete successfully. This applies to packages created with the IBM Update Assistant, InstallShield Package, or Microsoft Windows Installer Package wizards.
10. To limit the bandwidth that is used when copying packages, select the **Limit streaming bandwidth for system** check box.

    If you have selected **Always stream to system**, type a bandwidth value, in kilobytes per second (KBps), to define the bandwidth that is used to copy packages from IBM Director Server to the managed system or group. If a bandwidth limitation is also set in Server Preferences for streaming from the IBM Director Server, the lower value of the two settings is used as the limitation parameter.

    If you have selected **Use File Distribution Server Shares**, in the entry field, type a bandwidth value to define the bandwidth that is used to copy packages from the file-distribution server share to the managed system or group. If a file-distribution server share is unavailable at the time of distribution, the package is streamed from IBM Director Server; this bandwidth value is used to stream the package, unless a more restrictive bandwidth value has been set in Server Preferences for Software Distribution.

Discovering managed systems, devices, and objects

You can set preferences to discover managed systems in your IBM Director environment automatically. Additionally, you can add managed systems manually.

Discovery is the process by which IBM Director Server identifies and establishes connections with systems and devices that it can manage. The management server sends out a discovery request and waits for responses from managed systems. The managed systems listen for this request and respond to the management server that sent the request.

Before IBM Director can manage a system, that system must be discovered by IBM Director Server. After a system or device has been discovered, an icon that represents the object is displayed in the IBM Director Console window when the applicable group is selected.

Note: (Windows 2000, Server Edition only) The initial discovery performed by the management server is resource intensive. After the initial discovery is completed, the resource utilization returns to normal.

Configuring discovery preferences

This topic describes how to configure discovery preferences.

By default, IBM Director Server automatically discovers all managed systems that are on the same subnet as the management server. If you want to manage systems that are on a different subnet, you must configure discovery preferences.

Configuring discovery preferences for Level-2 managed systems

The Level 2: IBM Director Agents page allows you to customize the discovery parameters for Level-2 managed systems.

Complete the following steps to configure discovery preferences for Level 2: IBM Director Agents:

1. In IBM Director Console, click **Options** → **Discovery Preferences**.
2. In the Discovery Preferences window on the Level 2: IBM Director Agents page, configure the general discovery preferences for managed systems:
   a. In the **Auto-discover period (hours)** field, select how often IBM Director Server attempts to discover systems automatically. The possible range is one to 168 hours (seven days). By default, autodiscovery is disabled.
   b. In the **Presence Check period (minutes)** field, select how often IBM Director Server checks the status of each managed system. A presence check detects whether a managed system is online or offline. The possible range is one to 240 minutes (four hours); this feature also can be disabled. By default, a presence check is set to occur every 15 minutes.
   c. Select the **Automatically secure unsecured systems** check box to ensure that IBM Director Server automatically secures any unsecured managed systems that it discovers. When this feature is enabled, IBM Director Server prevents future management servers from managing the managed systems without first requesting access.
   d. Select the **Auto-add unknown agents which contact server** check box to ensure that IBM Director Server adds objects for newly-discovered managed systems to the Group Contents pane. This option might be useful if you are reinstalling IBM Director Server in an existing IBM Director environment, or if you have configured instances of IBM Director Agent to contact IBM Director Server directly.
3. At the top of the page, click **System Discovery (IP)**.

   Note: These discovery preferences are for Level-2 managed systems only.
4. To specify that IBM Director Server issues an IP broadcast to discover managed systems on the local subnet, on the System Discovery (IP) page, select the **Use TCP/IP general broadcasts** check box. By default, this feature is enabled.
5. To specify that IBM Director Server issues an IP broadcast to discover managed systems on a remote subnet, complete the following steps:

   a. Click **Add**.
   b. In the Add window, click **Broadcast**.
   c. Click **Next**.
   d. In the Add Broadcast Address window, in the **IP Address** and **Subnet Mask** fields, type the IP address and subnet mask.
   e. Click **OK**. The information about the broadcast operation is displayed in the **Address Entries** field.

6. To specify that IBM Director Server issues an IP multicast, complete the following steps:
   a. Select the **Use TCP/IP multicasts** check box.
   b. In the **Multicast Group** field, type the address of the multicast group address. By default, the multicast group address is set to 224.0.1.118.
   c. In the **Multicast TTL** fields, select the time to live for the multicast discovery packet. The time to live is the number of times that a packet is forwarded between subnets. By default, this is set to 32.

   Note: If you modify the multicast group address, you also must modify the multicast group address on each managed system.

7. To specify that IBM Director Server issues a broadcast relay request from a Level-2 managed system, complete the following steps:
   a. Click **Add**. The **Add** window opens.
   b. Click **Relay**.
   c. Click **Next**.
   d. In the Add Relay Address window, in the **IP Address** and **Subnet Mask** fields, type the IP address and subnet mask of an existing Level-2 managed system.
   e. Click **OK**. The information about the broadcast relay operation is displayed in the **Address Entries** field.

8. To specify that IBM Director Server issues a unicast, complete the following steps:
   a. Click **Add**.
   b. In the Add window, click **Unicast Address**.
   c. Click **Next**.
   d. In the Unicast Address window, in the **IP Address** field, type the IP address to which the packet is sent.
   e. Click **OK**. The information about the unicast operation is displayed in the **Address Entries** field.

9. To specify that IBM Director Server issues a unicast to a range of IP address, complete the following steps:
   a. Click **Add**.
   b. In the Add window, click **Unicast Range**.
   c. Click **Next**.
   d. In the Add Unicast Address Range window, in the **Start Address** and **End Address** fields, type the starting and ending IP addresses.
   e. Click **OK**. The information about the unicast range operation is displayed in the **Address Entries** field.

10. If you have IPX installed on the management server, at the top of the page, click **System Discovery (IPX)**.

11. To specify that IBM Director Server issues an IPX broadcast to discover managed systems on the local subnet, on the System Discovery (IPX) page, select the **Use IPX general broadcasts** check box. By default, this feature is enabled.
12. To specify that IBM Director Server issues an IPX broadcast to discover managed systems using a specific IPX address, complete the following steps:
    a. Click **Add**.
    b. In the Add IPX address entry window, type the IPX address.
    c. If you want to enable a broadcast relay, select the **Enable broadcast relay** check box.
    d. Click **OK**. The information about the broadcast operation is displayed in the **Address Entries** field.

Configuring discovery preferences for Level-1 managed systems

The Level 1: IBM Director Core Services Systems page allows you to customize the discovery parameters for Level-1 managed systems.

Complete the following steps to configure discovery preferences for Level-1 managed systems:

1. In IBM Director Console, click **Options** → **Discovery Preferences**.
2. In the Discovery Preferences window on the Level 1: IBM Director Core Services Systems page, define the discovery preferences that you want to use to find Level-1 managed systems.
3. Click **OK**.

Configuring discovery preferences for Level-0 managed systems

The Level 0: Agentless Systems page allows you to customize the discovery parameters for Level-0 managed systems.

Complete the following steps to configure discovery preferences for Level-0 managed systems:

1. In IBM Director Console, click **Options** → **Discovery Preferences**.
2. In the Discovery Preferences window on the Level 0: Agentless Systems page, define the discovery preferences that you want to use to find Level-0 managed systems.
3. Click **OK**.

Configuring discovery preferences for SNMP devices

The SNMP devices page allows you to customize the discovery parameters for SNMP devices.

Complete the following steps to configure discovery preferences for SNMP devices:

1. In IBM Director Console, click **Options** → **Discovery Preferences**.
2. In the Discovery Preferences window on the SNMP Devices page, configure the discovery preferences for SNMP devices.
    a. In the **Auto-discover period (hours)** field, select the frequency with which IBM Director Server attempts to discover SNMP devices automatically. The possible range is one to 168 hours (seven days). By default, autodiscovery is disabled.
    b. In the **Presence Check period (minutes)** field, select the frequency that IBM Director Server checks the status of each SNMP device. A presence check detects whether an SNMP device is online or offline. The possible range is

one to 240 minutes (four hours); this feature also can be disabled. By default, a presence check is set to occur every 15 minutes.

   c. Select the **Auto-add unknown agents which contact server** check box to ensure that IBM Director Server adds objects for newly-discovered SNMP devices to the Group Contents pane.

3. Configure seed addresses for SNMP discovery:
   a. In the **IP Address and Subnet masks** group box, specify the seed addresses. You can specify multiple IP addresses; the addresses are searched concurrently. By default, the IP address of the management server is added to this list. To optimize the chance of discovering all SNMP devices, be sure to specify the IP addresses for routers and DNS servers. During the discovery operation, IBM Director Server locates the address tables located on the specified devices and adds those addresses to the list of addresses to search. The process is repeated for every new SNMP device that is discovered from the new addresses. The discovery operation continues until no more addresses are found.
   b. In the **SNMP Version** field, select the SNMP version.
   c. In the **Community Names** field, type the community names and click **Add**. Continue until all community names are added. For SNMP versions 1 and 2c, be sure to order the community names appropriately.

Configuring discovery preferences for SMI-S storage devices

The SMI-S Storage Devices page allows you to customize the discovery parameters for SMI-S storage devices.

Complete the following steps to configure discovery preferences for SMI-S storage devices:

1. In IBM Director Console, click **Options** → **Discovery Preferences**.
2. In the Discovery Preferences window on the SMI-S Storage Devices page, select the preferences that you want IBM Director to use.
   a. In the **Auto-discover period (hours)** field, select the frequency with which IBM Director Server attempts to discover SMI-S storage devices automatically. The possible range is one to 168 hours (seven days). By default, autodiscovery is disabled.
   b. In the **Presence Check period (minutes)** field, select the frequency that IBM Director Server checks the status of each SMI-S storage device. A presence check detects whether a storage device is online or offline. The possible range is one to 240 minutes (four hours); this feature also can be disabled. By default, a presence check is set to occur every 15 minutes.
   c. Under **Service Location Protocol (SLP) Profiles**, select the SLP profiles that IBM Director Server will discover. The default is SNIA:Array.
   d. Under **Naming Conventions Template for SMI-S Devices**, add the parameters to include in the **Template Name** by selecting a parameter in the **Available Parameters** column, and then clicking **Add** to place it in the **Selected Parameters** column. To remove a parameter from the template name, select a parameter in the **Selected Parameters** column, and then click **Remove**. To restore the original default parameters, select the **Reset to default value** check box. The default is %MANUFACTURER% %HARDWARE_TYPE_MODEL% %HARDWARE_SERIAL_NUMBER%
3. To save your selections, click **OK**.

Configuring discovery preferences for BladeCenter Chassis

The BladeCenter Chassis page allows you to customize the discovery parameters for BladeCenter Chassis.

Complete the following steps to configure discovery preferences for BladeCenter Chassis:

1. In IBM Director Console, click **Options** → **Discovery Preferences**.
2. In the Discovery Preferences window on the BladeCenter Chassis page, configure the general preferences for discovering BladeCenter chassis:
   a. In the **Auto-discover period (hours)** field, select the frequency with which IBM Director Server attempts to discover BladeCenter chassis automatically. The possible range is one to 168 hours (seven days). By default, autodiscovery is disabled.
   b. In the **Presence Check period (minutes)** field, select the frequency that IBM Director Server checks the status of each BladeCenter chassis. A presence check detects whether an SNMP device is online or offline. The possible range is one to 240 minutes (four hours); this feature also can be disabled. By default, a presence check is set to occur every 15 minutes.
3. Specify the conventions that IBM Director Server uses when you rename BladeCenter chassis automatically. By default, this naming template is set to IBM %CHASSIS_MACHINE_TYPE_MODEL% %CHASSIS_SERIAL_NUMBER%.
   a. Click a parameter in the **Selected Parameters** field; then, click **Remove**. Continue until you have removed the default parameters.
   b. In the **Available Parameters** field, click a parameter; then, click **Add**. The parameter is added to the **Selected Parameters** list and the **Naming Template** field. Continue until you have selected all the parameters that you want to use.

   To restore the renaming conventions to the default setting, select the **Reset to default value** check box.

Configuring discovery preferences for physical platforms

The Physical Platforms page allows you to customize the discovery parameters for physical platforms.

Complete the following steps to configure discovery preferences for physical platforms:

1. In IBM Director Console, click **Options** → **Discovery Preferences**.
2. In the Discovery Preferences window on the Physical Platforms page, configure the general preferences for discovering service processors:
   a. In the **Auto-discover period (hours)** field, select the frequency with which IBM Director Server attempts to discover service processors automatically. The possible range is one to 168 hours (seven days). By default, autodiscovery is disabled.
   b. In the **Presence Check period (minutes)** field, select the frequency that IBM Director Server checks the status of each service processor. A presence check detects whether an SNMP device is online or offline. The possible range is one to 240 minutes (four hours); this feature also can be disabled. By default, a presence check is set to occur every 15 minutes.
3. Specify the conventions that IBM Director Server uses when you rename service processors automatically. By default, this naming template is set to IBM %CHASSIS_MACHINE_TYPE_MODEL% %CHASSIS_SERIAL_NUMBER%.

a. Click a parameter in the **Selected Parameters** field; then, click **Remove**. Continue until you have removed the default parameters.

b. In the **Available Parameters** field, click a parameter; then, click **Add**. The parameter is added to the **Selected Parameters** list and the **Naming Template** field. Continue until you have selected all the parameters that you want to use.

To restore the renaming conventions to the default setting, select the **Reset to default value** check box.

Discovering systems automatically

The Discover Systems task allows you to discover managed objects automatically. You can select to discover all managed objects or a particular type of managed object.

To discover systems, complete the following steps:

1. In the IBM Director Console, click **Tasks** → **Discover Systems**.
2. Click the type of system you want to discover:
   - All Managed Objects
   - BladeCenter Chassis
   - Level 0: Agentless Systems
   - Level 1: IBM Director Core Services Systems
   - Level 2: IBM Director Agents
   - Physical Platforms
   - SMI-S Storage Devices
   - SNMP Devices

Adding managed systems manually

You can manually define new systems for IBM Director to manage. You can add Level-0 managed systems, Level-1 managed systems, and Level-2 managed systems.

Complete the following steps to add a new managed system to IBM Director Console:

1. In the IBM Director Console, click **Console** → **New** → **Managed Objects** → **Systems**.
2. In the Add Systems window, specify the **System Name**, **Network Protocol** type, and the **Network Address**.
3. Click **OK**.

Discovering systems that use a mirrored image

If you have systems that are cloned or use a mirrored image, you must ensure that for Level-1 and Level-2 managed systems IBM Director has a unique identifier (UID) on each system. The SSH host key must also be unique for the mirrored systems, whether they are Level-0, Level-1, or Level-2 managed systems.

To discover cloned systems, complete the following steps:

1. If Level 1: IBM Director Core Services or Level 2: IBM Director Agent is installed on the mirrored system, delete the UID entry that was duplicated.

Platform	Location
AIX, i5/OS, Linux	/etc/ibm/director/twgagent/twgagent.uid

Platform	Location
Windows	Registry key: HKEY_LOCAL_MACHINE\SYSTEM\ CurrentControlSet\Control\Computer Name\ComputerName key Value: TWGMachineID

2. If you deleted the UID in the previous step, generate a new UID.

Platform	Instructions
AIX and Linux	Run the following command: `/opt/ibm/director/bin/genuid`

i5/OS	Restart IBM Director. In Qshell, complete the following commands:
	`/opt/ibm/director/bin/twgstop`
	`/opt/ibm/director/bin/twgstart`

Windows	From C:\Program Files\IBM\Director\bin, run GENUID.exe.

3. Regenerate the SSH host key.

Platform	Instructions

AIX and Linux	Run the following commands:
	`/etc/ssh/ssh_host_rsa_key`
	`/etc/ssh/ssh_host_dsa_key`
	Then, restart the SSH service.

i5/OS	Not Applicable
Windows	

4. In the IBM Director Console, click **Tasks** → **Discover Systems** → *System type* .

Adding a physical platform manually

You can manually define new physical platforms in IBM Director.

To add a new managed system to IBM Director Console, complete the following steps.

1. In the IBM Director Console, click **Console** → **New** → **Managed Objects** → **Physical Platforms**.
2. In the Add Physical Platform window, type the **Name** and **IP Address** for the new physical platform.
3. Click **OK**.

Adding a new SNMP device

This topic describes how to add a new SNMP device in IBM Director.

Complete the following steps to add a new SNMP device:

1. In the IBM Director Console, click **Console** → **New** → **Managed Objects** → **SNMP Devices**.

2. In the Add SNMP Devices window, click **SNMP Version** to select SNMPv1, SNMPv2c, or SNMPv3.
3. If you selected SNMPv1 or SNMPv2c, type the network address and the community name. If you selected SNMPv3, type the network address and select the profile name.
4. If you want to use this device address as an initial address for discovering additional SNMP devices, select the **Use as a discovery seed** check box.
5. Click **OK** to add the SNMP device to the Group Contents pane.

Adding a z/VM system manually

You can manually define new a new z/VM system for IBM Director to manage.

To add a new z/VM system to IBM Director Console, complete the following steps.

1. In the IBM Director Console, click **Console** → **New** → **Managed Objects** → **z/VM Systems**.
2. In the Add z/VM Systems window, complete the information for the new system:
   a. In the **IP Address** field, type the IP Address of the z/VM Management Agent that is running on the z/VM manageability access point (MAP).
   b. In the **Port Number** field, type the port for the z/VM Management Agent. The default value is 5989.
   c. In the **System Name** field, type the name of the z/VM system that you are adding.
   d. In the **Computer System Name** field, type the specification for the hardware that runs the z/VM system in the format: csname.LPAR_name, where csname represents the zSeries hardware and LPAR_name represents the logical partition that the z/VM system is running in. If you do not know the **Computer System Name**, then you can use the default value. The default value is replaced by the presence check after the managed object is unlocked.
   e. In the **Version** field, type the version of IBM Director that the z/VM system is running. If you do not know the version, use the default value. The default value is replaced by the presence check after the managed object is unlocked.
   f. In the **z/VM Management Profile Version** field, type the version of the z/VM Management Profile. You can accept the default value 1.1.0.

Discovering a BladeCenter chassis

This topic provides information about discovering a BladeCenter chassis. If the BladeCenter chassis is located on the same subnet as the management server, IBM Director can discover the BladeCenter automatically. If the BladeCenter chassis is located on a different subnet, you must discover the BladeCenter chassis manually.

IBM Director Server uses Service Location Protocol (SLP) to communicate out-of-band with the BladeCenter chassis. This communication occurs through the external Ethernet port on the BladeCenter management module. When the BladeCenter management module first is started, the management module attempts to acquire an IP address for the external management port using Dynamic Host Configuration Protocol (DHCP). If this attempt fails, the BladeCenter management module assigns an IP address (192.168.70.125) to the external management port.

Note: If you do not use a DHCP server to assign a temporary IP address to the BladeCenter chassis, introduce only *one* BladeCenter chassis onto the network at a time. IBM Director must discover and configure the chassis before another chassis is added to the LAN. Otherwise, an IP address conflict will occur.

Discovering a BladeCenter chassis automatically

This topic describes how to discover a BladeCenter chassis automatically.

The management server and the BladeCenter chassis must be connected to the network and on the same subnet. One of the following conditions must be true also:

- The network contains a DHCP server that has assigned an IP address to the management module.
- The default IP address of the management module has been changed to a valid IP address on the same subnet as the management server.

Complete the following steps to discover a BladeCenter chassis automatically:

1. Start IBM Director Console.
2. Click **Tasks** → **Discover** → **BladeCenter Chassis**. The discovery operation begins. When it is completed, the BladeCenter chassis managed object is displayed in the Group Contents pane.

   Note: The discovery operation might take several minutes, depending on the number of blade servers, management modules, and switch modules that are installed in the BladeCenter chassis.

Adding a BladeCenter chassis manually

This topic describes how to add a BladeCenter chassis manually.

Complete the following steps to add the BladeCenter chassis manually:

1. If the IP address of the management module is set to the default, manually change it.
2. From IBM Director Console, right-click in the Group Contents pane; then click **New** → **BladeCenter Chassis**.
3. In the Add BladeCenter Chassis window, in the **Chassis Name** field, type a name to identify the chassis. This name is displayed in the Groups pane of IBM Director Console.
4. In the **Network Address** field, type the IP address of the external port of the BladeCenter management module.
5. In the **User ID** field, type a valid user ID for the management module.
6. In the **Password** field, type the password that corresponds to the user ID that you typed in step 5.
7. Click **OK**. The BladeCenter chassis managed object is created. It is displayed in the Groups pane of IBM Director Console.

Discovering SMI-S storage devices

IBM Director supports storage managed objects, which represent storage-related devices that comply with the Storage Management Initiative Specification (SMI-S), including IBM System Storage DS family of storage systems.

Complete the following steps to discover SMI-S storage devices automatically:

1. Open IBM Director Console.

2. Click **Tasks** → **Discover** → **SMI-S Storage Devices**. The discovery operation begins. When it is completed, SMI-S storage devices are displayed in the Group Contents pane.

Encrypting interprocess communication

Encryption is enabled for interprocess communication by default using Advanced Encryption Standard (AES). You can change the encryption algorithm, disable and enable encryption, and manage encryption keys.

Changing the encryption algorithm

This topic describes how to change the algorithm used to encrypt communications between IBM Director Server and IBM Director Agent. When you change the algorithm, new encryption keys are sent to all managed systems. Encryption must be enabled previously.

1. From IBM Director Console, click **Options** → **Encryption Administration**.
2. In the **Encryption Suites** group box of the Encryption Administration window, select the encryption algorithm that you want to use.
3. Click **OK**.
4. In the Warning window, click **Yes**. The keys are sent to all managed systems.
5. In the Encryption Administration Success window, click **OK** to confirm that the encryption algorithm is changed and the keys are being synchronized.

Disabling encryption

This topic describes how to disable encryption on the management server. Encryption must be enabled previously.

1. From IBM Director Console, click **Options** → **Encryption Administration**. The Encryption Administration window opens.
2. Clear the **Enable encryption of data using** check box.
3. Click **OK**. The Warning window opens.
4. Click **Yes**. The Encryption Administration Success window opens.
5. Click **OK**.

Enabling encryption

This topic describes how to enable encryption on the management server.

1. From IBM Director Console, click **Options** → **Encryption Administration**.
2. In the Encryption Administration window, select the **Enable encryption of data** using check box.
3. Click **OK**.
4. In the Warning window, click **Yes**. The keys are sent to all managed systems.
5. In the Encryption Administration Success window, click **OK** to confirm that the encryption algorithm is changed and the keys are being synchronized.

Creating a new encryption key

This topic describes how a unique encryption key is generated for a managed server.

Encryption must be enabled.

1. In the IBM Director Console window, right-click the managed system and select **Reset Encryption Keys**.

2. IBM Director Server generates a new, unique key for the managed system.

Resending the encryption key to managed systems

This topic describes how to send the existing encryption keys to the managed systems.

1. From IBM Director Console, click **Options** → **Encryption Administration**.
2. In the **Resend encryption keys** group box in the Encryption Administration window, click **Resend**. When the operation is completed, the Encryption Administration Success window opens.
3. Click **OK**.
4. Close the Encryption Administration window.

Managing users

You can assign and edit privileges for users and groups, including task access and group access. By default, new users have no privileges.

Note: If you want to authorize a new IBM Director Console user, you must use the tools that are provided by the operating system to add a new user ID to one of the operating-system groups.

Authorizing IBM Director users

This topic describes how to authorize IBM Director users. On Windows platforms, you can also edit access privileges for a global group.

IBM Director Console uses the operating-system user accounts for user-logon security. When a user logs in to IBM Director, the user ID and password verification process used by the operating system is used to validate the user's authority to access IBM Director.

To use IBM Director, a user must have an operating-system account on the management server or the domain. In addition, a user must meet one of the following requirements, depending on the operating system running on the management server:

For i5/OS	Member of the IBM Director Administrators or IBM Director Super Administrators group
For Linux	Member of the diradmin or dirsuper group
For Windows	One of the following criteria: • Member of the DirAdmin or DirSuper group • Administrator privileges on the management server or the domain

Users' ability to perform tasks depends on which access privileges they have been assigned in the IBM Director environment. A super user can configure a default set of privileges for the administrator group. A super user also can edit user accounts on an individual basis.

Authorizing users for i5/OS

i5/OS users must have a user profile on the management server that is running i5/OS and be registered in a function usage group.

To initially connect to a managed system running i5/OS, a user must also have a user profile on the managed system. Additionally, a security administrator must authorize these users to IBM Director Server and IBM Director Agent functions.

IBM Director running on i5/OS has a set of associated function identifiers to use for authorizing users, configuring default users, and defining a specific user under which jobs can run. Users must be registered in one of the following functions:

- IBM Director Administrators
- IBM Director Super Administrators

IBM Director is shipped with the user profile QCPMGTDIR. QCPMGTDIR has *ALLOBJ special authority as well as *SECADM special authority. QCPMGTDIR is used to start all IBM Director jobs and is the default profile under which the jobs run. You can change the default profile from QCPMGTDIR to a user profile of your choice for the following function IDs:

- IBM Director Agent default user
- IBM Director Server default user
- IBM Director Agent run as user
- IBM Director Server run as user

The following table describes the three function usage groups to which a user can be authorized.

Function ID	Purpose
IBM Director Administrators	Perform management functions using tasks to which they are authorized.
IBM Director Agent access	Initially connect IBM Director Server to an IBM Director Agent. **Note:** By default, any user with *ALLOBJ authority has access to this function.
IBM Director Agent default user	By specifying a user profile other than the default profile, remote commands can be performed on a managed system using the specified user profile. No user ID and password are required when requesting the command.
IBM Director Agent run as user	By specifying a user profile other than the default profile, jobs on the managed system are performed under this profile. To complete all IBM Director tasks successfully, the user profile must have *ALLOBJ authority.
IBM Director Server default user	Allows a user profile to be registered as the default for tasks such as file transfer, software distribution, and event actions. To complete all IBM Director tasks successfully, the user profile must have *ALLOBJ authority.
IBM Director Server run as user	By specifying a user profile other than the default profile, jobs on the management server are performed under this profile. To complete all IBM Director tasks successfully, the user profile must have *ALLOBJ and *SECADM authority.
IBM Director Super Administrators	Configure a set of privileges for the administrator group, edit user accounts on an individual basis, and use the functions of the DIRCLI client.

Prerequisite:

To authorize users to these functions, you must have *SECADM authority.

Complete the following steps to authorize users to IBM Director functions:

1. In iSeries Navigator, right-click the server and click **Application Administration**.
2. On the Application Administration dialog, click the **Host Applications** tab.
3. Expand **IBM Director for iSeries**.
4. Select the function group to which you want to add users and click **Customize**. Complete the instructions on the dialog to grant authority.

You can also use the Work Function Usage (WRKFCNUSG) command in the character-based interface, WRKFCNUSG QIBM_QDIR*.

Creating user-account defaults

This topic describes how to set the default access privileges for new members of the administrators group.

A super user can use the User Defaults Editor to set the default access privileges for new members of the administrators group.

Complete the following steps to create user-account defaults:

1. In IBM Director Console, click **Options** → **User Administration**.

   This window contains a list of all users authorized to access IBM Director.
2. In the User Administration window, click **User** → **User Defaults**.

   In the User Defaults Editor window, you can set the default access privileges for new members of the DirAdmin group.

   Notes:

   a. For increased security, consider providing no default access privileges. You will have to set access levels for each user, but you can be sure that a user will not accidentally be able to access restricted groups or tasks.
   b. You can restrict access to the Event Action Plan wizard by removing users' access to the Event Action Plan Builder task.

Editing an individual user's access privileges

This topic describes how to edit an individual user's access privileges.

Complete the following steps to edit a user's access privileges:

1. In IBM Director Console, click **Options** → **User Administration**.

   This window contains a list of all users and groups that are authorized to access IBM Director.
2. In the User Administration window, select the user whose access privileges you want to modify. Click **Actions** → **User** → **Edit**.
3. In the User Editor window, click the **Privileges** tab.
4. To add a privilege, click the privilege in the **Available Privileges** pane, and then click **Add**. To remove a privilege, click the privilege in the **Privileges Granted to User** pane, and then click **Remove**.
5. To restrict the user's access to groups, click the **Group Access** tab.
6. To permit the user to access specific groups only, select the **Limit user access only to the groups listed** check box. To add a group, click the group in the **Available Groups** pane and click **Add**. To remove a group, click the group in

the **Groups User Can Access** pane and click **Remove**. To prevent the user from creating new groups or modifying existing groups, select the **Limit user to read-only access of groups** check box.

7. To restrict the user's access to tasks, click the **Task Access** tab.
8. To restrict the user to performing certain tasks only, select the **Limit user access only to the tasks listed** check box. To add a task, click the task in the **Available Tasks** pane and click **Add**. To remove a task, click the task in the **Tasks User Can Access** pane and click **Remove**.

   Note: You can restrict access to the Event Action Plan wizard by removing the user's access to the Event Action Plan Builder task.
9. When you have finished editing the user's privileges, click **OK**.

Editing group access privileges

This topic describes how to edit access privileges for a global group on Windows platforms.

Complete the following steps to edit privileges for a group:

1. In IBM Director Console, click **Options** → **User Administration**.

   This window contains a list of all users and groups that are authorized to access IBM Director.
2. In the User Administration window, click the **Groups** tab.
3. Select the group whose access privileges you want to modify. Click **Actions** → **Group** → **Edit**.
4. In the Group Editor window, click the **Privileges** tab.
5. On the Privileges page, to add a privilege, click the privilege in the **Available Privileges** pane and then click **Add**. To remove a privilege, click the privilege in the **Privileges Granted to User** pane and then click **Remove**.
6. To restrict access for the group to IBM Director groups, click the **Group Access** tab.
7. On the Group Access page, to permit the user to access specific groups only, select the **Limit user access only to the groups listed** check box. To add a group, click the group in the **Available Groups** pane and click **Add**. To remove a group, click the group in the **Groups User Can Access** pane and click **Remove**. To prevent the user from creating new groups or modifying existing groups, select the **Limit user to read-only access of groups** check box.
8. To restrict access for the group to tasks, click the **Task Access** tab.
9. On the Task Access page, to restrict the group to performing certain tasks only, select the **Limit user access only to the tasks listed** check box. To add a task, click the task in the **Available Tasks** pane and click **Add**. To remove a task, click the task in the **Tasks User Can Access** pane and click **Remove**.
10. When you have finished editing the privileges, click **OK**.

Creating a new user ID

This topic describes how to create a new user ID in IBM Director.

If you want to authorize a new IBM Director Console user, you must use the tools that are provided by the operating system to add a new user ID to one of the operating-system groups.

Changing user defaults

This topic describes how to change the defaults for new user IDs in IBM Director.

You can change the defaults for new IBM Director user IDs. You can specify the default information for the full name, description, privileges, group access limits, and task access limits for all new user IDs.

Note: These defaults affect only members of the Diradmin group; they do not limit the attributes of members of the Dirsuper group.

Complete the following steps to change the defaults for new IBM Director user IDs:

1. In IBM Director Console, click **Options** → **User Administration**.
2. In the User Administration window, click **User** → **User defaults**.
3. In the User Defaults Editor window, make the changes. Click **OK** to save the changes.

Editing an existing IBM Director user profile

This topic describes how to make changes to an existing IBM Director user profile.

Complete the following steps to edit an existing user profile:

1. In IBM Director Console, click **Options** → **User Administration**.
2. In the User Administration window, click the row of the user.
3. Click **User** → **Edit**.
4. In the User Editor window, make the changes.
5. Click **OK** when you are finished making all changes in the window.

Managing groups

You can create customized groups of managed objects that can be displayed in the Group Contents pane. You can also change the organization of the group using associations.

Creating a criteria-based dynamic group

This topic describes how to create a criteria-based dynamic group.

Complete the following steps to create a dynamic group:

1. In the IBM Director Console window, select **Console** → **New** → **Group** → **Dynamic Group...**.
2. In the Dynamic Group Editor window, in the Available Criteria pane, expand the tree that has the criterion you want to use to define the group. Click a criterion and click **Add**.

   The default operator is the equal sign (=). You can change the operator for any criterion by right-clicking the criterion and selecting another operator.

   Repeat this step to add more criteria. When you add criteria, the Choose Add Operation window opens. Click **All True** or **Any True**; then, click **OK**.
3. Click **File** → **Save As** to save the new dynamic group.
4. In the Save As window, type a descriptive name for the group. This name will be listed in the Groups pane.

   Note: The group name is case-sensitive.
5. Click **OK**. The group is displayed under **All Groups** in the Groups pane.
6. Click **File** → **Close Group Editor** to close the Dynamic Group Editor window.

Note: You cannot use a wild card (*.*) to create a dynamic group. To create a dynamic group using criteria that are not present in the IBM Director database, you must use DIRCLI.

Creating a group category

This topic describes how to create a group category.

You can organize large numbers of groups by creating group categories. Group categories are static, although the groups that are included in a category can be dynamic or static.

Complete the following steps to create a group category:

1. In the IBM Director Console window, select **Console** → **New** → **Group Category**.
2. In the **Name** field, type a descriptive name for the group category.
3. In the Group Category Editor, select the groups or group categories on the left that you want included in the category and click **Add**.
4. Click **OK**. The new group category is displayed in the Groups pane.

Creating a static group

This topic describes how to create a static group.

Complete the following steps to create a static group:

1. In the IBM Director Console window, click **Console** → **New** → **Group** → **Static Group**.
2. In the Static Group Editor: New window, select the managed objects that you want to add to the static group and click **Add**.
3. When you are finished adding managed objects, click **File** → **Save**.
4. In the Save As window, type a descriptive name for the group. This is the group name that will be listed in the Groups pane.

   Note: The group name is case-sensitive.
5. Click **OK**. The group is displayed under **All Groups** in the Groups pane.

Creating a task-based dynamic group

This topic describes how to create a task-based dynamic group.

Complete the following steps to create a task-based dynamic group:

1. In the IBM Director Console window, click **Console** → **New** → **Group** → **Task Based Group...**.
2. In the Task Based Group Editor window, in the Available Tasks pane, click a task you want to perform using this group; then, click **Add**.
3. When you are finished adding tasks, click **File** → **Save As**.
4. In the Save As window, type a descriptive name for the group. This name will be listed in the Groups pane.

   Note: The group name is case-sensitive.
5. Click **OK**. The group is displayed under **All Groups** in the Groups pane.
6. Click **File** → **Close Group Editor** to close the Task Based Group Editor window.

Exporting a dynamic group

This topic describes how to export a dynamic group. Exporting a group enables you to archive or back up the contents of the group.

Complete the following steps to export a group:

1. Right-click the **Groups** pane and click **Export Group**.
2. In the Group Export window, click the group that you want to export from the groups that are available for export.
3. Type a file name in the **Export Destination File** field, or click **Browse** to locate a file name.
4. Click **Export**. The group is exported to the file that you specified.

Importing a dynamic group

This topic describes how to import a dynamic group. Importing a previously exported group enables you to distribute a selected group to a new instance of IBM Director Server.

Complete the following steps to import a dynamic group:

1. Right-click the **Groups** pane and click **Import Group**.
2. In the Group Import window, select the group that you want to import by navigating the tree structure or typing the group name in the **File Name** field.
3. Click **OK**.
4. In the Group Import window, click one or more groups in the **Import Group Detail** pane.
5. Click **Function** and click the applicable action.
6. Click **Import** → **Import Selected Groups**. The group or groups are added, updated, or skipped.

Changing the organization of a group

This topic describes how you can use associations to change the organization of a group of managed objects that is displayed in the Managed Objects pane.

To display group contents according to an association:

1. Click **Associations**.
2. Select the radio button for the association type that you want to use.
3. **Optional:** Select one or more association options from the bottom half of the **Associations** menu.

The newly assigned association type for that group will persist the next time you display that group in the Group Contents pane.

Managing auditing

With the new auditing capabilities, you can track what takes place on IBM Director Server.

Audit records

The audit records provide information that is necessary to track basic, important, high-level actions, including the following data:

- Security

  - SSL turned on
  - SSL turned off
  - Encryption turned on (not audited on IBM Director Servers running on Windows due to a technical limitation)
  - Encryption turned off (not audited on IBM Director Servers running on Windows due to a technical limitation)
  - Request Access to Agent successful
  - Request Access to Agent unsuccessful
  - Director User Administration privileges altered
  - Successful login
  - Successful logout
  - Unsuccessful login
  - Auditing settings changed
- Task Activation or Deactivation
  - Task activated
  - Task deactivated
  - Task scheduled
- User Administration
  - Create a user or group on an agent system
  - Delete a user or group on an agent system
  - Modify a user or group on an agent system
  - Add or remove users from a group on an agent system
- File Read—audited only if the file read is on IBM Director Server and IBM Director Server is not running on Windows (due to technical limitations)
  - View a file
- File Write—only audited if the file written is on IBM Director Server and IBM Director Server is not running on Windows (due to technical limitations)
  - Add or Change a file—audited as File Transfer to IBM Director Server
  - Rename a file
  - Delete a file
  - Transfer a file
  - Rename a directory
  - Delete a directory
  - Add a directory
- Remote Access
  - Start a remote session
  - End a remote session
  - Start a remote control session
  - End a remote control session
- Configuration Changes
  - Change Network configuration (such as TCP/IP addresses and DNS servers)
  - Restart due to network configuration changes
  - Change SNMP configuration
  - Change Asset information
  - Change ASF configuration
  - Mass Config profile added, removed, or changed on a system or group
- Remote Command Execution
  - Run a remote command using a Process Task
- CIM
  - Create Instance (only selected Create Instance CIM calls are audited)
  - Modify Instance (only selected Modify Instance CIM calls are audited)
  - Delete Instance (only selected Delete Instance CIM calls are audited)
  - Invoke Method (only selected Invoke Method CIM calls are audited)
  - Set Property (only selected Set Property CIM calls are audited)

- Command Line Interface
  - Command was executed successfully
  - Command was not executed successfully

Enabling and disabling auditing

IBM Director Server Auditing Administration allows you to enable or disable auditing for IBM Director Server.

Complete the following steps to enable or disable auditing for IBM Director Server.

1. From IBM Director Console select **Options** → **Auditing Administration**.
2. On the IBM Director Server Auditing Administration window, select the Enable auditing check box. If you want to disable auditing, clear the check box.

   Note: If auditing is being enabled for the first time, all of the categories are selected for auditing by default.
3. Complete the fields on the window, and click **OK**.

Changing auditing settings

IBM Director Server Auditing administration allows you to select what information is audited as well as how the audit information is stored.

To change the auditing settings, complete the following steps:

1. From IBM Director Console select **Options** → **Auditing Administration**.
2. On the IBM Director Server Auditing Administration window, you can select which categories of information are audited by using the arrows to move selected categories between the **Available categories** list and the **Selected categories** list.
3. Specify the names of the audit files, the number of files to use for audit logs, and the maximum size for each file.

Deleting the audit log

IBM Director Server Auditing Administration allows you to delete the audit logs.

Attention: Audit files are not saved or archived automatically. Deleting the audit log is a task that cannot be undone.

Complete the following steps to delete the audit log:

1. From IBM Director Console select **Options** → **Auditing Administration.**.
2. On the IBM Director Server Auditing Administration window, click **Delete audit logs**.

Viewing the audit log

IBM Director Server Auditing Administration allows you to view the audit log, if you have enabled auditing for IBM Director Server.

Complete the following steps to view the audit log when auditing is enabled:

1. From IBM Director Console select **Options** → **Auditing Administration...**.
2. On the IBM Director Server Auditing Administration window, click **View audit log**. **View audit log** allows you to see the audit log results quickly, but the audit log is not formatted. The audit log is in CSV format; import the log to another tool for viewing and searching.

Note: If you have chosen to use more than one file for the audit log, you will only see the results for the most recently created log file.

Using mass configuration

You can use mass-configuration profiles to configure or run a single task on a group of managed objects. You can use Mass Configuration with the following tasks:

- Configure Alert Standard Format
- Asset ID
- Network Configuration
- Configure SNMP Agent

Creating a profile

This topic describes how to create a profile in IBM Director.

To use Mass Configuration, you must create a profile. The following procedure uses the Configure Alert Standard Format task as an example.

Complete the following steps to create a profile:

1. In IBM Director Console Tasks pane, right-click the **Configure Alert Standard Format** task and click **Profile Builder**.
2. In the Configure Alert Standard Format: Profile Builder window, click **New Profile**
3. In the Input window, type the new profile name in the field and click **OK**. The new profile name displays in the field in the upper left of the Configure Alert Standard Format: Profile Builder window.
4. In the right pane of the Profile Builder window, edit the information as applicable.
5. To permit an IBM Director user to modify the settings that you will apply with the profile, ensure that the **Enable Changes** check box in the notebook window is selected. By default, this check box is selected.

   Important: If you choose to clear the **Enable Changes** check box, you will be unable to modify the settings through IBM Director or by any other operating system method.
6. Click **Save Profile**, then **Yes**, to save the profile.
7. Click **File** → **Close** to close the Profile Builder window.

Applying a profile to a group

This topic describes how to apply a profile to a managed object or a group in IBM Director.

Profiles are saved in the IBM Director Console Tasks pane underneath the task with which they are associated. You can apply a profile to an individual managed object or a group.

Complete the following steps to apply a profile to a managed object or a group:

1. Expand the **Configure Alert Standard Format** task to display the task profiles.
2. Drag a profile onto a managed object or a group.
3. In the Status window, click **Close**.

Managing profiles

You can edit groups associated with a profile or delete the profile using the Profile Manager window.

1. Expand the **Configure Alert Standard Format** task to display the task profiles.
2. Right-click a profile and click **Profile Manager**.
3. In the Status window, to remove a profile, click the profile in the **Profile** field; then, click **Remove Profile**.
4. To remove a group from the profile, click the profile in the **Profile** list and click the group in the **Group** list; then, click **Remove Group**.
5. To view the status of the profile, click **Status**. The **Profile Status** field is displayed.
6. Click **Close** to close the Profile Status field and return to the Status window.
7. Click **X** in the right of the window bar to close the Status window.

Viewing license information

You can use license administration to view the total number of product licenses, the number of used licenses, and the number of available licenses.

To view license information, complete the following steps:

1. On the IBM Director Console click **Options** > **License Administration**.
2. View your license information.

Working with security states

You can utilize the security features of IBM Director to manage access to your managed systems.

Accessing a secured managed system

This topic describes how to access a secured managed system.

If a managed system is secure but the management server to which you are connected does not have authorization to access it, the managed system is displayed in the Group Contents pane of IBM Director Console with a padlock icon beside it.

Complete the following steps to access a secure managed system from an unauthorized management server:

1. In IBM Director Console, right-click the managed system to which you do not have access.
2. Click **Request Access**.
3. In the Request Access to Systems window, type the user ID and password of a user with administrator privileges on the managed system.
4. Click **OK**.

Automatically securing unsecured systems

This topic describes how to configure IBM Director to secure unsecured systems automatically.

Complete the following steps to configure IBM Director to secure unsecured systems automatically:

1. In IBM Director Console click **Options** → **Discovery Preferences**. The Discovery Preferences window opens and the IBM Director Systems page is displayed.
2. Select the **Automatically secure unsecured systems** check box.
3. Click **OK**.

   Note: When the **Automatically secure unsecured systems** feature is enabled, IBM Director Server prevents future management servers from managing the managed systems without first requesting access.

Removing access to a managed system manually

This topic describes how to remove access to a managed system manually. You can do this by removing the public key for the management server from the managed system.

Complete the following steps to revoke the ability of a management server to access a managed system:

1. From the managed system, change to the directory where the security information is stored. This is one of the following directories:

Operating system	Directory
Linux	/opt/ibm/director/data
i5/OS	/QIBM/UserData/Director/data
NetWare	*d*:\IBM\Director
Windows	*d*:\Program Files\IBM\Director\Data

   where *d* is the drive letter of the hard disk on which IBM Director is installed and IBM Director is installed in the default location.
2. Using a text editor, view each dsa*.pub file. The first characters in a dsa*.pub file are of the form DSA*xxxx*, where *xxxx* is the name of the management server.
3. Locate the dsa*.pub file for the management server that you want to unauthorize, and delete it.
4. To stop IBM Director Agent, from a command prompt, type one of the following commands and press Enter:

For i5/OS	`/qibm/userdata/director/bin/twgend`
For Linux	`/opt/ibm/director/twgstop`
For NetWare	`unload twgipc`
For Windows	`net stop twgipc`

5. To restart IBM Director Agent, type one of the following commands and press Enter:

For i5/OS	`/qibm/userdata/director/bin/twgstart`
For Linux	`/opt/ibm/director/twgstart`
For NetWare	`load twgipc`
For Windows	`net start twgipc`

After IBM Director Agent starts, the management server whose public key you removed is no longer able to access the managed system.

Securing a managed system manually

This topic describes how to secure a managed system manually.

Use this procedure in the following situations:

- You suspect that a rogue management server was introduced into an IBM Director environment before all managed systems were secured, and you want to resolve any possible security risks.
- You want to establish trust relationships between a managed system and multiple management servers.

You can use this procedure to secure either an unsecured or secured system. You also can automate this procedure by using logon scripts or other automated execution mechanisms.

Complete the following steps to secure a managed system manually:

1. If you have not done so already, install and start IBM Director Server. IBM Director Server creates a dsa*.pub and dsa*.pvt file, as well as a secin.ini file set to secure.

   Note: The secin.ini file only exists for Windows platforms.
2. Copy the dsa*.pub and secin.ini files to a file server or other accessible location.

   Note: If you want to authorize more than one IBM Director Server to manage a system, copy the dsa*.pub files from each. Only one copy of secin.ini is necessary.
3. If IBM Director Agent installed on the managed system has not been started yet, go to step 5. Otherwise, stop IBM Director Agent. From a command prompt, type the following command and press Enter:

Operating system	Command
i5/OS	`/qibm/userdata/director/bin/twgend`
Linux	`/opt/IBM/director/twgstop`
NetWare	`unload twgipc`
Windows	`net stop twgipc`

4. Delete all existing dsa*.pub files from the managed system.
5. Place the dsa*.pub and secin.ini files (that you copied in step 2) into one of the following directories:

Operating system	Directory
i5/OS	`/QIBM/UserData/Director/data`
Linux	`/opt/ibm/director/data`
NetWare	*c*`:\IBM\Director`
Windows	*c*`:\Program Files\IBM\director\data`

   c is the hard disk where IBM Director Agent is installed, and IBM Director Agent is installed in the default directory.
6. To restart IBM Director Agent, type one of the following commands and press Enter:

Operating system	Command
i5/OS	`/qibm/userdata/director /bin/twstart`
Linux	`/opt/IBM/director/twgstart`
NetWare	`load twgipc`
Windows	`net start twgipc`

After IBM Director Agent starts, the managed system is secure; it permits *only* authorized IBM Director Servers (that is, the ones whose dsa*.pub file you copied to the managed system) to manage it.

Viewing alerts in the Message Browser

You can use the Message Browser to view events (alerts) that are sent to IBM Director Console.

The Message Browser is displayed automatically whenever an alert is sent to the management console. You can opt to be notified in this manner when an event occurs by configuring an event action plan with the Send an Event Message to a Console User event action.

The Message Browser displays all alerts, including management console ticker-tape alerts. However, the Message Browser does not display any ticker-tape messages. (A ticker-tape message can display, for example, resource-monitor data.)

You can start the Message Browser to view all active messages that are received and clear any previous messages. To start the Message Browser, click **Tasks** → **Message Browser**.

Enabling the Wake on LAN feature

If your server supports the Wake on LAN feature, you can enable it after IBM Director is installed. See your server documentation to determine whether or not your server supports this feature.

Enabling the Wake on LAN feature for Linux or AIX

This topic describes how to enable the Wake on LAN feature for IBM Director Agent.

Complete the following steps to enable Wake on LAN for IBM Director Agent:

1. To stop IBM Director Agent, from a command prompt, type the following command and press **Enter**:
   ```
   /opt/ibm/director/bin/twgstop
   ```
2. Open an ASCII text editor and edit the ServiceNodeLocal.properties file. This file is in the /opt/ibm/director/data directory.
3. Modify the value of ipc.wakeonlan to read as follows:
   ```
   ipc.wakeonlan=1
   ```
4. Save and close the ServiceNodeLocal.properties file.
5. To start IBM Director Agent, type the following command and press **Enter**:
   ```
   /opt/ibm/director/bin/twgstart
   ```

Enabling the Wake on LAN feature on Windows

This topic describes how to enable Wake on LAN on a managed system running windows.

Note: To determine whether your server supports the Wake on LAN feature, see your server documentation.

Complete the following steps to enable Wake on LAN:

1. Click **Start Settings > Control Panel**. The ″Control Panel″ window opens.
2. Double-click **Add/Remove Programs**. The ″Add/Remove Programs″ window opens.
3. Click the IBM Director software component that you want to modify; then, click **Change**. The InstallShield wizard starts, and the ″Welcome to the InstallShield Wizard″ window opens.
4. Click **Next**. The ″Program Maintenance″ window opens.

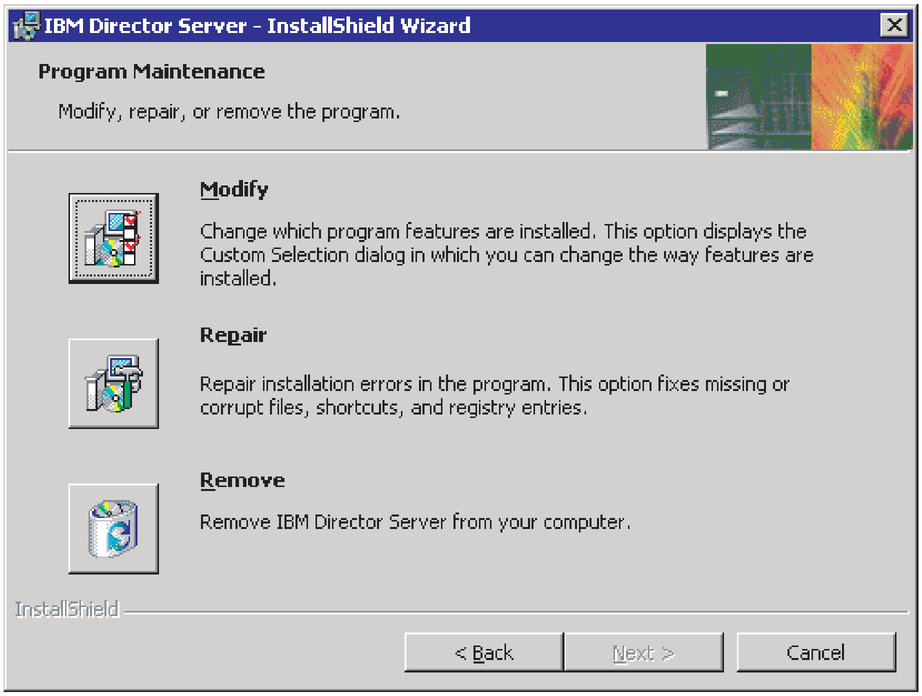

IBM Director Server - InstallShield Wizard

Program Maintenance

Modify, repair, or remove the program.

Modify

Change which program features are installed. This option displays the Custom Selection dialog in which you can change the way features are installed.

Repair

Repair installation errors in the program. This option fixes missing or corrupt files, shortcuts, and registry entries.

Remove

Remove IBM Director Server from your computer.

Figure 11. ″Program Maintenance″ window

5. Click **Modify**; then, click **Next**.
6. Continue through the wizard until you reach the ″Network driver configuration″ window.
7. Select the Enable Wake on LAN check box.
8. Complete the wizard.

Configuring the database on the management server

This topic describes how to configure the database after IBM Director Server is installed.

Configuring the database on Linux or AIX

This section describes how to configure your database application for use with IBM Director after IBM Director Server is installed. These instructions are applicable for all Linux platforms and AIX.

Configuring the database on a Linux or AIX management server using the cfgdb command

This topic describes how to configure the database after IBM Director Server is installed using the cfgdb command.

Note: The cfgdb command must be used in a graphical environment. To configure the database from a command line, see "Configuring the database on a Linux or AIX management server using the cfgdbcmd command."

Complete the following steps to install and configure a database after you have installed IBM Director Server:

1. To stop IBM Director Server, from a command prompt, type the following command and press **Enter**:
   ```
   /opt/ibm/director/bin/twgstop
   ```
2. Type the following command and press **Enter**:
   ```
   /opt/ibm/director/bin/cfgdb
   ```
3. Follow the instructions on the screen.
4. To restart IBM Director Server, type following command and press **Enter**:
   ```
   /opt/ibm/director/bin/twgstart
   ```

Configuring the database on a Linux or AIX management server using the cfgdbcmd command

This topic describes how to configure the database from a command line after IBM Director Server is installed.

Complete the following steps to install and configure a database from the command line after you have installed IBM Director Server:

1. Open the cfgdbcmd.rsp file in an ASCII text editor and modify the settings. The cfgdbcmd.rsp file is located in the /opt/ibm/director/data directory and is fully commented.
2. Save the modified response file with a new file name.
3. To stop IBM Director Server, from a command prompt, type the following command and press **Enter**:
   ```
   /opt/ibm/director/bin/twgstop
   ```
4. Type the following command and press **Enter**:
   ```
   cfgdbcmd -rspfile response.rsp
   ```
   where *response.rsp* is the name of the response file as saved in step 2.
5. When the configuration is completed, restart IBM Director Server. Type following command and press **Enter**:
   ```
   /opt/ibm/director/bin/twgstart
   ```

Configuring the database on a Windows management server

This topic describes how to configure the database after IBM Director Server is installed.

Complete the following steps to configure a database after you have installed IBM Director Server:

1. Stop IBM Director Server. From a command prompt, type the following command and press Enter:

   ```
   net stop twgipc
   ```
2. Type the following command and press Enter:

   ```
   cfgdb
   ```

   The "IBM Director database configuration" window opens.
3. Follow the instructions on the screen.

Note: For the Database Configuration window to accept Windows color settings after installation, you must first go to the **Accessibility Preferences** tab in the **Console Preferences** window, and select the Windows option in the **Colors** field.

Enabling SNMP access and trap forwarding for Linux

This topic describes how to enable SNMP access and trap forwarding for Linux.

To enable SNMP access and trap forwarding on Linux, you must be running one of the following operating systems:

- i5/OS, Version 5 Release 3
- Red Hat Linux, version 3.0
- Red Hat Linux, version 4.0
- SUSE LINUX Enterprise Server 8 for x86
- SUSE LINUX Enterprise Server 9 for x86
- Red Hat Enterprise Linux AS, version 4.0, for IBM System z9, zSeries and S/390
- SUSE LINUX Enterprise Server 9 for IBM System z9, zSeries and S/390

Note: IBM Director 5.10 supports SNMP access and trap forwarding using NetSNMP version 5.2.1 (all releases). You must remove all NetSNMP libraries older than version 5.2.1 from your system before completing the steps in this task. Ensure that no net-snmp, net-snmp-libs or net-snmp-devel rpms older than version 5.2.1 remain on your system.

Complete the following steps to enable SNMP access and trap forwarding for managed systems running Linux:

1. Download the net-snmp-5.2.1.tar.gz file from the Net-SNMP Web site at http://www.net-snmp.org/download.html.

   Note: Net-SNMP is not supported on VMware console operating systems.
2. Build and install Net-SNMP. Refer to the INSTALL and README files included in the net-snmp-5.2.1.tar.gz package for instructions. If you are not running an AMD64 or EM64T distribution, go to step 4 on page 119.
3. (AMD64 and EM64T distributions only) SNMP functionality on these systems requires 32-bit versions of the NetSNMP libraries for use by the IBM Director SNMP SubAgent. The NetSNMP Master Agent can use the native 64-bit libraries.

   You can request the 32-bit libraries from your Linux distribution provider or compile them yourself. If you compile the libraries yourself, it is strongly recommended that you run the configuration using the `––without-rpm` command option.

Compile the libraries in one of the following ways:

- Using 32-bit compiler flags (not tested with NetSNMP).
- Compile on a 32-bit version of the Linux distribution and then move the libraries to your 64-bit system.

Copy the 32-bit libraries to the ibm/director/cimom/lib directory or to another location on your system library path.

Important: If you are using a 64-bit installation of the NetSNMP Master Agent, ensure that the 32-bit libraries do not interfere with the 64-bit libraries used by the Master Agent. You must only include one library in the library search path.

4. Configure the Net-SNMP agent using one of the following methods:

   Note: The Agent Operating mode must be set to AgentX Master Agent.

   - Use the **snmpconf** utility to change the Agent Operating mode to AgentX Master Agent and to configure Net-SNMP for access groups and trap destinations.
   - To configure Net-SNMP manually, open the snmpd.conf file in an ASCII text editor and locate the following section of text:

     ```
     # master: Should the agent operate as a master agent or not.
     # Currently, the only supported master agent type for this token
     # is "agentx".
     #
     # arguments: (on|yes|agentx|all|off|no)
     master agentx
     ```

     Uncomment the `master agentx` string by removing the hash mark (#) if necessary and save the modified file.

5. Ensure that the LD_LIBRARY_PATH environment variables in the dacimlist and dasnmp start-up scripts include the following path information:

   - /usr/local/lib
   - /opt/ibm/director/lib
   - /opt/ibm/director/cimom/lib

   You can configure this information in either of the following ways:

   - Create a *filename*.sh file in the /etc/profile.d/ directory, where *filename* is a name of your choice. Include the following text in the file (all on one line):

     ```
     export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib:/opt/ibm/director/lib:
      /opt/ibm/director/cimom/lib
     ```

   - Create a *filename*.conf file in the /etc/ld.so.conf.d/ directory, where *filename* is a name of your choice. Include the following text in the file

     ```
     /usr/local/lib
     /opt/ibm/director/lib
     /opt/ibm/director/cimom/lib
     ```

6. Ensure that the SNMPCONFPATH environment variables in the dacimlist and dasnmp start-up scripts include the path for the location of the snmpd.conf file. You can configure this variable by creating a *filename*.conf file in the /etc/profile.d/ directory, where *filename* is a name of your choice. Include the following text in the file:

   ```
   export SNMPCONFPATH=$SNMPCONFPATH:/path
   ```

   where *path* is the path for the snmpd.conf file.

Chapter 3. Managing systems

You can use IBM Director to manage your systems by performing the following tasks.

Accessing the IBM Director interface

You can access IBM Director functions using the IBM Director Console graphical user interface (GUI) or the IBM Director command-line interface.

Starting IBM Director Console

This topic describes how to start IBM Director Console.

Complete the following steps to start IBM Director Console:

1. Perform one of the following actions:

Operating system	Instructions
For Linux	From a command prompt, type the following command and press **Enter**: `twgcon`
For Windows	Click **Start** → **Programs** → **IBM Director Console**.

2. In the IBM Director Login window, type the name of the management server in the **IBM Director Server** field.
3. In the **User ID** field, type one of the following strings:

Operating system	Instructions
For Linux	*UserID*
For Windows	*ComputerName\UserID*

   where:
   - *UserID* is a valid IBM Director user ID.
   - *ComputerName* is either the local computer name or the domain, depending on whether the IBM Director service account has domain or local privileges
4. In the **Password** field, type the password that corresponds to the user ID.
5. If you want to use Secure Socket Layer (SSL) to encrypt data communications between IBM Director Console and IBM Director Server, select the **Use SSL** check box.

   Note: SSL console connections must be enabled. This is the default option.
6. Click **OK** to start IBM Director Console.

Starting the IBM Director command-line interface

Use the command-line interface to directly access many IBM Director functions.

The IBM Director command-line interface (CLI) is an important primary interface into IBM Director which may be used either as an efficient way to accomplish simple tasks directly or as an embeddable and scriptable framework for achieving higher level goals.

Note: The IBM Director CLI provides support in version 5.10 for a subset of the commands previously available through the deprecated dircmd executable.

Complete the following steps to access the IBM Director command-line interface:

1. Log in to the management server as an IBM Director super user. Log in directly, or remotely via telnet, SSH, or RSH. Command line access is solely limited to IBM Director super users, and the IBM Director command-line interface executable (dircli) is only installed with IBM Director Server.
2. Type the command you want to execute.

   The general command syntax is described below:

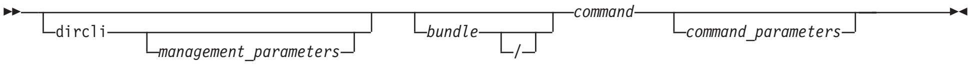

dircli
The **dircli** keyword is optional if both of the following are true:
- the command is being run on a UNIX or Linux platform
- the command is unique on that platform

When these conditions are met, symbolic links are defined on installation of IBM Director Server which allow users to omit the **dircli** keyword.

management_parameters
management_parameters are principally retained for syntax compatibility with command-line execution through the dircmd executable which was implemented in earlier versions of IBM Director.

bundle
bundle specifies a command bundle, and is typically used with the retained commands formerly available through dircmd, or for extensions to the CLI.

command
command is one of the dircli commands, or one of the retained commands formerly available through dircmd. Specific syntax for each command is included in the reference topic for the command.

command_parameters
command_parameters are specific to the command being performed.

CLI command outputs are sent to stdout. Errors are sent to stderr.

Every CLI command that is executed is audited with the following information:
- Executing user
- Time stamp
- Command
- Command arguments
- Return code
- How long the command took to execute

The log is stored as ASCII text in the *IBM_Director_Base_Directory*\log\CLI\Log.txt file. Every time the IBM Director Server is restarted the Log.txt file is backed up to log.bk and a new log is created.

Examples

The following examples all execute the **accessmo** command on a UNIX platform:

```
dircli accessmo -n webserver
dircli managedobject/accessmo -n webserver
dircli managedobject accessmo -n webserver
accessmo -n webserver
```

Related reference

"IBM Director CLI (dircli)" on page 531
This topic lists the commands available in the IBM Director command-line interface, **dircli**.

Launching tasks in IBM Director

The IBM Director interface allows you the flexibility to start tasks using a variety of methods.

You can launch tasks in IBM Director from the task menu, from the toolbar, from the context menu of a managed object, a group, or a task, and by dragging a task to a target object, or by double-clicking a task.

Task menu

Complete the following steps to launch tasks from the Task menu:

1. In the Groups Content pane of IBM Director Console, select a target object for the task, if applicable.
2. In the IBM Director Console window, click **Tasks** → *Task name*. *Task name* is a variable that represents that task you want to start.

Toolbar

Complete the following steps to launch tasks from the toolbar:

1. In the Groups Content pane of IBM Director Console, select a target object for the task, if applicable.
2. Near the top of the IBM Director Console window, below the menu bar, locate the icon on the toolbar that represents the task you want to complete.
3. Click the menu arrow to view the options available for the task, and select the task you want to complete.

Context menu

Complete the following steps to launch tasks from the context menu of a managed object, a task or a group:

1. In the IBM Director Console window, right-click a managed object, a group, or a task.
2. In the context menu that appears, click the task you want to complete.

Drag and drop

Certain tasks that you perform, such as scheduling a process task, support dragging the task onto a managed object or group to launch the task.

1. In the Tasks pane of IBM Director Console locate the noninteractive task you want to perform.
2. If the Tasks pane does not appear, click the triangle along the right side of IBM Director Console, or from the menu bar click **View** → **Tasks Pane**.

3. Drag the task from the Tasks pane to the target object or group in the Groups Content pane or the Groups pane.

Double-click a task

1. In the Tasks pane of IBM Director Console locate the task you want to perform.
2. If the Tasks pane does not appear, click the triangle along the right side of IBM Director Console, or from the menu bar click **View** → **Tasks Pane**.
3. Double-click the task you want to complete.

Launching the command-line interface for a BladeCenter chassis or service processor

This topic describes how to launch the command-line interface for a BladeCenter chassis or a service processor from the IBM Director Console.

To launch the command-line interface, follow these steps:

1. Open **IBM Director Console**.
2. In the Group Contents pane, right click a managed object.
3. If the managed object is a BladeCenter chassis, click **BladeCenter Management** → **Launch Command Line Interface**. If the managed object is a physical platform for a service processor, click **Management Processor Assistant** → **Launch Command Line Interface**.

For more information about the IBM Management Processor command-line interface, see *IBM Management Processor Command-Line Interface User's Guide*.

Managing Hardware Management Console

IBM Director allows you to discover a Hardware Management Console (HMC) for IBM eServer i5 and eServer p5 models. You can view the servers and hardware resources that are managed by the HMC, collect inventory information, perform power control, and launch management tools.

Accessing the Hardware Management Console

When an HMC is discovered initially, it is in a secured state. You must request access to unlock it.

To access an HMC after its initial discovery, complete the following steps:

1. In the Groups pane of IBM Director Console, select **HMC Systems and Members**.
2. In the **HMC Systems Membership** pane, right-click the HMC and select **Request Access**.
3. Specify the user ID and password to access the HMC.

Launching Hardware Management Console tools

With the HMC Manager Tools, you can perform HMC management tasks such as accessing the Information Center and Setup Wizard, diagnosing problems with Service Focal Point, configuring the HMC, managing servers and frames, updating code for your HMC or operating system, and managing users.

You must have access to the HMC to complete this task.

To perform HMC management tasks, complete the following steps:

1. From the IBM Director Console window, select **Tasks** → **HMC Manager Tools** → *management task*.
2. Log in to the HMC with the ID and password. You can select the option to save the ID and password whenever the HMC console is launched.
3. Complete the management task.

Viewing Hardware Management Console membership

You can use IBM Director to view the topology and hardware resources associated with the HMC.

You must have access to the HMC to complete this task.

To view the topology of systems managed by the HMC, complete the following steps:

1. In the Groups pane of IBM Director Console, select **HMC Systems and Members**.
2. In the **HMC Systems Membership** pane, expand the HMC you want to view.
3. Expand the associated server to view logical partitions that are created on that server.
4. To view resource attributes of a managed object, double-click the managed object.

Monitoring performance

Use the Capacity Manager task to monitor performance across your systems. You can generate reports and graphs for performance data.

Activating a performance monitor

This task describes how to activate a performance monitor.

Tip: You cannot select a group of monitors by clicking the group name. You must select each monitor individually. If you have deactivated a monitor, it will not be reactivated until you reactivate it.

Complete these steps to activate a monitor:

1. In the IBM Director Console Tasks pane, expand the **Capacity Manager** task.
2. Drag the **Monitor Activator** subtask onto a managed system or group on which the Capacity Manager Agent is installed.
3. In the left pane of the Monitor Activator window, click the monitor that you want to activate.
4. Click **On**.
5. Click **Apply**.

Changing performance-analysis report settings

This task describes how to change the performance-analysis settings.

Complete these steps to change the performance-analysis report settings:

1. In the Report Viewer window, click **Edit icon (✎)** → **Settings**.
2. In the Settings window, click the **Graph** tab to configure the appearance of the graph in the Graph pane.
3. Click the **Window** tab to configure the appearance of the viewer.

4. Click the **Monitors** tab to configure the threshold settings for each monitor.
5. Click **Apply**.
6. Click **OK**.

Changing the graph display settings

This task describes how to configure the graph display settings in the Report Viewer window.

Complete these steps to configure display options for the graphs in a report:

1. In the Report Viewer window, click the **Graph** tab.
2. Type the maximum number of systems to display on the graph before combining results into a trend in the **Trend graph** field.
3. Type the dimensions of the grid size in the **Maximum grid size** fields.
4. Select the appropriate graph options from the **Graph options** box, including those for displaying maximum and minimum lines for averaged values, displaying the legend on the screen, and setting the thickness of threshold lines.
5. If you decide to return to the default settings, click **Return to defaults**.

Changing the monitor display settings

This task describes how to configure the monitor display settings in the Report Viewer window.

Complete these steps to configure monitor settings in a report:

1. From the Report Viewer window, click the **Monitors** tab.
2. On the Monitors page, click a monitor in the left pane to select it.
3. **Optional:** You can adjust the threshold settings on performance-analysis monitors to conduct resource planning to discover whether a bottleneck appears when capacity is set to a given value. However, changing the default settings will impair the performance-analysis function. If you want to adjust the threshold settings, type custom values for the warning threshold and critical threshold settings in the **Threshold settings** fields.
4. Select the **Show thresholds as a percent of maximum value** check box to display thresholds as a percentage of the maximum value on the report.
5. Select the **Adjust range to peak value** check box to set the peak value that is reported as the vertical range value of the graph.
6. If you decide to return to the default settings, click **Return to defaults**.

Changing the report window display settings

This task describes how to configure the window display settings in the Report Viewer window.

Complete these steps to configure the window display settings for the Report Viewer window:

1. In the Report Viewer window, click the **Windows** tab.
2. Decide whether to abbreviate column headings. If you choose to abbreviate column headings, type the maximum number of characters in the field.
3. Select **Small icons** or **Large icons**.
4. If you decide to return to the default settings, click **Return to defaults**.

Creating an event filter for bottlenecks

This task describes how to create an event filter for bottlenecks.

Complete the following steps to create an event filter specifically for bottlenecks:

1. In IBM Director Console, click **Tasks** → **Event Action Plan Builder**.
2. In the Event Action Plan Builder window, click **File** → **New** → **Simple Event Filter**.
3. In the left pane of the Event Type page in the Simple Event Filter Builder window, clear the **Any** check box.
4. In the right pane, expand **Capacity Manager**.
5. Expand **Bottleneck**.
6. Click **Recommendation**.
7. Click the **Extended Attributes** tab.
8. Clear the **Any** check box.
9. In the **Keywords** list, click **Hours since bottleneck first started**.
10. In the **Operator** list, click **Equal to**.
11. In the **Values** field, type **2**.
12. Click **File** → **Save As**.
13. In the Save Event Filter window, type the name of the filter.
14. Click **OK** to save the filter. The new filter is displayed in the Event Filters pane under **Simple Event Filter**.

Creating a performance-analysis report

This task describes how to create a performance-analysis report using a predefined or customized report definition.

To create a report using a predefined or customized report definition:

1. In the IBM Director Console Tasks pane, expand **Capacity Manager**.
2. Expand **Report Generator**.
3. Drag the report definition that you want to use onto one or more managed systems or group. A status window opens to indicate the progress. If the report definition specifies that the report is sent to the report viewer, the Report Viewer window opens. If the report definition specifies that the report is sent to a file, the report is saved automatically. Unless you specify a directory in the report definition, the file is saved to the default directory, IBM\Director\reports..
4. If you selected a report definition that is saved to a file, specify one of these options:
   - Click **Execute Now** to generate the report now. A status window opens to indicate progress. The report is saved to the directory specified in the report definition or to the default directory, \IBM\Director\reports.
   - Click **Schedule** to set a time to generate the report.

If the report definition specifies that the report is to be displayed in the report viewer, the Report Viewer window opens.

Creating a performance-analysis report definition

After you create a customized report definition, you can generate a performance-analysis report that includes only those parameters that you have specified.

Complete these steps to create a new performance-analysis report definition:

1. Run the **Monitor Activator** subtask on a managed system or group to activate the monitors for that system or group.
2. Expand **Report Generator**
3. Double-click **New Report Definition**.
4. In the Report Definitions window, specify the report parameters:
   a. Click the **Report Parameters** tab.
   b. Select the report duration, the global sampling frequency, and whether to collect minimum and maximum values.

      Note: Selecting the **Collect min and max values** check box specifies that the minimum and maximum data points for each sample are collected. An advantage of collecting the minimum and maximum data points is that you can use a slower sampling frequency, which collects data less frequently; using this option reduces the size of the report while still providing informative managed-system performance data. If memory usage is an issue, consider using a slower sampling frequency. Note that the average is always collected.
   c. Select number of minutes Capacity Manager waits for a system to respond before considering the system unable to provide the data from the **Timeout per system** drop-down list.
   d. Click **New**. In the New Time Interval window, specify the time of the report.
5. Specify the methods you want to use to generate the report:
   a. Click the **Method of Generating a Report** tab.
   b. Select **Generate to Viewer** or **Generate to File**.
   c. Select the appropriate file-format check boxes to generate the file in the selected formats. The default file format is XML.
   d. If you selected **Generate to SQL**, type the database Web address and root table name in the applicable fields.

      Note: The ability to generate a report in SQL format is available on a Window platform only.
   e. If you use SQL authentication, type the user ID and password for the SQL connection in the applicable fields.
   f. Select **Generate Bottleneck events** to generate an event in the IBM Director event log.
   g. Select **Back up existing files** to archive saved reports.
   h. Select **Number of backups** to set the number of reports to keep.
   i. Select **Set destination directory** to set the destination directory.

      Note: The default destination directory is IBM\Director\reports.
6. Specify the monitor selection:
   a. Click the **Monitor Selection** tab.

b. On the Monitor Selection page, click **Include all activated monitors** to include all active monitors in the report, or click **Select individual monitors** to select specific monitors.

c. If you chose to set individual monitors:

1) Click the appropriate monitors in the **Monitors** list.

2) Click **Include** or **Exclude** to include or exclude the selected monitors.

d. Select the **Override global settings** check box to use a different sampling frequency than the default.

7. Specify the thresholds:

a. Click the **Threshold Settings** tab.

Note: A threshold setting applies to all managed systems included in the report definition.

b. In the **Critical threshold** field, type the value of the critical threshold.

c. In the **Warning threshold** field, type the value of the warning threshold.

d. Click **Return to defaults** to set the threshold values to default values.

e. Click **Save As**.

f. In the Save As window, type the name of the report definition and click **OK**.

Deactivating a performance monitor

This task describes how to deactivate a performance monitor.

Note: You cannot select a group of monitors by clicking the group name. You must select each monitor individually. If you have deactivated a monitor, it will not be reactivated until you reactivate it.

Complete these steps to deactivate a performance monitor:

1. In the IBM Director Console Tasks pane, expand the **Capacity Manager** task.
2. Drag the **Monitor Activator** subtask onto a managed system or group on which the Capacity Manager Agent is installed.
3. In the left pane of the Monitor Activator window, click the monitor that you want to deactivate.
4. Click **Off**.
5. Click **Apply** to close the Monitor Activator window.

Identifying bottlenecks

This task describes how to identify bottlenecks in a managed system or group.

Use one of the following methods to identify bottlenecks in a managed system or group:

- Schedule performance analysis to check for bottlenecks and generate an event when a threshold is exceeded or met.
- Use the Report Generator function to generate a report immediately.

If a bottleneck is found, the monitor name is shown in bold and in red in the performance-analysis section of the report, and recommendations for correcting the bottleneck are made. If no bottleneck is found, the performance-analysis icon indicates that no bottlenecks were found.

Printing a performance-analysis report

This task describes how to print a performance-analysis summary and graph.

Perform one of the following procedures to print either a performance-analysis summary or a graph:

- To print a performance-analysis summary, in the Report Viewer window, click the **File icon (📁) → Print → Performance analysis report**
- To print a performance-analysis graph, in the Report Viewer window, click the **File icon (📁) → Print → Graph Print**

Saving a performance-analysis report

This task describes how to save a performance-analysis summary and graph.

You can save a performance-analysis report summary in HTML format for later viewing and printing from a Web browser. The report in HTML format includes the monitor and managed-system parameter information from the Table view in Report Viewer window.

You can also save a performance-analysis graph in a GIF file. This report includes the managed-systems information from the Graph pane in the Report Viewer window.

- Complete these steps to save a performance-analysis summary on the management console as an HTML file:
  1. Click **File icon (📁) → Export report to local HTML**.
  2. Type a new file name.
  3. In the Export Report to Local HTML window, click **Save**.
- Complete these steps to save a performance-analysis summary on the management server as an HTML file:
  1. Click **File icon (📁) → Export report to remote HTML**.
  2. Type a new file name.
  3. In the Export Report to Remote HTML window, click **Save**.
- Complete these steps to save a performance-analysis graph on the management console as an HTML file:
  1. Click **File icon (📁) → Export graph to local HTML**.
  2. Type a new file name.
  3. In the Export Graph to Local HTML window, click **Save**.
- Complete these steps to export a performance-analysis graph on the management server as a GIF file:
  1. Click **File icon (📁) → Export graph to remote GIF** .
  2. Type a new file name.
  3. In the Export Graph to Remote GIF window, click **Save**.

Scheduling to check for bottlenecks

You can set up a schedule to have IBM Director check for bottlenecks on a regular basis. When a bottleneck is detected, an event is added to the event log, and IBM Director generates a report.

You must select the **Generate bottleneck events** check box for the report definition that you are using. Otherwise, an event action plan cannot notify you that a bottleneck has occurred, because event action plans depend on events to trigger event actions.

Complete these steps to schedule regular checks for bottlenecks:

1. In the IBM Director Console Tasks pane, expand the **Capacity Manager** task.
2. Expand the **Report Generator** subtask. Drag **Hourly Bottleneck Events** onto the managed system or systems or group that you want to monitor for bottlenecks.
3. Click **Schedule**.
4. In the New Scheduled Job window, type a job name.
5. Select the date and time for the initial check.
6. Click **Advanced** to schedule the job to repeat at regular intervals.
7. On the **Date/Time** page of the New Scheduled Job window, select the **Repeat** check box.
8. In the **Repeats** group box of the Repeat window, select how often you want to check for performance bottlenecks.
9. Click **OK**.
10. Click **File** → **Save As**.
11. In the Save Job window, type a descriptive name for the scheduled job. Click **OK**. A confirmation message is displayed, indicating that you have saved the job.
12. Click **OK** to close the message window.

Setting up automatic notification of bottlenecks

This task describes how to sets up automatic notification when bottlenecks occur.

Capacity Manager can determine where and when bottlenecks occur. Complete these steps to set up automatic notification when a bottleneck occurs:

1. Set up a schedule to periodically check for performance bottlenecks and to generate an event when a performance threshold is met or exceeded, indicating a bottleneck.
2. Create an event filter to notify you of the event.
3. Create an event action plan that uses the event filter.
4. Apply the event action plan to the managed systems or groups that you want to monitor.

Viewing a performance forecast graph

This task describes how to view a performance forecast graph.

To view the forecast graph for a selected managed system, go to the Report Viewer window. Click the Forecast icon ( ) in the lower-right corner of the lower-right pane. Capacity Manager displays the forecast graph for the selected monitor.

Notes:

1. You cannot use the Zoom tool and the Forecast tool at the same time.
2. The forecast data is more meaningful for managed systems that are individually graphed rather than shown in a trend graph. To change from a trend graph to a graph of individual managed systems, either

set your trend graph threshold to a higher number or select fewer managed systems to graph at one time.

Viewing previously generated reports

This task describes how to view reports that were previously generated.

Note: If you use the Report Viewer to display a report that was saved in XML format, you can adjust the threshold settings on performance-analysis monitors.

Complete these steps to view a previously generated report:

- To view a report that has been saved to the management server:
  1. In the IBM Director Console Tasks pane, expand **Capacity Manager**.
  2. Right-click **Report Viewer**.
  3. Click **Open**.
  4. In the Open Remote Report window, select a file.
  5. Click **Open**.
- To view a report that has been saved to the management console:
  1. In the IBM Director Console Tasks pane, expand **Capacity Manager**.
  2. Right-click **Report Viewer**.
  3. Click **File icon (📂)** → **Open local report**.
  4. In the Open Local Report window, select a file.
  5. Click **Open**.

Viewing performance statistics

This task describes how to view a performance statistics for active performance monitors.

In the IBM Director Console Tasks pane, expand the **Capacity Manager** task. Drag the **Monitor Activator** subtask onto a managed system or group on which the Capacity Manager Agent is installed.

In the left pane, all monitors are displayed in a tree structure; each monitor has an icon to indicate its status. The names of performance-analysis monitors are displayed in bold.

In the Systems pane, an icon is displayed beside each managed system or group to indicate its status. In the Legend pane, the monitor and managed-system icons and their descriptions are displayed.

Monitoring system availability

You can use the System Availability task to analyze the availability of a managed system or group. You can also use this task to view statistics about managed-system uptime and downtime through reports and graphical representations.

Viewing system availability

This topic describes how to start the System Availability task in IBM Director.

In the IBM Director Console Tasks pane, drag the **System Availability** task onto a managed system or group that supports System Availability.

The list on the toolbar in the System Availability window has four options:

Distribution of System Outages
A pie chart representing the percentage of all system outages.

Distribution of System Uptime
A pie chart representing the percentage of all system uptime.

System Outages by Day of Week
A bar chart measuring the frequency of outages by day of the week, with planned and unplanned outages differentiated.

System Outages by Hour of Day
A bar chart measuring the frequency of outages by hour of the day, with planned and unplanned outages differentiated.

To see the value of a specific pie chart or bar chart section, move the cursor over a specific section.

Notes:

1. (Windows operating systems that support IBM Director and are configured to adjust automatically for daylight saving time only) The event times that are specified in the system-availability report might vary by 1 hour from the event times in the Windows event viewer, because the Windows event viewer adds or subtracts one hour to adjust for daylight saving time. Because this adjustment can cause duplicate entries in the system-availability database when the time adjustment is made, System Availability does not use the daylight saving time adjustments.
2. (Linux only) On managed systems where compression of message logs is the default, turn off compression of message logs to view system-availability reports.
3. System Availability reads the message logs only if the message logs are in their default directory.
4. System Availability should run as or more often than the message logs are archived to avoid losing availability information.

The availability report is a snapshot of system availability. It provides an overall statistical summary of event and problematic details and measurements for the currently selected managed systems in a tree structure, or all managed systems if the root of the tree is selected. Systems identified as problematic are listed in the detail section and are flagged with a red X. There are two types of reports that you can view by following these steps:

To view the availability report, click **View** → **Availability Report**.

To view a more detailed view of the availability report, right-click the graph, and then click **Detailed List of Record**.

In the System Availability window, you can detach the current view to compare and contrast different system-availability views and time frames. Click **View** → **Detach View**. The current view is separated as an independent window that does not reflect subsequent changes to the report. Closing the System Availability task closes any detached view windows.

With the exception of a detached view, you can print any window that is displayed in the System Availability task by clicking **File** → **Print**.

Changing the graph dates

This topic describes how to specify the time period for which data is graphed in IBM Director.

Complete the following steps to specify the time period for which data is graphed:

1. In the System Availability window, click **File** → **Set Time**.
2. In the Customization of Graph Dates window, in the **Select Date Range** field, select one of the following time ranges for which you want to view data.

   All Displays system-availability data from the time that System Availability was loaded on the target system up to the present day. This selection is the default.

   1 week
   Displays system-availability data from one previous week up to midnight of the present day.

   1 month
   Displays system-availability data from one previous month up to midnight of the present day.

   3 months
   Displays system-availability data from three previous months up to midnight of the present day.

   1 year Displays system-availability data from one previous year up to midnight of the present day.

   Customize
   Customizes the range of time for which to display system-availability data.

   Note: If you select **Customize**, type the From and To dates in the applicable fields.
3. Click **Update**.

   Note: These customized settings apply only to the System Availability report that is currently open; they are not global settings applicable to all System Availability reports.

Changing the settings criteria for System Availability

This topic describes how to change the System Availability settings criteria in IBM Director.

System Availability scans for problematic systems within a range of time. The time begins a specified number of days in the past (the default is 30) and ends with the current time. The number of unplanned outages that occur in this time frame is counted, and if the total number meets or exceeds the specified count, the managed system is marked as problematic. You can also specify a percentage of time in which the managed system has unplanned outages, instead of a specific number of outages, by selecting the **Percentage** check box.

1. To specify the settings criteria, click **File** → **Settings**.
2. In the Settings window, change any of the criteria; then, click **Save**.

Note: Select **Use all available data** to evaluate all persistent data available in the IBM Director Server database.

All system-availability reports that are run after you click **Save** use the new settings.

Saving the system-availability report

This topic describes how to save the current system-availability report in IBM Director.

You can save the current report as a series of HTML files to a directory on the management console. Then, you can view the report in a Web browser at a later time. You also can save the current report in XML format.

- Complete the following steps to export and save a report in HTML format:
  1. Follow the steps in the Starting the System Availability topic to generate a system-availability report.
  2. After the report is generated, click **File** → **Export Availability Report** → **Export HTML Report**.
  3. In the “Select a directory to save report files” window, type a file name and click **Select**.
  4. In the “Confirm Directory window, click **OK**. The files are saved to the location that you specified.
  5. (Windows only) In the Open saved file window, in the **File name** field, type a file name; then, click **Select** to save the report to the specified location.
  6. (Windows only) Click **Yes** to open the exported report in a Web browser immediately.
- Complete the following steps to export and save a report in XML format:
  1. Follow the steps in the Starting the System Availability topic to generate a system-availability report.
  2. After the report is generated, click **File** → **Export Availability Report** → **Export HTML Report**.
  3. In the “Select a directory to save report files” window, type a file name and click **Select**.
  4. In the Confirm Directory window, click **OK**. The files are saved to the location that you specified.

Monitoring system status

This topic describes the System Status task in IBM Director.

Clearing system status flags

This topic describes how to clear a system status flag on a managed object in IBM Director.

Complete the following steps to clear a system status flag on a managed object:

1. In the Groups pane, click **All Systems and Devices**. The Group Contents pane displays the managed objects.
2. In the Group Contents pane, right-click the managed object to which you added the Disk Information system status, and then click **System Status**. The System Status menu is displayed.
3. Click **Disk Information**. The menu closes and the Disk Information icon is removed from the managed object.

Setting system status flags

This topic describes how to set a system status flag on a managed object.

Complete the following steps to set a system status flag on a managed object:

1. In the Groups pane, click **All Systems and Devices**.
2. Right-click a managed object in the Group Contents pane and click **System Status**.
3. In the System Status menu, click **Disk Information**. The menu closes and a Disk Information icon is displayed next to the managed object icon.

Scheduling tasks

This topic provides information about using the Scheduler task.

You can use Scheduler to run a single noninteractive task or set of noninteractive tasks at a later time. (Only *noninteractive tasks*, which are defined as tasks that do not require any user input or interaction, can be scheduled.) You can specify an exact date and time you want the task to be started, or you can schedule a task to repeat automatically at a specified interval. Scheduled tasks are referred to as jobs.

IBM Director does not allow saving changes to an existing job; you must always save changes to an existing job as a new job.

Scheduling a task

This topic describes how to schedule a task.

Complete the following steps to schedule a task:

1. In IBM Director Console, click **Tasks** → **Scheduler**.
2. In the Scheduler window, double-click the date on which you want the new job to start.
3. In the New Scheduled Job window, specify information about when the job will run:
   a. Select the **Schedule the task to execute on a date and time** check box to assign a date and time to the job. If this check box is cleared, the job is added to the jobs database, but it is not activated automatically.
   b. In the **Date** field, type the date that you want the job to run.
   c. In the **Time** field, select the time that you want the job to run.
4. Click the **Task** tab.
5. On the Task page, in the Available pane, double-click the task that you want the job to perform. The task is added to the Selected Task pane. You can select multiple tasks for a single job. Each task is processed in the order in which it is displayed in the Selected Tasks pane.
6. Click the **Targets** tab.
7. If you want to run the job on a managed system, on the Targets page, in the Available pane, double-click a managed system. The managed system is added to the Selected Group pane. Repeat this step until you have added all of the managed systems on which you want to run the job.
8. If you want to run the job on a group, click **Use a group as the target**. The window is populated with group information. In the Available pane, double-click the group. The group is added to the Selected Group pane. You can select only one group as a target for any job.

9. Click the **Options** tab.
10. In the Options page, select job options. You can choose from the following options:

    Delay execution on unavailable systems
    If you select this check box, the task will be performed on any offline targeted managed objects as soon as they are online again.

    If you do not select this check box and a targeted object is offline at the time of job activation, the job returns an error status.

    Execute on systems that are added to the target group
    (This option is available only if the job is targeted against a group.) If you select this check box, the job remains active until it is cancelled. IBM Director detects any new managed objects that are added to the target group, and the scheduled job then is activated on those new managed objects.

    Execute in client time zone
    If you select this check box, the task is run at the designated time in the local time zone of the target managed object.

    Note: You cannot schedule a job to repeat hourly and be run in the local time zone. Also, if the start date of the first scheduled time zone occurs before the target managed object date, the job cannot be created.
11. Select execution history options:
    a. To limit the number of job records that are stored in the execution history, select the **Limit execution status** check box and specify the number of job records that are stored.
    b. To delete job records of a certain age, select the **Delete execution status** check box and specify the number of days after which job records are deleted.
12. Specify when events are generated. You can choose from the following options:

    Generate an event when a job execution completes successfully

    Generate an event when a job execution completes with errors

    Generate an event when a target system completes successfully
    An event is generated when a job is successfully run on a target system.

    Generate an event when a target system completes with errors
    An event is generated when errors occur as a job is run on a target system. For example, if a target object does not respond, the target object is completed with errors
13. Click **File** → **Save As**.
14. In the Save Job window, type a descriptive name for the scheduled job. Click **OK**.
15. Click **OK** to close the confirmation message window.

Viewing information about scheduled jobs

This topic describes how to view information about previously scheduled jobs.

To open the Scheduler window and view information about previously scheduled jobs, in IBM Director Console, click **Tasks** → **Scheduler**.

The Scheduler window has these four pages:
- **Month Calendar**
- **Week Calendar**
- **Day Calendar**
- **Jobs**

The first three pages are calendar pages; the Jobs page lists all the scheduled jobs.

Using the Calendar pages

This topic describes how to use the Calendar pages in IBM Director. The three calendar pages, Month, Week, and Day, display the schedules for the execution of all jobs.

Note: The calendars are independent of each other. Changing the date on one calendar does not change the date on another calendar. Also, selecting a job on one calendar does not select it on other calendars.

To view the execution history for a job, right-click a job, and then click **Open Execution History**.

Viewing job information

This topic describes how to view job information in IBM Director.

The Jobs page displays a list of all scheduled jobs and status information for job executions. When you click a scheduled job type in the left pane, information about that job type is displayed in the right pane. The information includes the number of executions that are active or complete, the next date the job will be executed, the tasks that the job will perform, and any options that have been specified for the job.

When you click a specific execution of a scheduled job in the left pane, information about that job execution is displayed in the right pane. The information that is displayed is identical to the information in the Execution History window.

Viewing job properties

This topic describes how to view job properties in IBM Director.

To view the properties of a scheduled job in the Scheduler window, right-click a job and click **Open Job Properties**. The Scheduled Job window opens for the job, with four pages, Date/Time, Task, Targets, and Options.

You can use the Scheduled Job window to change the properties of a job and save it as another scheduled job. IBM Director does not allow saving changes to an existing job; you always must save it as a new job.

Viewing scheduled job history information

This topic describes how to view scheduled job history information in IBM Director. The Scheduler maintains the execution history information for immediate executions and scheduled jobs.

To view information about the execution of a scheduled job in the Scheduler window, right-click a job and click **Open Execution History**.

The Execution History window displays the overall status of the job. The top pane shows a summary of the status (for example, Complete) for the target objects. Target objects are grouped together according to the status of each target for an execution and are displayed in the bottom pane of the window.

Viewing execution history logs

This topic describes how to view execution history logs in IBM Director.

To view the entire log for an execution history in the Scheduler window, right-click a job and click **View Log**.

Launching external applications

The External Application Launch task allows third party management software and other programs that function outside of IBM Director to be integrated with IBM Director. This new mechanism uses easy-to-create files, called CMDTask files, to add tasks that can be launched for managed objects. After a task is added, you can launch the task by opening it or by dragging and dropping it onto a managed object.

After you create a CMDTask file, it appears as a subtask of the **External Application Launch** task in the Task pane of IBM Director Console. You can create single tasks or groups of related tasks.

Creating a command task file

This topic describes how to create CMDTask files that are used to integrate third party management software and other programs with IBM Director.

CMDTask files must follow specific style requirements, as described below:

- All CMDTask files must reside in the *director_install_dir*\classes\extensions directory and must have an extension of ".CMDExt".
- The extension .CMDExt is not case sensitive.
- If creating a group of related tasks (a parent with child tasks), all .CMDExt filenames must have the same prefix.
- All child task filenames must end with _*n* (for example, TELNET_1.CMDExt, TELNET_2.CMDExt, TELNET_3.CMDExt...).

When IBM Director Server is started, all files with an extension of .CMDExt in the *director_install_dir*\classes\extensions directory are processed. In IBM Director Console the title of each task, which is specified within the file with the Title parameter, appears as a subtask of the **External Application Launch** task in the Tasks pane.

For example, if you have a file called TELNET.CMDExt (with Title=telnet command) residing on the IBM Director Server in the directory *director_install_dir*\classes\extensions, the task label **telnet command** is displayed as a subtask of the **External Application Launch** task.

You can also create a group of related tasks (one parent with a child tree). For example, you could create the following group:

1. A parent task with a file name of FILESYSTEM.CMDExt and Title=File System commands

2. A child task with a file name of FILESYSTEM_1.CMDExt (the "_1" is required) and a) Title=Volume command and b) ParentTaskFilename=FILESYSTEM
3. A child task with a file name of FILESYSTEM_2.CMDExt (the "_2" is required) and a)Title=Dir command and b) ParentTaskFilename=FILESYSTEM
4. A child task with a file name of FILESYSTEM_3.CMDExt (the "_3" is required) and a) Title=DiskCopy command and b) ParentTaskFilename=FILESYSTEM

As noted in the example, you must append _*n* to the name of each child task file, and include both the Title and ParentTaskFilename parameters in each file. Note that when you launch a parent task, any child task beneath the parent task is not executed. The parent-child relationship is only for display and organizational purposes within External Application Launch. No other programmatic connections exist.

The task titles are sorted within a group. For example, the tasks specified in the examples would display in the following order:

```
External Application Launch
    File System commands
        Dir command
        DiskCopy command
        Volume command
    telnet command
```

To refresh the list of subtasks, right-click the **External Application Launch** task, and then click **Refresh**.

To restrict the use of the External Application Launch task to specific users, use the **Task Access** menu options under **User Administration**.

.CMDExt file parameters:

Note: An * indicates parameters that you are strongly encouraged to set. Parameters and values are case sensitive. Each parameter must be defined on a separate line, and when editing the .CMDExt file (properties file), ISO 8859-1 character encoding must be used.

Parameter name	Description	Value information
Title*	Title of the task or NLS key for the title that will appear as a subtask of **External Application Launch**	It is recommended that you set this parameter. If not set, the base file name of the .CMDExt file is used.
ResourceBundle	NLS bundle used for Title	Optional Parameter—This is a Java ResourceBundle. If specified, these files contain the Title key and the translated Title key value. Specify the full class name.

Parameter name	Description	Value information
ParentTask Filename	When creating a child task file, the name of the parent task file	Specify this parameter only if you are creating a child task. Specify the file name of the parent task without the .CMDExt extension. For example: ParentTaskFilename=FILESYSTEM The task label for the child task appears below the task label for the parent task. The existence of the ParentTaskFilename is checked.
CommandString. Windows*	The command string to be executed on a Windows system.	*Important:* .CMDExt files use Java properties file formatting; thus, any backslash must be coded as a double backslash. Environment variable values should have only one backslash. For example: `dir c:\\*exe /s`
CommandString. Unix*	Command string to be executed on a UNIX/Linux system	*Important:* .CMDExt files use Java properties file formatting; thus, any backslash must be coded as a double backslash. Environment variable values should have only one backslash.
Cwd.Windows	The path name of the current working directory on a Windows system	Specify only if the command needs to start in a specific directory. *Important:* .CMDExt files use Java properties file formatting; thus, any backslash must be coded as a double backslash. Environment variable values should have only one backslash.
Cwd.Unix	The path name of the current working directory on UNIX/Linux systems	Specify only if the command needs to start in a specific directory. *Important:* .CMDExt files use Java properties file formatting; thus, any backslash must be coded as a double backslash. Environment variable values should have only one backslash.
ShellRequired	Indicates whether a shell window is required.	Specify one of the following values: **true** Shell window is required. **false** Shell window is not required (default).
Icon.Small	The path to the CMDTask's small icon, which appears next to the task label in the subtasks list under **External Application Launch**	This path assumes you are in *director_install_dir*/classes. Start the path with a slash (/) before typing in any path name. For example, if the icon is located under *director_install_dir*/classes/icondir, the value for the parameter is: Icon.Small = /icondir/iconname16.gif

Parameter name	Description	Value information
Icon.Large	The path to CMDTask's large icon, which appears in the Tasks pane on Director console.	This path assumes you are in *director_install_dir*/classes. Start the path with a slash (/) before typing in any path name. For example, if the icon is located under *director_install_dir*/classes/icondir, the value for the parameter is: Icon.Large = /icondir/iconname32.gif
Targeted	Indicates whether the CMDTask needs to be dragged and dropped onto a target system.	Specify one of the following values: **none** Task can be opened using a right mouse click or double-clicked **one** Task needs to be dragged and dropped onto a target system (default). **none|one** Task can be opened or dragged and dropped onto a target system.
Timeout	The number of seconds to wait for the task command to return after execution. IBM Director Console will be busy during the timeout period.	Specify a numeric value from 1 through 60. Default is 5 seconds. Maximum value allowed is 60 seconds.

Example .CMDExt files:

These examples are provided as common usages to help avoid syntax confusion. Use these examples as guidance for creating simple .CMDExt files.

The following example shows how to create a command task to open a telnet session and hold the window open while users are typing in their userid and password. With the –hold option, you also receive an error message if the telnet command is not executed successfully.

Example 1: **Telnet.CMDExt**

```
Title= Telnet Command
#On UNIX/Linux:
CommandString.Unix = xterm –hold –e telnet  $CMDTASK_IP_ADDRESS0
#On Windows:
CommandString.Windows = telnet  %CMDTASK_IP_ADDRESS0%
ShellRequired=true
```

The following example shows how to create a command task to net use into a system.

*Example 2:***NetUse.CMDExt**

```
Title= Net Use Command
#On Windows:
CommandString.Windows = net use * \\\\%CMDTASK_COMPUTERNAME%\\c$ /u:userid pwd
ShellRequired=false
```

Launching a command task file

After you create a CMDTask file and the task is displayed on the IBM Director Console under the **External Application Launch** task, you can start the task in the same ways as you start any other IBM Director task.

For targeted tasks, drag the task onto a managed system. For non-targeted tasks, double-click on the task or open it with a right mouse click. When you start the task, the IBM Director Console launches the program specified on the appropriate “CommandString” line in the CMDTask file. However, if the task is targeted, several special environment variables can be set before the program is launched – this allows the targeting information to be passed to the program.

Special environment variables:

Note: These variables can be used with CommandString.Windows and CommandString.Unix parameters.

Variable Name	Description
CMDTASK_IP_ADDRESS*i*	This variable retrieves an IP Address from the TCP/IP Addresses attribute of the targeted system. Because a targeted system could have multiple IP addresses, the *i* should be replaced with 0, 1, 2, ... beginning with zero as the first address in the attribute list for the targeted system. Example: CMDTASK_IP_ADDRESS0
CMDTASK_MAC_ADDRESS*i*	This variable retrieves a MAC Address from the MAC Addresses attribute of the targeted system. Because a targeted system could have multiple MAC addresses, the *i* should be replaced with 0, 1, 2, ... beginning with zero as the first address in the attribute list for the targeted system. Example: CMDTASK_MAC_ADDRESS0
CMDTASK_COMPUTERNAME	This variable retrieves the ComputerName attribute of the targeted system.

Usage Recommendations:

- Execute CommandString.Windows and CommandString.Unix values from a command line first to ensure that they execute successfully outside of the CMDTask file. This saves debugging efforts. Do not chain any commands or use pipes or redirection as only a single command can be executed.
- If the launched application requires an interactive dialog or advanced command line syntax, such as wildcard expansion, or special characters, set the ShellRequired parameter to true. This will start a persistent shell on Windows via ‘start cmd.exe /k’ and on UNIX/Linux via ‘bash –c’. Without the ShellRequired parameter set to true (the default is false), the CommandString is just executed. Any output is shown in a window that you can scroll and later close. Also, if it takes awhile to launch the application, specify a Timeout value in seconds to handle the time it takes to launch the application.
- When using the environment variables, CMDTASK_IP_ADDRESS*i* and CMDTASK_MAC_ADDRESS*i*, always remember to append a number at the end. If it is for the first address, append 0 to the variable.

- When a CMDTask is defined on the IBM Director Server, it is displayed on IBM Director Consoles. However, to launch the command (application) successfully, the necessary applications or files must also reside on the IBM Director Console machine as well.
- Because .CMDExt files are Director properties files, if any parameter is defined more than once, only the last definition found is used. Additionally, any misspelled parameter is not detected, and is ignored.
- When editing the .CMDExt file (properties file), ISO8859-1 character encoding is used. For characters that cannot be represented directly in this encoding, Unicode escapes are used. However, only a single "u" character is allowed in an escape sequence.

Managing processes and services

This topic provides information about managing processes and services.

Stopping an application

This topic provides information about stopping an application on a managed system.

Complete the following steps to stop an application:

1. In the IBM Director Console Tasks pane, drag the **Process Management** task onto a managed system.
2. On the Applications page in the Process Management window, right-click the application that you want to close, and click one of the following values, depending on your operating system:

For a system running i5/OS	**End Job** or **End Job Immediate**
For a system running Linux	**Kill process** or **Terminate signal**
For a system running NetWare	**Unload Module**
For a system running Windows	**Close Application**

3. A confirmation window is displayed. Click **Yes**.

Viewing processes and services

This topic provides information about viewing processes and services.

To view processes, services, and device-services information, do the following:

In the IBM Director Console Tasks pane, drag the **Process Management** task onto a managed system.

The Process Management window contains one or more of the following pages:

Applications
Displays all processes that are running on the managed system or group.

Services
(Windows only) Displays the status and description of all Windows services that are installed on the managed system or group.

Device Services
(Windows only) Displays all hardware device drivers that are installed on the managed system or group.

Subsystems
(i5/OS only) Displays the status of i5/OS subsystems.

Servers
(i5/OS only) Displays the status of servers that are installed.

Working with i5/OS processes and services

This topic provides information about working with i5/OS processes and services.

Starting servers, ending servers, and showing jobs on i5/OS

This topic describes how to start or stop installed servers running on i5/OS or show jobs on the installed server.

Complete the following steps to start or stop installed servers running on i5/OS or show jobs on the installed server:

1. In the IBM Director Console Tasks pane, drag the **Process Management** task onto a managed system.
2. In the Process Management window, click the **Servers** tab. Right-click the server that you want to manage and click **Start Server**, **End Server**, or **Show Jobs**.

   Note: IBM Director does not currently support starting and stopping all types of servers. Some servers might be displayed that you cannot start or stop through IBM Director.

Starting subsystems, ending subsystems, and showing jobs on i5/OS

This topic describes how to start or stop i5/OS subsystems or show jobs on an i5/OS subsystem.

Complete the following steps to start or stop i5/OS subsystems or show jobs on an i5/OS subsystem:

1. In the IBM Director Console Tasks pane, drag the **Process Management** task onto a managed system.
2. In the Process Management window, click the **Subsystems** tab. Right-click the subsystem that you want to manage and click **Start Subsystem**, **End Subsystem**, or **Show Jobs**.

Working with Windows processes and services

This topic provides information about working with Windows processes and services.

Changing the state of a Windows service

This topic provides information about changing the state of a Windows service. You can start, stop, pause, or resume a Windows service.

Complete the following steps to change the state of a Windows service:

1. In the IBM Director Console Tasks pane, drag the **Process Management** task onto a managed system.
2. In the Process Management window, click the **Services** tab and right-click the service that you to want to start, stop, pause, or resume; then, click the applicable choice.

Setting the priority of a Windows application

This topic provides information about setting the priority of a Windows application.

Complete the following steps to set the priority of a Windows application:

1. In the IBM Director Console Tasks pane, drag the **Process Management** task onto a managed system.
2. In the Process Management window, click the **Applications** tab.
3. Right-click the application, and click **Set Priority**; then click the priority level. A confirmation window is displayed.
4. Click **Yes**.

Managing systems with ASF

This topic provides information about managing systems with Alert Standard format (ASF).

Configuring ASF

This topic describes how to configure ASF on a managed system. The available configuration options vary, depending on the type of network interface card (NIC) and the level of ASF that it supports.

To configure ASF:

- The managed system must contain an ASF-capable NIC, and the applicable device drivers must be installed.
- IBM Director must have performed an inventory collection on the managed system.

Complete the following steps:

1. In IBM Director Console, drag the **Configure ASF** task onto the managed system.
2. In the Alert Standard Format window, complete the following steps to enable ASF on the NIC:
   a. Select the **Enable ASF Hardware** check box.
   b. Select the **Enable all Platform Event Traps** check box if it is displayed. (This check box displays only for certain types of NICs.)
   c. (ASF 2.0 systems only) To enable remote power management, select the **Enable Remote Management** check box.
3. Click the **Configuration** tab.
4. In the Configuration page, configure the ASF settings:
   a. In the **Management Server (IP address)** field, type the IP address of the management server to which the PET alerts are sent.
   b. To ensure that heartbeat alerts are sent, select the **Enabled** check box and type the number of seconds in the **Heartbeat frequency (seconds)** field.
   c. In the **Minimum watchdog timer (seconds)** field, type the minimum number of seconds for the watchdog timer.
   d. In the **Minimum ASF Sensor Inter Poll Wait Time (5 ms units)** field, type the minimum time to wait for the ASF Sensor Inter Poll.
5. (ASF 2.0 systems with remote management enabled only) Click the **Remote Management** tab.

6. In the Remote Management page, create or modify authentication keys:
   a. If you have not created authentication keys previously, click **Generate Keys**. Three authentication keys are generated. If you want to copy the authentication keys, do so now. When you click **Apply**, the key values are replaced by asterisks.
   b. If you want to create new authentication keys, select the **Overwrite the existing keys used by IBM Director Console for authentication** check box.
7. Click **Apply** . If you created or modified authentication keys, the keys are written to both IBM Director Server and the managed system.

Performing power-management operations on ASF 2.0 systems

This topic describes how to perform a power-management operation on an ASF 2.0 system. You can power the system on, power the system off, and restart the system.

Complete the following steps:

1. In the IBM Director Console Groups pane, click the **Systems with ASF Secure Remote Management** group. The managed systems are displayed in the Group Contents pane.
2. **Optional:** Click multiple managed systems in the Group Contents pane.
3. Right-click the managed system or systems in the Group Contents pane and click **Power Management**; then, click the operation that you want to perform on the managed system or systems.

Testing secure communication with an ASF 2.0 system

This topic describes how to test secure communication with an ASF 2.0 system. This procedure allows you to verify that the authentication keys saved on the management server match the authentication keys saved on the managed system.

The managed system must be configured for ASF. The authentication keys necessary for secure management must have been created previously.

Complete the following steps to test secure communication with an ASF 2.0 system:

1. In IBM Director Console, drag the **Configure ASF** task to the managed system.
2. In the Alert Standard Format window, click the **Remote Management** tab.
3. Click **Test**.

Managing asset information

This topic provides information about working with asset information. You can use the Asset ID task to configure asset information for a managed system or a group; then, you can inventory the asset information and use it as a basis for event action plans.

Configuring asset information for a group

This topic provides information about configuring asset information for a group of systems. You can create a profile that contains the asset information and run the profile against the group of systems.

Complete the following steps:

1. Create an Asset ID profile (see "Creating an Asset ID profile" on page 149).
2. Drag the profile onto the group of systems that you want to configure.
3. When the profile is applied successfully to all managed systems in the group, close the Status window.

The information that you provided is written to the EEPROM. For systems that are not enabled for Enhanced Asset Information Area (EAIA), this information is written to the asset.dat file in the data directory. The asset.dat file is deleted when IBM Director is uninstalled.

Note: For AIX, i5/OS, and Linux, the data is stored in a file.

Configuring asset information for a managed system

This topic describes how to configure asset information for a managed system. You can configure information about the system user, lease, and warranty; you also can provide data in up to five user-defined fields. If you specify an end date for a lease or warranty, a warning event is generated on that date.

Complete the following steps to configure asset information:

1. In the IBM Director Console Tasks pane, drag the **Asset ID** task onto the managed system. The serial number information for major system components is displayed in the Serialization page in the Asset ID window.
2. To configure system information, click the **System** tab. The following information is displayed on the System page:
   - Host name of the system
   - MAC address
   - Operating system
   - Universal unique identifier (UUID)
   - Remote Deployment Manager (RDM) profile associated with the system, if applicable

   The **RDM profile** field is the only field that you can modify.
3. To configure user information, click the **User** tab. You can provide the following information about the system user on the User page:
   - Name
   - Phone
   - Location
   - Department
   - Position
4. To configure lease information, click the **Lease** tab. You can specify the following information on the Lease page:
   - Lease start date
   - Lease end date
   - Term of the lease (in months)
   - Amount of the lease
   - Name of the company that leased the system

   If you specify an end date for the lease, a warning event is generated when the lease ends.
5. To configure asset information, click the **Asset** tab. You can provide the following information on the Asset page:
   - Date the system was purchased
   - Date the system was last inventoried
   - Asset number

If the system has a radio-frequency identification (RFID) number, it is displayed on this page.

6. To configure custom data fields, click the **Personalization** tab. You can provide custom information in up to five user-defined fields on the Personalization page.
7. To configure warranty information, click the **Warranty** tab. The information you specify here is collected during an inventory collection. You can view it by using the Inventory task, or manage the warranty information by creating a dynamic group. You can provide the following information on the Warranty page:
   - Duration of the warranty (months)
   - Cost of the system
   - End date for the warranty

   If you specify an end date for the warranty, a warning event is generated when warranty expires.
8. Click **Apply**.
9. Close the Asset ID window.

The information that you provided is written to the EEPROM. For systems that are not enabled for Enhanced Asset Information Area (EAIA), this information is written to asset.dat in the data directory. Asset.dat is deleted when IBM Director is uninstalled.

Creating an Asset ID profile

This topic describes how to create an Asset ID profile. You can apply an Asset ID profile to a group of managed systems. Applying this profile enables you to configure asset information, such as lease and warranty end dates, on multiple systems simultaneously.

Complete the following steps to create an Asset ID profile:

1. In the IBM Director Console Tasks pane, right-click **Asset ID**; then, click **Profile Builder**.
2. In the Asset ID: Profile Builder window, click **New Profile**. The Input window opens.
3. Type a descriptive name for the profile; then, click **OK**.
4. To permit an IBM Director user to modify the settings that you will apply with the profile, ensure that the **Enable Changes** check box in the Asset ID notebook window is selected. By default, this check box is selected.

   Important: If you choose to clear the **Enable Changes** check box, you will be unable to modify the settings through IBM Director or by any other operating system method.
5. To configure system information, click the **System** tab. In the **RDM Profile** field, type the profile name.
6. To configure user information, click the **User** tab. You can provide the following information about the system user:
   - Name
   - Phone
   - Location
   - Department
   - Position

To add data, select the check box next to the field; then, type the applicable information.

7. To configure lease information, click the **Lease** tab. You can specify the following information:
   - Lease start date
   - Lease end date
   - Term of the lease (in months)
   - Amount of the lease
   - Name of the company that leased the system

   To add data, select the check box next to the field; then, type the applicable information.

   Note: If you specify an end date for the lease, a warning event is generated when the lease ends.
8. To configure asset information, click the **Asset** tab. You can provide the following information:
   - Date the system was purchased
   - Date the system was last inventoried
   - Asset number

   To add data, select the check box next to the field; then, type the applicable information.
9. To configure custom data fields, click the **Personalization** tab. You can provide custom information in up to five user-defined fields. To add data, select the check box next to the field; then, type the applicable information.
10. To configure warranty information, click the **Warranty** tab. You can provide the following information:
    - Duration of the warranty (months)
    - Cost of the system
    - End date of the warranty

    To add data, select the check box next to the field; then, type the applicable information.

    Note: If you specify an end date of the warranty, a warning event is generated when warranty expires.
11. Click **Apply** to save all settings to the profile.
12. Click **Save Profile**.
13. In the Save profile window, click **Yes**.
14. Close the Profile Builder window. The profile is displayed in the Tasks pane under the Asset ID task.

Inventorying asset information

This topic describes how to create a custom inventory query that targets asset information.

Complete the following steps:

1. In the IBM Director Console Groups pane, right-click **Systems with Asset ID**; then, click **Perform Inventory Collection**. This procedure ensures that the IBM Director database contains the most recent Asset ID information.
2. In the Tasks pane, right-click **Inventory**; then, click **Create Custom Query**.
3. In the Inventory Query Builder window, in the Available Criteria pane, expand the following tree nodes:

- Asset ID
- Lease
- Personalized Data
- Serial Number
- User Details
- Warranty

4. Click the fields that you want to query and click **Add**. The criteria are displayed in the Selected Criteria pane.
5. Click **File** → **Save**.
6. In the Save Query window, type a name for the query and click **OK**.
7. Close the Inventory Query Builder window.
8. Drag the **Inventory** task onto **Systems with Asset ID**.
9. In the Available Queries pane, click the custom query that you created. The asset information is displayed in the Query Results pane.

Modifying an Asset ID profile

This topic describes how to modify an Asset ID profile.

You must have created an Asset ID profile previously.

Complete the following steps to modify an Asset ID profile:

1. In the IBM Director Console Tasks pane, right-click **Asset ID**; then, click **Profile Builder**.
2. In the Asset ID: Profile Builder window, select the profile that you want to modify in the top field of the left pane. The Asset ID notebook window opens in the right pane.
3. Modify the profile; then, click **Apply**.
4. Click **Save Profile**.
5. In the Save Profile window, click **Yes**.
6. Close the Profile Builder window.

Configuring a BladeCenter unit

This topic provides information about using the BladeCenter Configuration Manager to configure a BladeCenter unit.

You can use the BladeCenter Configuration Manager to create a profile that contains BladeCenter-unit configuration information. Once created, the profile can be reused and modified. You also can use the BladeCenter Configuration Manager to generate an XML configuration file that can be used with the IBM Director command-line interface, DIRCLI.

Applying a BladeCenter configuration profile to a chassis

After you configure a BladeCenter configuration profile through the BladeCenter Configuration Manager, you can apply it to a chassis.

To apply a BladeCenter configuration profile, complete the following steps:

1. In the Tasks pane of IBM Director Console, expand **BladeCenter Management** → **BladeCenter Configuration Manager**. If the Tasks pane does not appear, click the triangle along the right side of IBM Director Console, or from the menu bar click **View** → **Tasks Pane**.

2. From the Tasks pane, drag the profile that you want to apply to a chassis in the Group Contents pane or a group in the Groups pane.

Creating a BladeCenter configuration profile

Use the BladeCenter Configuration Manager to create a profile that includes configuration settings for the components of BladeCenter chassis, such as the IP addresses to use, as well as settings for switches and management modules.

To create a BladeCenter configuration profile, complete the following steps:

1. In IBM Director Console, click **Tasks** → **BladeCenter Management** → **BladeCenter Configuration Manager** → **Open**.

   Note: If you are creating the profile for a specific BladeCenter chassis, you can first select the BladeCenter chassis in the Group Contents pane. The configuration options are filtered for the selected hardware.
2. In the Quick Start window, click the **Create a new profile** icon.
3. In the Profile Parameters window, select the type of chassis that you are using, and type a name for the new profile.
4. Click **OK**.
5. In the Select Components window, select the check boxes next to the components that you want to include in the profile.
6. Click **OK**.
7. In the Configuration Manager Editor window, click a component to configure its settings.
8. After you have configured the settings for each component of the chassis, click **File** → **Save Profile**.

You can access the profile in IBM Director Console, by clicking **Tasks** → **BladeCenter Management** → **BladeCenter Configuration Manager** → *New Profile*. *New Profile* is the name you assigned the profile in the Configuration Manager Editor.

Editing a BladeCenter Chassis profile

Use the BladeCenter Configuration Manager to edit a profile that includes configuration settings for the components of BladeCenter chassis, such as the IP addresses to use, as well as settings for switches and management modules.

To edit a BladeCenter configuration profile, complete the following steps:

1. In IBM Director Console, click **Tasks** → **BladeCenter Management** → **BladeCenter Configuration Manager** → *Profile Name* → **Edit**. *Profile Name* is the name of the profile that you want to edit.
2. In the Configuration Manager Editor window, click a component to configure its settings.
3. After you make the configuration changes to the components that you want to edit, click **File** → **Save Profile**.

You can access the profile in IBM Director Console, by clicking **Tasks** → **BladeCenter Management** → **BladeCenter Configuration Manager** → *Profile Name*.

Exporting a profile to an XML file

You can use the BladeCenter Configuration Manager to export a BladeCenter configuration profile to an XML file. You might want to export the configuration profile so that you can back up the profile, transfer the profile to another IBM Director management server, or to further customize the information in an ASCII text editor.

Complete the following steps to export a BladeCenter configuration profile to an XML file:

1. In IBM Director Console, click **Tasks** → **BladeCenter Management** → **BladeCenter Configuration Manager** → **Open**.
2. In the Quick Start window, click the **Open an existing profile** icon.

   Note: You can also create a new profile to export.
3. On the Select a Profile window, select the profile you want to export to XML and click **OK**.
4. In the Configuration Manager Editor window, click **File** → **Export to XML**.
5. In the Save window, select a directory location, and type a file name for the XML configuration file. Click **Save**.

Importing an XML file to the BladeCenter Configuration Manager

You can use the BladeCenter Configuration Manager to import an XML configuration file and save it as a profile to the IBM Director Console.

To import an XML configuration file, complete the following steps:

1. In IBM Director Console, click **Tasks** → **BladeCenter Management** → **BladeCenter Configuration Manager** → **Open**.
2. In the Quick Start window, click the **Import an XML file** icon.
3. On the Open window, locate the XML file and click **Open**.
4. In the Configuration Manager Editor window, verify that your components are configured correctly by clicking each component.
5. To create a profile that is based on the XML file configuration settings, click **File** → **Save Profile**.

You can access the profile in IBM Director Console, by clicking **Tasks** → **BladeCenter Management** → **BladeCenter Configuration Manager** → *Profile Name*.

Managing event action plans

This topic provides information about managing event action plans.

Building an event action plan

This topic describes how to build an event action plan.

Creating an event action plan

This topic describes how to create an event action plan.

Complete the following steps to create a new event action plan:

1. In IBM Director Console, click **Tasks** → **Event Action Plan Builder**.

   The Event Action Plan Builder window contains three panes:

   Event Action Plans pane

   Lists event action plans. One default event action plan, Log All Events, is included with IBM Director. Also, if you used the Event Action Plan wizard to create an event action plan, that plan is listed.

Event Filters pane
Lists event filter types, with customized filters that are displayed under the applicable filter types. Expanding the **Simple Event Filter** tree displays, in addition to any customized simple event filters that were created, the preconfigured event type filters.

Actions pane
Lists event action types, with customized actions that are displayed under the event action types.

2. In the Event Action Plans pane, right-click **Event Action Plan**; then, click **New**.
3. In the Create Event Action Plan window, type a name for the plan and click **OK** to save it. The event action plan is displayed in the Event Action Plans pane.

Creating event filters

This topic describes how to create event filters. An event filter processes only the events that are specified by the filter and ignores all other events.

Complete the following steps to create event filters:

1. In IBM Director Console, click **Tasks** → **Event Action Plan Builder**.
2. In the Event Action Plan Builder window, double-click an event filter type in the Event Filters pane.

   Note: Alternatively, you can create an event filter for an event that has already occurred. In the IBM Director Tasks pane, double-click the **Event Log** task. In the Events pane, right-click an event; then, click **Create** and select one of the four event filter types.
3. In the Event Filter Builder window, complete the applicable fields for the event filter that you want to create. By default, the **Any** check box is selected for all filtering categories, indicating that no filtering criteria apply. Depending on the event filter type that you selected, the Event Filter Builder window contains different pages.
4. Click **File** → **Save As**.
5. In the Save Event Filter window, type a name for the filter. When you are naming an event filter, the name should indicate the type of events for which the filter is targeted and any special options that you have configured for the filter, including the time the filter is active and event severity. For example, an event filter for unrecoverable storage events that occur on a weekend should be named to reflect that.
6. Click **OK** to save the filter. The new filter is displayed in the Event Filters pane under the applicable filter type.
7. **Optional:** Create additional event filters for use in a single event action plan. Repeat step 1 through step 6.
8. In the Event Filters pane, drag the event filter onto the event action plan in the Event Action Plans pane. The event filter is displayed under the event action plan.
9. If you have created additional event filters that you want to use in this event action plan, repeat step 8.

When the event filter is completed, customize the event actions.

Activating an event action plan

This topic describes how to activate an event action plan.

Complete the following steps to associate the event filter and event actions with the event action plan and then activate it:

1. In the IBM Director Console Tasks pane, expand the **Event Action Plan** task. The event action plan that you created is displayed in the Event Action Plan tree.
2. Drag the event action plan from the Tasks pane onto the applicable managed object or objects or managed group. A confirmation message is displayed indicating that you have successfully applied the event action plan to the target object or group.

Enabling and viewing an event action history

This topic provides information about using the Event Action Builder task to enable and view an event action history.

By default, the event action history is disabled.

1. To enable the event action history, in the Event Action Plan Builder Actions pane, right-click the customized event action and click **Enable**.
2. To view the event action history, right-click the event action again and click **Show**.

Exporting an event action plan

This topic describes how to export an event action plan. You can export an event action plan for use on another management server or for viewing in HTML or XML.

With the Event Action Plan Builder, you can export event action plans to files. You can export event action plans from IBM Director Server to three types of files:

Archive
Copies the selected event action plan to a file that you can import to any management server.

You might export event action plans in archive format for two reasons:

- To move event action plans from one management server to another
- To back up event action plans on a management server

HTML
Creates a detailed listing of the selected event action plans, including their filters and actions, in a Hypertext Markup Language (HTML) format.

XML
Creates a detailed listing of the selected event action plans, including their filters and actions, in an XML format.

Complete the following steps to export an event action plan:

1. In IBM Director Console, click **Tasks** → **Event Action Plan Builder**.
2. In the Event Action Plan pane of the Event Action Plan Builder window, click the event action plan that you want to export.
3. Click **File** → **Export**, and select the type of file to which you want to export. Depending on which type of file you select, the applicable window opens (for example, if you select **Archive**, the Select Archive File for Export window opens).

4. Type a file name and, if necessary, change the location where you want to save the file. Click **OK** to export.

Importing an event action plan

This topic describes how to import an event action plan.

You can import event action plans using an archive file of an event action plan from another management server.

Complete the following steps to import an event action plan:

1. Transfer the archive file that you want to import to a drive on the management server.
2. In IBM Director Console, click **Tasks** → **Event Action Plan Builder**.
3. In the Event Action Plan Builder window, click **File** → **Import** → **Archive**.
4. In the Select File for Import window, select the archive file that you transferred in step 1.
5. Click **OK** to begin the import process.
6. In the Import Action Plan window, click **Import** to complete the import process. If the event action plan was previously assigned to managed objects or groups, you can preserve those assignments during the import process.

Restricting an event action plan

This topic describes how to restrict an event action plan. If an event action plan is restricted, all managed objects in a group to which the plan is applied must receive the event for the event action to occur. The default setting is **Unrestricted**.

You can restrict whether an event action plan applies both to events that are received by all managed objects in a group and to events that are received by one or more managed objects in the group, or just to the events that are received by all managed objects in the group.

Complete the following steps to restrict an event action plan:

1. In IBM Director Console, click **Associations** → **Event Action Plans**.
2. Expand the tree for the managed object or group that has the event action plan that you want to restrict applied to it.
3. Right-click the event action plan and click **Restricted**.

Using the Event Action Plan wizard

The Event Action Plan wizard allows you to combine event filters and event actions to easily create an event-action plan.

Complete the following steps to use the Event Action Plan wizard:

1. In IBM Director Console, click **Tasks** → **Event Action Plans** → **Event Action Plan Wizard**.
2. In the Event Action Plan Wizard, on the Name page, type a descriptive name for the event action plan you are creating.
3. Click **Next**.
4. In the Event Action Plan Wizard, on the Systems page, choose the systems to which you want to apply the event-action plan. Select the managed objects in the left pane, and click **Add** to move them to the **Selected** pane.
5. Click **Next**.

6. In the Event Action Plan Wizard, on the Event Filters page, select the check boxes adjacent to the types of events you want to monitor. You can select the following event filters:

   Hardware Predictive Failure Analysis® (PFA) events
   These events are sent when failure of a computer subsystem is imminent. Some of the subsystems for which IBM Director sends PFA events include hard disk drives, voltage regulation modules, power supplies, and thermal sensors.

   Environmental sensor events
   These events are sent when environmental sensors built into a system board detect that a manufacturer-defined threshold has been exceeded. For example, a fan shaft might have stopped.

   Storage events
   Storage events are sent when there has been a change in the status of a storage subsystem. Examples might be the failure of a hard disk drive or a logical drive that has been rebuilt.

   Security events
   These events are sent when there has been a change in the status of physical security. An example might be that someone has removed the system from the LAN or opened the system enclosure.

   IBM Director Agent offline
   An IBM Director Agent offline event is sent when IBM Director Agent stops running because the program either failed or was intentionally stopped.

   CPU Utilization
   CPU utilization events are sent when a user-defined threshold has been met or exceeded. For example, a group of processes that were not stopped by a parent process correctly, either in error or because of a virus, might cause long-term, increased CPU utilization. When you choose this event filter, the wizard creates a resource monitor to track microprocessor use in one or more selected systems.

   Note: The wizard cannot apply a resource monitor to groups or locked systems. You can create and apply a resource monitor using the Resource Monitor task in IBM Director Console. You must request access to the locked systems before you apply the resource monitor. After you apply the resource monitor to these groups or previously locked systems, the event filter will start to filter these events.

   Memory use
   Memory-use events are sent when a user-defined threshold has been met or exceeded. For example, a memory leak, caused by a process that has been stopped unexpectedly without freeing resources, might occur. When you choose this event filter, the wizard creates memory-used and memory-available resource monitors to track the memory use in one or more selected systems.

   Note: The wizard cannot apply a resource monitor to groups or locked systems. You can create and apply a resource monitor using the Resource Monitor task in IBM Director Console. You must request access to the locked systems before you apply the resource monitor. After you apply the resource monitor to these groups or previously locked systems, the event filter will start to filter these events.

You cannot apply memory-used and memory-available resource monitors to systems running NetWare or i5/OS.

7. Click **Next**.
8. In the Event Action Plan Wizard, on the Actions page, choose the event actions that you want performed when an event occurs. If you want to be notified by e-mail when an event occurs, select the **E-mail** check box in the Select the notification window. Then, configure the e-mail notification:
   a. In the **E-mail address** field, type the e-mail address to which the notification will be sent.
   b. In the **Reply-to address** field, type the e-mail address that will be displayed in the reply-to field of the e-mail.
   c. In the **SMTP server** field, type the host name or IP address of the Simple Mail Transfer Protocol (SMTP) server.
   d. In the **SMTP port** field, type the port number of the SMTP server. By default, the SMTP port is set to 25.
   e. In the **Subject** field, type the message that will be displayed in the subject-line of the e-mail. You can use variable such as *&type* and *&system*. For example, you might want to type the following string: `IBM Director alert:` *`&system &type`*. When the e-mail is generated, the name of the managed system is substituted for *&system*, and the type of event that occurred is substituted for *&type*.
   f. In the **Body of message** field, type the message that will be displayed in the body of the e-mail. You can use variable such as *&text*. For example, you might want to type the following string: *`&time &date &text`*. When the e-mail is generated, the body will contain the time and date the event occurred, as well as details about the event.

   The e-mail notification is configured. Note that *&type*, *&system*, *&time*, *&date*, and *&text* are event-data substitution variables.
9. If you want to be notified by pager when an event occurs, select the **Pager** check box in the Select the notification window. Then, configure the pager notification:
   a. From the **Serial port device name list**, select the name of the serial port device.
   b. In the **Network access number** field, type the telephone number that will be dialed when an event occurs.
   c. In the **Pager ID or PIN number** field, type the pager ID or personal identification number (PIN).
   d. In the **Message** field, type the message that will be sent when an event occurs.
   e. In the **Modem initialization string (optional)** field, type the modem initialization string.

   The pager notification is configured.
10. If you want to start a program as a result of the event, select the **Start program** check box. Then configure the start program instructions.
    a. Select the location where the program should be started.
    b. Type the **Host name** of the selected system.

    c. Type the working directory of the program that needs to start.
    d. Type the **Program name** of the application that needs to start.
11. Click **Next**.
12. In the Event Action Plan Wizard, on the Time Range page, choose the period of time over which you want to collect the events. You can select **All day** to enable the plan to be active all the time. Or you can select **Custom** to choose the time range for the plan to be active during specific days of the week.
13. Click **Next**.
14. In the Event Action Plan Wizard, on the Summary page, verify the details of the event action plan. If you need to make changes, click **< Back**.
15. Click **Finish**.

The event action plan is saved. To locate the new plan, in the IBM Director Console, click **Tasks** → **Event Action Plans** → *New plan*. *New plan* is the name you assigned in the wizard. The plan also appears in the Event Action Plan Builder, and under the systems or groups that the plan is applied to when Event Action Plans is selected in the IBM Director Console Associations menu.

After completing the wizard, you can edit the plan with the wizard, or you can use the Event Action Plan Builder to add additional actions not found in the wizard.

Note: If you add actions that are not available in the Event Action Plan Wizard to the plan using Event Action Plan Builder, the plan will become non-editable with the wizard.

Viewing event-action-plan associations

This topic describes how to view event-action-plan associations.

You can view which event action plans are applied to which managed objects and groups. In IBM Director Console, click **Associations** → **Event Action Plans**. If a managed object or group has an event action plan assigned to it, you can expand the managed object or group and expand the **Event Action Plan** folder to view the specific event action plans that are applied to the managed object or group.

To view which managed objects have event action plans applied to them, click **All Systems and Devices** in the Groups pane. If a managed object has an event action plan applied to it, you can expand the managed object in the Group Contents pane and expand the **Event Action Plan** folder to view the plans that are applied to the managed object.

To view which groups have event action plans applied to them, click **All Groups** in the Groups pane. If a group has an event action plan applied to it, you can expand the group in the Group Category Contents pane and expand the **Event Action Plan** folder to view the plans that are applied to the group.

Managing events

This topic provides information about working with the event log.

Changing the display options

This topic describes how to change the display options for the event log. You can set the time range, the number of events displayed, the color of specific events, and the total number of entries that are stored.

Changing the color of event types

This topic describes how to change the color of events displayed in the event log. You can set the color based either on severity or category.

Complete the following steps to change the color of an event type:

1. In the IBM Director Console Tasks pane, double-click the **Event Log** task.
2. In the Event Log window, click **Options** → **Customize Color** → **Severity** → **Critical**.
3. On the Swatches page, click the color in which you want Critical events to be displayed.
4. Click **OK**.

Changing the number of events displayed

This topic describes how to change the maximum number of events that are displayed in the event log. By default, this value is set to 100.

Complete the following steps to change the number of events that are displayed:

1. In the IBM Director Console Tasks pane, double-click the **Event Log** task.
2. In the Event Log window, click **Options** → **Set Log View Count**.
3. In the Set Log View Count window, type the number events to display in the event log in the **Change maximum number of entries** field.
4. Click **OK**.

Changing the number of events stored

This topic describes how to change the number of events that are stored in the event log.

Complete the following steps to change the number of events that are stored in the event log:

1. In IBM Director Console, click **Options** → **Server Preferences**.
2. In the Server Preferences window, click the **Event Management** tab.
3. In the Event Management page, in the **Enter the maximum number of event log entries** field, type the maximum number of event log entries.
4. Click **OK**.

Changing the time range

This topic describes how to change the time range of the event log. By default, the time range is set to 24 hours.

Complete the following steps to change the time range:

1. In the IBM Director Console Tasks pane, double-click the **Event Log** task.
2. In the Event Log window, click **Options** → **Set Time Range**.
3. In the Set Time Range window, type the number of units of time in the left field.
4. Select **hour(s)**, **day(s)**, or **week(s)** from the right list.
5. Click **OK**.

Exporting events

This topic describes how to export events.

You can export events that are displayed in the event log to an HTML, XML, or comma-separated value (CSV) file.

Complete the following steps to export an event from the event log:

1. In the Event Log window, click the event or events that you want to export to a file.
2. Click **File** → **Export**, and click the file format to which you want to export the event or events. The applicably named window opens.
3. In the **File Name** field, type a file name.
4. Click **OK**.

Viewing events

This topic describes how to view events that are stored in the event log.. You can view all events, events for a specific managed system or group, or events that are filtered using a specific criteria.

To view all events in the event log, in the IBM Director Console Tasks pane, double-click the **Event Log** task.

To view the events for a specific managed system or group, in the Event Log window, drag the **Event Log** task onto the managed system or group; the Event Log window for that managed system or group opens.

To view events by filter criteria, in the IBM Director Console Tasks pane, expand the **Event Log** task tree, and then double-click the filter for which you want to see all the events; the Event Log window opens, and only those events are displayed.

Managing files

You can use the File Transfer task to transfer files from one location to another and to synchronize files, directories, and drives. By default, the File Transfer task uses TCP. If you disable TCP session support on a managed system, the File Transfer task uses User Datagram Protocol (UDP).

Disabling TCP session support

This topic provides information about disabling TCP session support on a managed system. By default, the File Transfer task uses TCP. If you disable TCP session support on a managed system, the File Transfer task uses User Datagram Protocol (UDP).

Complete the following steps to disable TCP session support on a managed system:

1. Using a text editor, open one of the following files:

For managed systems running i5/OS	TCPIP.NETExt in the /QIBM/UserData/Director/classes/extensions directory
For managed systems running Linux	TCPIP.NETExt in the /opt/ibm/director/classes/extensions directory

For managed systems running UNIX	TCPIP.NETExt in /opt/ibm/director/classes/extensions directory
For managed system running Windows	TCPIP.INI file in the data directory

2. Modify the file:

For managed systems running i5/OS, Linux, or UNIX	1. Locate the following line in the file: `net.session.classname=com.tivoli.twg.netipc.TWGTCPSocketImplFactory` 2. Insert the following character at the beginning of this line: `#`
For managed system running Windows	Add the following line to the file: `SESSION_SUPPORT=0`

3. Save the file.
4. Stop and restart IBM Director Agent on the managed system.

   Note: (Managed systems running Windows) The file name is TCPIP.INI if TCPIP (All Adapters) is enabled in the Network Driver Configuration window. If an individual adapter is enabled, for example TCPIP1, you must create or edit a file named TCPIP1.INI for that adapter. Repeat the procedure for each individual adapter.

Synchronizing files, directories, or drives

This topic describes how to synchronize files, directories, or drives. When you synchronize files, directories, or drives, you replace the contents of the target file, directory, or drive with the contents of the source file, directory, or drive.

You can synchronize a source file, directory, or drive with as many target-managed system files, directories, or drives as you choose, but you must synchronize the file, directory, or drive on each managed system individually. You cannot synchronize multiple target managed systems from a source managed system at the same time.

Attention: Files or directories that are present only in the selected files, directories, or drives on the target managed system, but are not present in the selected files, directories, or drives on the source managed system, are deleted after synchronizing.

Complete the following steps to synchronize files, directories, or drives:

1. In the IBM Director Console Tasks pane, drag the **File Transfer** task onto the managed system (the target system) to which you want to transfer files. IBM Director takes a few seconds to query the files on the source system and on the target system.
2. In the File Transfer window, if you want the source to be identical to the target, in the Source File System pane, right-click the source; then, click **Synchronize from Target**. If you want the target to be identical to the source, in the Target File System pane, right-click the target; then, click **Synchronize from Source**.

3. If you receive a message indicating that the selected names are different, click **Yes** to continue. The selected files, directories, or drives are now synchronized.

Transferring files

This topic describes how to transfer files from one system to another.

Note: You can use the File Transfer task with only one system at a time. You cannot transfer files to multiple systems or to a group.

Complete the following steps to transfer files:

1. In the Group Contents pane of IBM Director Console, select the managed system (the target system) to which you want to transfer files.
2. In the IBM Director Console window, click **Tasks** → **File Transfer** → **Transfer File to / from** *target*. IBM Director takes a few seconds to query the files on the source system and on the target system.
3. To select another managed system, click **Other**.
4. In the Choose Target window, select the managed system that you want to transfer files to or from and click **OK**. The managed system is added to the target system list and is selected as the target system.

   Note: The Choose Target window does not display locked managed systems.
5. In the Source File System pane, expand the applicable hard disk drive. The contents of that drive are displayed, showing subdirectories and files.

   Note: You can switch between **Local** (management console) and **Director Server** (management server) by clicking the arrow in the Source File System pane and selecting your choice.
6. Select the item that you want to transfer; then, drag it onto the applicable location in the Target File System pane.

   Note: Using the wildcard function, you can filter which files are displayed in the Source File System or Target File System panes. When the File Transfer window opens, the **Name** field contains *.* by default.

Viewing hardware status

This topic provides information about using the Hardware Status task. You can view hardware status for a managed system or hardware-status group; you also can disable hardware-status events for a specific managed system or event type.

Clearing hardware-status events

This topic describes how to clear hardware-status events.

The managed system must have generated a hardware-status event.

Complete the following steps to clear a hardware-status event:

1. In the Group Contents pane, click the hardware-status icon that is displayed next to the managed system.
2. In the Status Groups pane of the Hardware Status window, perform one of the following actions:

To clear all events for the managed system	Right-click the managed system; then, click **Clear all Events**.

To clear a specific event type	Right-click the event type; then, click **Clear all Events**.

3. Close the Hardware Status window. In the Group Contents pane, the hardware-status icon is no longer displayed next to the managed system.

Disabling hardware-status events

This topic describes how to disable hardware-status event. When hardware-status events are disabled, whether for a managed system or a specific event type, IBM Director does not track hardware status for the specified managed system or event type.

Complete the following steps to disable hardware-status events:

1. In the Group Contents pane, click the hardware-status icon that is displayed next to the managed system.
2. In the Status Groups pane of the Hardware Status window, perform one of the following actions:

To disable all events for the managed system	Right-click the managed system; then, click **Ignore Events**.
To disable a specific event type	Right-click the event type; then, click **Ignore Events**.

3. Close the Hardware Status window. IBM Director will ignore all future events generated for the specified event type or managed system.

Enabling hardware-status events

This topic describes how to enable hardware-status events that you previously chose to ignore. By default, all hardware-status events are enabled.

Complete the following steps to enable previously-disabled hardware-status events

1. In the Group Contents pane, click the hardware-status icon that is displayed next to the managed system.
2. In the Status Groups pane of the Hardware Status window, perform one of the following actions:

To enable all events for the managed system	Right-click the managed system; then, click **Enable Events**.
To enable a specific event type	Right-click the event type; then, click **Enable Events**.

3. Close the Hardware Status window. IBM Director will monitor the future events generated for the specified event type or managed system.

Viewing hardware-status-group information

This topic describes how to view hardware-status information for a hardware-status group.

The hardware-status group must contain at least one managed system.

Complete the following steps to view hardware-status information for a hardware-status group:

1. In the lower-right corner of IBM Director Console, click the icon for the hardware-status group.
2. In the Status Groups pane of the Hardware Status window, expand the status-group for which you want to view information.

Viewing hardware-status information for a managed-system

This topic describes how to view hardware-status information for a managed system.

The managed system must have generated a hardware-status event.

The management server determines the hardware status of a managed system in one of two ways:

- It receives an alert from a managed system, which notifies it of a change in the health of a hardware component in that managed system. This is the normal mode of operation.
- It queries the managed system for its current hardware status. This occurs when the management server has been offline and unavailable to receive alerts or when a new managed system with non-normal status is first discovered. In some cases, the management server is unable to query a managed system with non-normal status that is first discovered due to platform-interface limitations. In either case, the managed system is monitored from discovery onward, but no initial baseline is gathered. This scenario is evident if the Hardware Status task is invoked against a managed system, but the task pane displays only the high-level categories of alert information (for example Environmental) with no component nodes displayed for any of the categories (for example Environmental > Power Supply).

Complete the following steps to view hardware-status information for a managed system:

1. In the Groups pane, click the hardware-status group that contains the managed system.
2. In the Groups Contents pane, click the hardware-status icon that is displayed next to the managed system.
3. In the Results pane of the Hardware Status window, click the row for the event for which you want information. The information is displayed in the Event Details section of the pane.

   Note: If the Events Detail section is not displayed, click **View** → **View detailed panel**.

Inventorying hardware and software

This topic provides information about using the Inventory task. You can use this task to inventory both hardware and software.

Adding an entry to the Inventory Software Dictionary

This topic describes how to add an entry to the Inventory Software Dictionary.

Complete the following steps to add an entry to the inventory-software dictionary:

1. In the IBM Director Console Tasks pane, right-click the **Inventory** task; then, click **Edit Software Dictionary**.

2. In the Inventory Software Dictionary Editor window, type a name to identify the entry in the **Title** field, which is located in the Entry Definition pane. In the **Entry Type** field, select the folder in the Available Entries pane in which the entry will be displayed. In the other fields, type the information that you want to use to identify the application.

   The **Title** and **Entry Type** fields are the only required fields. However, any information that you type in the Entry Description pane is displayed when you use the Inventory Query Browser window to view software information. It is not used as search criteria when you are collecting inventory data. The information that is entered in the **Associated Files** group box is used as the search criteria.
3. In the **Associated Files** group box, click **Add**.
4. In the Associated File Attributes window, click **Enter File Information Manually** or **Select File From List**; then, click **OK**.
5. In the Associated File Attributes window, if you clicked **Enter File Information Manually**, type the file name for which you want the inventory-software scanner to search. To further qualify the file, you can type a specific file size, range of file sizes, file date, or range of file dates. Click **OK**.

   If you clicked **Select File from List**, type the file name in the **File Name** field, or select the file. Click **OK**. The corresponding attributes are displayed in the **Associated Files** group box.
6. Optional: In the **Associated Files** group box, click **Edit** to change any of the attributes.
7. Optional: If you want to add more files to the software-dictionary entry definition, repeat step 3 through step 6.
8. Click the **Save Entry** icon to save the entry. The definition is added immediately to the software dictionary. The next time inventory data is collected, the data that you have provided in the Associated Files pane is used as criteria in locating the file.

Creating a custom inventory collection

You can create a customized inventory collection to collect only a specific set of inventory and apply it to a managed object or group of managed objects.

Complete the following steps to create a customized inventory collection:

1. From the IBM Director Console window, select **Tasks** → **Inventory** → **Custom Collections**.
2. On the Custom Collection: New window, complete the required information to define the collection criteria and post-collection actions for your customized inventory collection.
3. Click **OK**.

Creating a custom inventory query

This topic describes how to create a custom inventory query.

Complete the following steps to create and use a custom query to view inventory data:

1. In the IBM Director Console Tasks pane, right-click the **Inventory** task; then click **Build Custom Query**.
2. In the Available Criteria pane of the Inventory Query Builder: New window, drag the data items that you want to add to the query onto the Selected

Criteria pane. The order of the criteria in the Selected Criteria pane is the order in which the criteria will be displayed in the Inventory Query Browser window.

3. Click **File** → **Save As** to save the query. The new query is displayed under the Custom folder in the Available Queries pane of the Inventory Query Browser window.

Editing a custom inventory query

This topic describes how to edit a custom inventory query.

Complete the following steps to edit a custom query:

1. In the IBM Director Console Tasks pane, double-click the **Inventory** task.
2. In the Available Queries pane of the Inventory Query Browser window, expand the **Custom** folder to view the list of custom queries. Right-click the query that you want to edit, and click **Modify**.
3. Add or delete criteria in the Selected Criteria pane.
4. Click **File** → **Save** to save your changes and update the query.

   Note: If you edit and save a custom query, the Inventory task might not be able to interpret the new query, and the saved query might not be displayed in the Available Queries pane of the Inventory Query Browser window. To view the query, restart the Inventory task to open the Inventory Query Browser: All Systems and Devices window. The saved query is displayed in the Available Queries pane.

Exporting inventory query results to a file

This topic describes how to export the results of an inventory query to a file. You can export inventory query results in CSV, HTML, or XML format.

Complete the following steps to export inventory query results:

1. In the IBM Director Console Tasks pane, double-click the **Inventory** task.
2. In the Inventory Query Browser window, click the query that you want to export.
3. Click **File** → **Export** and click the format to which you want to export the results.
4. Type a file name and specify the location where you want to save the file; then, click **OK**.

Managing inventory monitors

You can enable, disable, or remove inventory monitors on managed systems in the Applied Inventory Monitors window.

Complete the following steps to manage inventory monitors that are applied to managed systems:

1. From the IBM Director Console window, select **Tasks** → **Inventory** → **Inventory monitors** → **View Applied Inventory Monitors**.
2. In the Applied Inventory Monitors window, select the monitor that you want to change, and click one of the following management options:
   - Remove
   - Enable
   - Disable
3. Click **OK**.

Monitoring inventory changes

You can use inventory monitors to track when hardware or software is added, changed, or removed on managed systems.

Complete the following steps to monitor inventory changes:

1. From the IBM Director Console window, select **Tasks** → **Inventory** → **Inventory Monitors** → **Create**.
2. On the Inventory Monitor: New window, name the monitor, and select the inventory criteria changes for which you want to be notified.
3. Click **OK**.
4. When the inventory monitor is created, you can apply it to a system or group by dragging the monitor to the system or group for which you want to monitor inventory.

Tracking software suites

This topic describes how to track software suites, such as Microsoft Office. To track software suites, you must create entries in the Inventory Software Dictionary, and then create a dynamic group that contains the managed systems on which the software suite is installed.

The Inventory Software Dictionary finds a match for an entry definition only if all associated files for the entry are in the same directory. To locate product suites (such as Microsoft Office) that might not have all applications in the same directory, you can create separate dictionary entries for each application in the suite and then create a dynamic group to display all managed systems and devices that are found with the specified application files.

Complete the following steps to create separate dictionary entries and to create a dynamic group:

1. In the IBM Director Console Tasks pane, right-click the **Inventory** task; then, click **Edit Software Dictionary**.
2. In the Inventory Software Dictionary Editor window, in the Entry Definition pane, use the **Title** and **Entry Type** fields to identify and classify each entry that you create in the inventory-software dictionary. You also can complete the other fields as needed.
3. In the **Associated Files** group box, click **Add**.
4. In the Associated File Attributes window, click **Enter File Information** or **Select File From List**; then, click **OK**. The easiest method is to select the file from a list. When you finish selecting the file name, the corresponding attributes are displayed in the **Associated Files** group box.
5. **Optional:** Click **Edit** to change any of the attributes.
6. **Optional:** If you want to add more files to the definition, repeat steps 3 through 5.
7. Click the **Save Entry** icon to save your software-dictionary entry. You have now created one entry that identifies the file (or set of files, if you specified more than one file) that corresponds to one application in a single directory.
8. Click **File** → **New** to add another software-dictionary entry. Repeat steps 2 through 7 for each software-dictionary entry you want to create, and then click **File** → **Close** to close the Inventory Software Dictionary Editor window.
9. To ensure detection of the installed software packages, perform an inventory collection on the managed system or device with the specific software that is installed on it.

10. In the IBM Director Console Groups pane, right-click anywhere *except* on an entry and click **New Dynamic**.
11. In the Available Criteria pane of the Dynamic Group Editor window, expand the **Inventory** tree. To display the list of software-dictionary entries from which you can create a new dynamic group, expand the **Software** tree, and then expand the **Program Title** tree.
12. Locate and click the first software-dictionary entry that you created; then, click **Add** to add the entry to the Selected Criteria pane.
13. Locate and click the second software-dictionary entry that you created; then, click **Add** to add it to the Selected Criteria pane.
14. In the Choose Add Operation window, click **All true (AND)** to create a group that includes a managed system or device only if all of the software-dictionary entries that you selected are located on that managed system or device.
15. Locate and add the rest of the entries that you created. For each subsequent entry that you add to the Selected Criteria pane, select the **All true (AND)** option when prompted.
16. When you have finished building your group of entries, click **File** → **Save As**.
17. In the Save As window, type the name that you want to display in the Groups pane. Click **OK**.
18. Click **File** → **Close Group Editor** to close the Dynamic Group Editor window.
19. Click the new group in the IBM Director Console Groups pane. The managed systems and devices that meet the search criteria for the software entries that you created are displayed in the Group Contents pane. All entries must be present on the managed system or device for the managed system or device to be displayed.

Viewing inventory data

This topic describes how to view inventory data using either a predefined or a custom query.

Complete the following steps to view inventory data:

1. In the IBM Director Console Tasks pane, drag the **Inventory** task onto a managed system or group.

   The Inventory Query Browser window has two panes: Available Queries and Query Results. The Available Queries pane automatically displays predefined queries that are included in IBM Director and any custom queries that you have created previously. In the Query Results pane, you can view the details of the query for each selected managed system.
2. In the Available Queries pane, expand one of the following folders:

To use a custom query	Custom
To use a predefined query	Standard

3. Click a query. The results for each managed system are displayed in a table in the Query Results pane. If no information is currently available about that query, a message is displayed.

Managing service processors

You can use Management Processor Assistant and the Server Configuration Manager to manage service processors in xSeries and Netfinity servers.

Applying a server configuration profile to a managed object

After you configure a server configuration profile through the Server Configuration Manager, you can apply it to an xSeries or Netfinity server.

To apply a server configuration profile, complete the following steps:

1. In the Tasks pane of IBM Director Console, expand **Server Configuration Manager**. If the Tasks pane does not appear, click the triangle along the right side of IBM Director Console, or from the menu bar click **View** → **Tasks Pane**.
2. From the Tasks pane, drag the profile that you want to apply to a managed object in the Group Contents pane or a group in the Groups pane.

Creating a server configuration profile

Use the Server Configuration Manager to create a profile that includes configuration settings for service processors in xSeries and Netfinity servers. You can also reserve IP addresses for service processors.

To create a server configuration profile, complete the following steps:

1. In IBM Director Console, click **Tasks** → **Server Configuration Manager** → **Open**.

   Note: If you are creating the profile for a specific server, you can first select the server in the Group Contents pane. The configuration options are filtered for the selected hardware.
2. In the Quick Start window, click the **Create a new profile** icon.
3. In the Profile Parameters window, select the type of chassis that you are using, and type a name for the new profile.
4. Click **OK**.
5. In the Select Components window, select the check boxes next to the components that you want to include in the profile.
6. Click **OK**.
7. In the Configuration Manager Editor window, click a component to configure its settings.
8. After you have configured the settings for each component of the chassis, click **File** → **Save Profile**.

You can access the profile in IBM Director Console, by clicking **Tasks** → **Server Configuration Manager** → *New Profile*. *New Profile* is the name you assigned the profile in the Configuration Manager Editor.

Editing a server configuration profile

Use the Server Configuration Manager to edit a profile that includes configuration settings for service processors in xSeries and Netfinity servers. You can also edit the IP addresses that are reserved for the service processor.

To edit a server configuration profile, complete the following steps:

1. In IBM Director Console, click **Tasks** → **Server Configuration Manager** → *Profile Name* → **Edit**. *Profile Name* is the name of the profile that you want to edit.
2. In the Configuration Manager Editor window, click a component to configure its settings.
3. After you make the configuration changes to the components that you want to edit, click **File** → **Save Profile**.

You can access the profile in IBM Director Console, by clicking **Tasks** → **Server Configuration Manager** → *Profile Name*.

Exporting a server configuration profile to an XML file

You can use the Server Configuration Manager to export a server configuration profile to an XML file. You then can edit the XML file in an ASCII text editor and use DIRCLI, the IBM Director command-line interface, to edit the profile. You also can use DIRCLI to apply the profile to one or more xSeries or Netfinity servers.

To export a server configuration profile to an XML file, complete the following steps:

1. In IBM Director Console, click **Tasks** → **Server Configuration Manager** → **Open**.
2. In the Quick Start window, click the **Open an existing profile** icon.
3. On the Select a Profile window, select the profile you want to export to XML and click **OK**.
4. In the Configuration Manager Editor window, click **File** → **Export to XML**.
5. In the Save window, select a directory location, and type a file name for the XML configuration file. Click **Save**.

Importing an XML file to the Server Configuration Manager

You can use the Server Configuration Manager to import an XML configuration file and save it as a profile to the IBM Director Console.

To import an XML configuration file to the Server Configuration Manager, complete the following steps:

1. In IBM Director Console, click **Tasks** → **Server Configuration Manager** → **Open**.
2. In the Quick Start window, click the **Import an XML file** icon.
3. On the Open window, locate the XML file and click **Open**.
4. In the Configuration Manager Editor window, verify that your components are configured correctly by clicking each component.
5. To create a profile that is based on the XML file configuration settings, click **File** → **Save Profile**.

You can access the profile in IBM Director Console, by clicking **Tasks** → **Server Configuration Manager** → *Profile Name*.

Powering on and off scalable partitions

You can use IBM Director Console to power on and power off scalable partitions on xSeries 460 servers.

Power operation that are performed on managed objects that represent scalable partitions use out-of-band communication. Power operations that are performed on managed-system objects created from powered-on scalable partitions use in-band communication to power off the scalable partition.

Restriction: The out-of-band power operations in IBM Director 4.22 are only for use by xSeries 460 servers. Other supported servers (such as xSeries 455 and xSeries 445 servers) should install and use Scalable Systems Manager (SSM) 4.20 if needed.

Powering on a scalable partition

When you power on a scalable partition, the servers that are part of the scalable partition are powered on, and the operating system is started on the scalable partition. If IBM Director Agent is installed on the scalable partition, it is started also.

Before you power on a scalable partition, make sure that the primary scalable node is unlocked in IBM Director and that all scalable nodes in the scalable partition have been discovered by IBM Director.

Complete the following steps to power on a scalable partition:

1. Request access to the primary scalable node and unlock the physical platform:
   a. In the Group Contents pane of IBM Director Console, right-click the physical-platform icon and click **Request Access**.
   b. Type a valid login ID and password and click OK.
2. In the Group Contents pane, expand the tree structure for the scalable system that contains the scalable partition.
3. Right-click the inactive scalable partition that you want to power on; then, click **Power On**.

Shutting down and powering off a scalable partition

When a scalable partition is powered on, you can use the IBM Director Console to shut down and power off the scalable partition.

IBM Director notifies the operating system that the scalable partition will be shut down and powered off in one of the following ways:

- If the device driver for the service processor on the primary scalable node is available, the operating system attempts to exit from running applications before it shuts down. Then, IBM Director powers off the servers that are represented as scalable nodes.
- If the device driver for the service processor on the primary scalable node is not available, the scalable partition is immediately powered off and there is no attempt to exit from running applications. In this case, the operating system can display only a message that it is shutting down. It then flushes its disk caches before the servers that are represented as scalable nodes in the scalable partition are physically powered off. Application processes running on the system are not shut down.

Complete the following steps to shut down the operating system and power off all servers that are represented as scalable nodes:

1. From IBM Director Console, in the Group Contents pane, expand the tree structure for the scalable system that contains the scalable partition.
2. Right-click the active scalable partition that you want to power off; then, click **Shutdown and Power Off**.
3. Click **Execute Now**.

Powering off a scalable partition

When a scalable partition is powered on, you can use IBM Director Console to immediately power off the scalable partition.

When you power off a scalable partition, IBM Director notifies the operating system that the scalable partition will be powered off immediately. The operating system then displays a message that it is shutting down. It flushes its disk caches before the servers that are represented as scalable nodes in the scalable partition are physically powered off. Application processes that are running on the system are not shut down in an orderly way. After approximately 10 seconds, IBM Director physically powers off the server.

Complete the following steps to immediately power off a scalable partition:

1. From IBM Director Console, in the Group Contents pane, expand the tree structure for the scalable system that contains the scalable partition.
2. Right-click the scalable partition that you want to power off; then, click **Power Off Now**.
3. Click **Execute Now**.

Managing process monitors

This topic provides information about managing process monitors. You can use process monitors to generate events when an application process starts, stops, or fails to start.

Applying a process monitor

This topic provides information about applying a process monitor.

To apply a process monitor, complete the following steps:

1. Drag the process monitor onto the managed system that has a process that you want to monitor.
2. In the Process Monitor window, click **Execute Now** or click **Schedule** to schedule it for a later time.

Creating a process monitor

This topic provides information about creating a process monitor.

Complete the following steps to create a process monitor:

1. In the IBM Director Console Tasks pane, expand the **Process Management** task.
2. Double-click the **Process Monitors** subtask.
3. In the Process Monitors window, type the executable file name of the application process that you want to monitor.
4. Select any combination of the **Start**, **Stop**, and **Fail** check boxes, to specify which action or actions you want to monitor.
5. If you selected the **Fail** check box, type a timeout setting. This setting is the number of seconds that the process monitor will wait for the application process to start before generating a fail event.
6. To monitor additional processes with the same Process Monitors subtask, click **Edit** → **New Row**.
7. Repeat steps 3 through 6 until you have listed the executable file names of all the processes that you want to monitor.
8. Click **File** → **Save As** to save the process monitor.
9. In the Save As window, type a name to identify the process monitor; then, click **OK**. The new process monitor is displayed as a subtask under the Process Monitors task in IBM Director Console.

Removing all process monitors

This topic provides information about removing all process monitors from a managed system.

Complete the following steps to remove all process monitors:

1. In the IBM Director Console Tasks pane, expand the **Process Management** task.
2. Drag the **Remove Process Monitors** subtask onto the managed system from which you want to remove all process monitors.
3. Click **Execute Now** or click **Schedule** to schedule the removal for a later time.

Removing an individual process monitor

This topic provides information about removing an individual process monitor.

Complete the following steps to remove an individual process monitor:

1. In the IBM Director Console tasks pane, expand the **Process Management** task.
2. Drag the managed system from which you want to remove the process monitor onto the **Process Monitors** task.
3. In the Process Monitor window, right-click the process monitor that you want to remove and click **Delete Row**.
4. Click **File** → **Save**.
5. Click **Yes** to confirm. The monitor is removed from the managed system.

Viewing process monitors

This topic provides information about viewing process monitors.

To view a list of the process monitors that are running on a managed system, drag the **Process Monitors** task onto the managed system. The Process Monitors window opens, and the list of process monitors that are running on that managed system is displayed.

Monitoring Microsoft clusters

You can use the Microsoft Cluster Browser task to view the structure, nodes, and resources that are associated with a Microsoft Cluster Server (MSCS) cluster.

You can determine the status of a cluster resource and view the associated properties of the cluster resources. The Microsoft Cluster Browser does not display the status of a cluster as a whole but displays the individual cluster resource statuses.

To monitor a cluster, complete the following steps:

1. In the Group Contents pane of IBM Director Console click the title and select **Clusters and Cluster Members**.

   Note: Only cluster names can be browsed. If a managed system is a member node of a cluster, the following message displays:

   ```
   The targeted system does not support this task.
   ```
2. Select the cluster for which you want to view information.
3. Click **Tasks** → **Microsoft Cluster Browser** → **Microsoft Cluster Browser:** *Target*. *Target* is the cluster you selected.

4. The Cluster Browser window is displayed with the cluster you selected in the Clusters pane. The cluster is displayed as the root of a hierarchical structure.
5. To view cluster status and description, in the Clusters pane, double-click the cluster.
6. To view information about the resources that are assigned to the cluster, in the Clusters pane, expand the **Properties** tree and double-click the applicable resource.

Configuring network settings

This topic provides information about configuring IP addresses for managed systems and for the BladeCenter chassis.

Changing the IP address of a BladeCenter chassis manually

This topic describes how to change the IP address of the BladeCenter chassis manually.

Complete the following steps to change the IP address of the BladeCenter chassis manually:

1. Using a crossover cable, connect a system to the external port of the management module.
2. Change the IP address of the non-chassis system to an address on the 192.168.70.0 subnet.
3. Using the non-chassis system, open a Web browser.
4. In the **Address** field or the **Location** field, type the following address and press **Enter**:

   ```
   http://192.168.70.125
   ```

   A password window opens.
5. In the applicable fields, type the default user name (`USERID`) and password (`PASSW0RD`) for the BladeCenter management module. (Use uppercase letters and a zero, not the letter O.)
6. Click **OK**.
7. In the BladeCenter Management Module window, click **Continue**.
8. In the left pane of the System Status Summary window, click **Network Interfaces**.
9. In the **DHCP** field of the Management Module Network Interfaces window, click **Disabled—Use static IP configuration**.
10. In the **IP address** field, type a valid IP address on the same subnet as the management server.
11. In the **Subnet mask** field, type a valid subnet mask.
12. In the **Gateway address** field, type a valid gateway address.
13. Click **Save**.
14. In the left pane, click **Restart MM**.

Configuring IP addresses

This topic describes how to configure IP addresses for a managed system.

Complete the following steps to view and configure IP addresses:

1. In the IBM Director Console Tasks pane, drag the **Network Configuration** task onto a managed system or group.
2. In the Network Configuration window, click the **IP Address** tab.
3. In the IP Address page, in the **Adapter** field, select the network adapter.
4. Obtain an IP address either automatically or manually. To obtain an IP address automatically from a DHCP server, click **Use DHCP for automatic configuration**. To manually configure the IP address, complete these steps:
   a. Click **Configure manually**.
   b. In the **IP Address** field, type the IP address of the managed system.
   c. In the **Subnet Mask** field, type the subnet mask that is used by the managed system.
   d. In the **Default Gateway** field, type the gateway address that is used by the managed system.
5. Click **Apply** to save the changes.
6. Click **File** → **Close**.

Running a command-line program on a managed system

This topic provides information about using the Process Management task to run a command-line program on a managed system.

You can run a command in the following ways:

- As a process task
- From the Process Management window by clicking **Actions** → **Execute Command**

Anonymous command restrictions apply to each method of running the commands.

Creating a process task

This topic provides information about creating a process task. You can use a process task to run a command-line program on a managed system. A process task can be run immediately, scheduled to run at a specific time, or scheduled to run at regular intervals.

Note: Remember that anything a system-account user can do from a command line can be done to the system, regardless of the user who is logged in to the managed system.

Complete the following steps:

1. In the IBM Director Console Tasks pane, expand the **Process Management** task.
2. Double-click the **Process Tasks** subtask.
3. In the Process Task window, specify information about the command-line operation:
   a. In the **Command** field, type the fully qualified file name and command syntax. Consider the following information:
      - For managed systems running i5/OS, the command is run in the QShell environment.
      - For managed systems running Windows, preface the command with the following string to ensure that it runs in a Windows command-shell window:

```
cmd /c
```

b. Select the **Log** check box if the command produces text-based output, for example, a directory listing; then type a timeout value, in seconds. Make sure that the timeout value is long enough.

c. Optional: If you want to run the process using an alternate user account and override the default user ID, you can specify a user ID and password in the **Login** group box.

Note: (Managed systems running i5/OS only) Commands cannot be performed using the default user ID that is shipped with IBM Director. For a process task to run successfully to i5/OS systems, the process task must specify a valid user ID and password, or the managed system to which it is being run must be configured to specify a user that has the authority to run the process task. You can register that user in the **IBM Director Agent default user** function ID on the managed system using Application Administration in iSeries Navigator.

Note: You either must specify an alternate user ID, or remove the default user ID from the registered function and add a new default user ID that has the required authority to perform the command.

4. Click **File** → **Save As** to save the process task.
5. In the Save As window, type a name.

   Note: The name for a process task should include the following information:
   - Type of process task that is to be run
   - Name of the process task that is to be run
   - Types of managed systems with which the process task will work correctly
6. Click **OK**. The new process task is displayed under **Process Tasks** in IBM Director Console.

Running a process task

This topic provides information about running a process task.

Complete the following steps to run a process task:

1. In the IBM Director Console Tasks pane, expand the **Process Management** task.
2. Drag the process task onto the managed system on which you want to run the process task.
3. In the Process Task window, click **Execute Now** or click **Schedule** to run the process task at a later time.

Restricting anonymous command execution

This topic provides information about restricting anonymous command execution. By default, commands are run on the managed system as either system account (Windows) or root (Linux). You can restrict anonymous command execution by disabling this function and always requiring a user ID and password.

Restricting anonymous command execution on Linux

This topic provides information about restricting anonymous command execution on managed systems running Linux.

Complete the following steps:

1. Change to the directory in which IBM Director Agent is installed. If you installed IBM Director in the default directory, the directory name is opt/ibm/director/data.
2. From a command-line prompt, type the following string, and then press **Enter**:
   ```
   vi ProcMgr.properties
   ```
3. Change the value of `RestrictAnonCmdExe` to `False`.
4. Save the file. The change takes effect immediately.

Restricting anonymous command execution on Windows

This topic provides information about restricting anonymous command execution on managed systems running Windows.

Complete the following steps:

1. On a Windows system, type `regedit` at a command line, and then press **Enter**:
2. Navigate to the registry entry HKEY_LOCAL_MACHINE\SOFTWARE\IBM\Director\CurrentVersion.
3. Double-click **RestrictAnonCmdExec**.
4. In the **Value data** field, change **0** to **1**.
5. Click **OK**. The change takes effect immediately.

Managing CIM classes and methods

You can use the Common Information Model (CIM) Browser task to find in-depth information for problem determination, or to develop a systems-management application using the CIM layer.

Changing the value of a CIM-class property

This task describes how to change the value of a CIM class property.

Attention: Do not change the value of a property unless you are thoroughly familiar with the structure and manipulation of CIM data. Setting a property value incorrectly can cause unpredictable results on the target system.

Complete these steps to change the value of a CIM-class property:

1. In the CIM Browser window, navigate to the class instance for which you want to change a property value. In the lower-right pane, the Properties page displays the class-instance properties.
2. Right-click the property row that you want to change.
3. Click **Set value**.
4. In the Set Value window, type the new value.
5. Click **OK**. If IBM Director cannot change the value on the target system, a message indicates the failure.

Creating a CIM-class method shortcut

This task describes how to create a shortcut for a CIM-class method.

Complete these steps:

1. In the CIM Browser window, navigate to the class that has the method for which you want to create a shortcut. In the lower-right pane, click the **Methods** tab to display the associated methods.
2. Right-click a method and click **Execute**.

3. If the method has any input arguments, one or more **Input** fields are displayed in the Execute Method window. Type the arguments in these fields.
4. Click **Save**.
5. Either type a new name or keep the default name.
6. Click **OK**. The new shortcut is displayed under **CIM Browser** in the IBM Director Console Tasks pane.

Creating a CIM-class shortcut

This task describes how to create a shortcut for a CIM class.

Complete these steps to create a shortcut for a specific CIM class:

1. In the CIM Browser window, navigate to the class for which you want to create a shortcut.
2. Right-click the class name and click **Create browser task for class**. A window opens with the name of the class entered as the default name.
3. Type a new name, or keep the default name.
4. Click **OK**. The new subtask is displayed under **CIM Browser** in the IBM Director Console Tasks pane.

Running a CIM-class method

This task describes how to run a CIM-class method.

Attention: Do not execute a method unless you are thoroughly familiar with the structure and manipulation of CIM data. Executing a method incorrectly can cause the connection to the target system to be lost.

Complete these steps to execute a method for a CIM-class instance:

1. In the CIM Browser window, navigate to the class instance that has the method that you want to execute.
2. In the lower-right pane, click the **Methods** tab. The associated methods are displayed.
3. Right-click a method and click **Execute**.
4. In the Execute Method window, if the method has any input arguments, type the arguments in the input fields.
5. Click **Execute** to run the method. If IBM Director cannot run the method on the target system, a message indicates the failure.

Running a CIM-class method shortcut

This task describes how to run a shortcut for a CIM-class method.

To run a CIM-class method using a shortcut, drag the shortcut onto a CIM-enabled managed system that supports the method that you want to run.

Running a CIM-class shortcut

This task describes how to run a shortcut for a CIM class.

To run a CIM class using a shortcut, drag the shortcut onto a CIM-enabled managed system that has the instance, properties, and methods that are associated with those in the shortcut.

Viewing CIM classes and methods

This task describes how to display a list of CIM classes and methods.

If one or more managed systems are not configured for CIM data, a message is displayed indicating that the target systems do not support the task. If a managed system is inaccessible (for example, if it is offline) the CIM Browser window opens, but you cannot expand the CIM tree of the managed system.

Complete these steps to view CIM classes and methods that are available for one or more managed system from the IBM Director Console:

1. If you want to display information for multiple managed systems, select the managed systems for which you want to view information.
2. From the Tasks pane, drag the **CIM Browser** task to any selected managed systems.
3. In the CIM Browser window, turn on or turn off the displaying of managed-system classes by right-clicking a managed system and clicking **Show System Classes**. A managed-system class is indicated by a double underscore that precedes the class name. Also, you can expand the managed-system tree to display the CIM name spaces of the managed system and then expand a name space to display its classes. The name space that contains the IBM-specific classes is root\IBMSD.
4. View an instance of a class by clicking the class name.

   If an instance of the class is found, the right pane splits. In the lower-right pane, the associated properties and methods are displayed under the **Properties** and **Methods** tabs. All classes can have associated properties or methods.

   Note: Displaying instances of some CIM classes might cause excessive resource usage on the managed system. The resource usage continues until all instances have been opened, even if the request is canceled. Therefore, avoid attempting to view instances of root\cimv2:CIM_DirectoryContainsFile and root\cimv2:Win32_Subdirectory on managed systems running Windows, or root/ibmsd for IBMPSG on managed systems running Linux.

Creating a CIM dynamic group

You can use the Task Based Group Editor to create new dynamic group filters that are based on combinations of tasks that apply to managed systems that are CIM-enabled.

Complete the following steps to create a dynamic group for CIM-enabled systems:

1. Right-click in the Groups pane of the Director Console to display the context menu.
2. Click **New Task Based Group**.
3. In the Available Resources pane of the Task Group Editor window, click **CIM Browser** and then click **Add to add the selection to the Selected Criteria pane**. Selecting CIM Browser creates a filtering criteria for managed systems that are CIM-enabled.
4. Click **Save As** to save the new group with a name of your choosing.
5. Type a descriptive name for the group (for example, “CIM-enabled systems”).
6. Click **Close Group Editor** to save your group and exit the dialog.
7. Refresh the IBM Director Console with a discovery operation, and the new group is displayed in the Groups pane.

8. Click your new group to see which managed systems match the CIM criteria. CIM-enabled systems, if discovered, are listed in the Groups pane.

Managing racks

You can use the Rack Manager task to group your xSeries and Netfinity equipment in rack suites, virtual racks, and associate an unrecognized managed system or device with a predefined component of a similar size.

Starting the Rack Manager task

This topic describes how to start the Rack Manager task in IBM Director.

To start the Rack Manager task, in the IBM Director Console Tasks pane, drag the **Rack Manager** task onto a managed system or group.

The left pane displays the Topology view by default. You can change the left pane view by clicking the list above the left pane. Four views are available:

Topology
Displays the Racks tree, which contains any racks that have been created, and the Floor tree, which contains all managed systems and devices that have not been added to a rack. A BladeCenter unit is displayed as a Chassis tree. Expanding a **Chassis** tree displays all blade servers in that chassis.

Components
Displays the predefined components that are available for association and for inclusion in a rack.

Cluster
Displays clusters and cluster members, if any cluster components exist. If there are no cluster components, this option is disabled.

Multi-Node Systems
Displays complexes, partitions, virtual nodes, and I/O expansion units, if any exist. If there are none, this option is disabled.

Information in all of these views is displayed in a tree structure.

The information in the right pane can be displayed in two ways:

Rack view
The right pane is subdivided into two subpanes. The information in the upper-right subpane displays rack information graphically. For example, if a rack component has a hardware-status alert, the rack component is outlined in red (for critical alert), yellow (for warning alert), or blue (for informational alert). The lower-right subpane displays the properties of the component that is selected in the upper pane or the left pane. If the inventory-collection function of IBM Director does not recognize the managed system or device that is selected in the left pane, `Unknown` is displayed for some of the properties that are displayed in the lower-right pane.

Table view
The right pane displays rack information, such as position in rack, hardware status, and state, in a table structure.

To view the rack information graphically, click **View** → **Rack view**. To view the rack information in table structure, click **View** → **Table view**.

Starting a component association

This topic describes how to associate an unknown managed system or device with a predefined component in IBM Director.

Some managed systems and devices are not rack mountable until they are associated with predefined components. This association occurs when the inventory-collection function of IBM Director does not recognize the managed system or device.

Complete the following steps to associate an unknown managed system or device with a predefined component:

1. From the **Floor** tree in the Topology view, right-click the managed system or device and click **Associate**.
2. In the Associate window, expand the applicable tree and click the predefined component type that most closely resembles the managed system or device in size.
3. Click **OK**. The properties of the component that was associated with that managed system or device are displayed in the lower-right subpane.

You can change the association of a component by first canceling the component association and then associating it with a different predefined component.

Canceling a component association

This topic describes how to cancel a component association in IBM Director.

You might want to cancel a component association in any of the following situations:

- You have made an incorrect component association.
- Inventory collection on the component has been performed successfully.
- The association is no longer valid.

To cancel the association of a managed system or device with a predefined component, in the Topology view left pane, right-click the component that you want to disassociate, and click **Disassociate system**. The component information in the lower-right subpane reverts to the information that was received initially through the inventory-collection function of IBM Director.

Creating and configuring a rack

This topic describes how to create and configure a rack in IBM Director.

You must first create a rack and then add components to the rack.

Complete the following steps to create a rack and add components to the rack:

1. In the Topology view, click **File** → **New Rack**.
2. In the Add Rack window, type a name and description for the rack. Select the type of rack from the list.
3. Click **OK**. The new rack is displayed in the right pane.
4. To add a component to the rack, in the left pane, expand the **Floor** tree.

5. From the **Floor** tree, drag a managed system or device onto a rack that is displayed in the right pane. If the inventory-collection function of IBM Director does not recognize the managed system or device, a message is displayed, asking whether you want to associate the managed system or device with a predefined component. Click **OK**.
   a. In the Associate window, expand the applicable tree and click the predefined component type that most closely resembles the target managed system or device in size.
   b. Click **OK**. The properties of the component that was associated with that managed system or device are displayed in the lower-right subpane.
   c. From the left pane, drag the managed system or device onto a rack.

   The managed system or device is displayed in the right pane as a component of the rack.
6. **Optional:** In the Components view, expand the applicable category of components.
7. Drag the predefined component onto a rack in the right pane. The component is displayed in the rack.

Adding components to an existing rack

This topic describes how to add components to an existing rack in IBM Director.

Complete the following steps to add components to an existing rack:

1. In the left pane of the Topology view of the Rack Manager window, expand the **Floor** tree.
2. Drag a managed system or device onto a rack. If the inventory-collection function of IBM Director does not recognize the managed system or device, a message is displayed, asking whether you want to associate the managed system or device with a predefined component. Click **OK**.
   a. In the Associate window, expand the applicable tree and click the predefined component type that most closely resembles the managed system or device in size.
   b. Click **OK**. The properties of the component that was associated with that managed system or device are displayed in the lower-right pane.
   c. From the left pane, drag the managed system or device onto a rack. The managed system or device is displayed in the right pane as a component in the rack.
3. **Optional:** In the left pane, select the **Components** view from the list.
4. Expand the applicable category of components.
5. Drag the predefined component onto a rack in the right pane. The component is displayed in the rack.

Removing a rack component

This topic describes how to remove a rack component in IBM Director.

To remove a rack component, in the right pane of the Topology view, right-click the rack component that you want to delete, and then click **Delete**.

This action deletes the managed system or device from the rack and displays the managed system or device in the left pane in the **Floor** tree.

Managing remote systems

You can use the Remote Control task to manage a remote Level-2 system by displaying the screen image of the Level-2 managed system on a management console.

Note:

- You can use Remote Control on Level-2 managed systems running Windows only. You cannot use Remote Control on SNMP devices.
- By default, Remote Control uses TCP. If you disable TCP-session support on a managed system, Remote Control uses UDP.

Changing remote-control states

This topic describes how to change a remote-control state in IBM Director.

To change the remote-control state, in the Remote Control window, click **Session**; then, click the state to which you want to change. The state is displayed at the top of the Remote Control window.

Changing the refresh rate

This topic describes how to change the rate at which the screen image refreshes in the Active and Monitor remote-control states in IBM Director.

You can change the rate at which the screen image refreshes in the Active and Monitor remote-control states. The following refresh rates are available:

Fastest	Screen refresh with no delay
Fast	Screen refresh every 2 seconds
Medium	Screen refresh every 10 seconds
Slow	Screen refresh every 30 seconds

To change the refresh rate, in the Remote Control window, click **Session** → **Refresh rate**; then, click the refresh rate that you want.

Playing a recorded remote-control session

This topic describes how to play a recorded remote-control session in IBM Director.

To play a recorded remote-control session, double-click the recorded remote-control session that was saved in the IBM Director Console Task pane under the Remote Control task. The remote-control session player opens. Use the controls at the bottom of the window to play, stop, and pause.

Recording a remote-control session

This topic describes how to record a remote-control session as a file and replay it later on IBM Director Console.

You can record a remote-control session as a file and replay it later on IBM Director Console. Complete the following steps to record a remote-control session:

1. In the Remote Control window, click **File** → **Start Session Logging**.
2. In the Save Session As window, type a name for the session log file. Click **OK**. Recording begins immediately.

3. When you want to stop recording, click **File → Stop Session Logging**. The session log file is saved in the IBM Director Console Task pane under the Remote Control task.

Restricting remote-control usage

This topic describes how to restrict remote-control usage in IBM Director.

You can restrict remote-control usage by using either of two methods:
- Remote-access authorization
- User administration

Remote-access authorization

This topic describes remote-access authorization in IBM Director.

Using this method, the user of the remote system can accept or reject a remote-control session when another user attempts to start the Remote Control task. If the user does not respond to the request within 15 seconds, the attempt is rejected. You can configure this option during installation of IBM Director Agent by enabling the **Require user authorization for screen access** option in the Network Driver Configuration window. This setting must be enabled on each managed system for which you want to require local authorization. See the *IBM Director Installation and Configuration Guide* for more information.

Sending key combinations

This topic describes how to send key combinations in IBM Director.

When you are using the Remote Control task, nearly all key combinations are automatically passed through to the remote managed system. However, operating-system requirements restrict the use of certain key combinations, for example, Ctrl+Alt+Del. The following key combinations cannot be used during a remote-control session because they interfere with the operating system on which the management console is running:
- Alt+Esc
- Alt+Tab
- Ctrl+Esc
- Ctrl+Alt+Del

However, in the Remote Control window, you can click **Keystrokes** and then click the applicable option to enter those key combinations for the remote managed system.

Additionally, you can set Accessibility Preferences within Console Preferences. You can set a preference for special keystrokes to be processed locally or remotely. This enables you to navigate through the Remote Control window using the keyboard.

Starting a remote-control session

This topic describes how to start a remote-control session in IBM Director.

To start a remote-control session, in the IBM Director Console Tasks pane, drag the **Remote Control** task onto a managed system.

You can select to start a remote-control session in either the Active or Monitor state.

Complete these steps to set the startup mode of a remote-control session:

1. In IBM Director Console, click **Options** → **Server Preferences**.
2. In the Server Preferences window, click the **Remote Control** tab.
3. In the **Default session state when connecting to an agent** field, select either **Active** or **Monitor**.
4. Click **OK**.

The user of the managed system can regain control at any time by pressing Alt+T on the managed system.

Transferring the clipboard

This topic describes how to copy from the management console to the managed system in IBM Director.

Note: Clipboard transfer is also supported from the managed system to the management console.

Complete the following steps to copy from the management console to the managed system:

1. From the management console desktop, select and copy the text.

   Note: This function supports text only.
2. In the Remote Control task window, click **Edit** → **Transfer Clipboard** to transfer the contents of the management console clipboard to the managed system clipboard.
3. Using the Remote Control task, open a text file and click **Edit** → **Paste** in the application window.

Establishing a remote session

This topic describes how to use the Remote Session task in IBM Director.

You can use the Remote Session task to run command-line programs on a remote managed system just as you would use the Remote Control task. Remote Session creates less network traffic and uses fewer system resources than the Remote Control task and therefore is useful in low-bandwidth situations.

Note: You can have multiple remote sessions active at the same time, but you can have only one remote session through a management server to a single managed system.

To start the Remote Session task, in the IBM Director Console Tasks pane, drag the **Remote Session** task onto a managed system. A window similar to a command-prompt window opens.

When you are targeting a managed system that is running UNIX or Linux, Remote Session uses the SSH protocol. If the SSH server on the managed system does not respond, the Remote Session task attempts to use the Telnet protocol to connect to the managed system.

Note: (Managed systems running i5/OS only) The Remote Session task uses the Telnet protocol only.

You can select text within the Remote Session window and click **Edit** → **Copy** to copy the selected text. You also can import text into a remote session by clicking **Edit** → **Paste**.

Monitoring system resources

This topic describes the Resource Monitors task in IBM Director.

You can use the Resource Monitors task to view statistics about critical system resources, such as processor, disk, and memory usage. With resource monitors, you also can set thresholds to detect potential problems with managed systems or devices. When a threshold is met or exceeded, an event is generated. You create event action plans to respond to resource-monitor events. You can apply resource monitors to individual managed systems and devices and to groups.

In IBM Director Console, under the **Resource Monitors** task, two subtasks are displayed:

All Available Recordings
View information about previously configured resource-monitor recordings.

All Available Thresholds
View information about previously configured resource-monitor thresholds.

Changing a resource-monitor threshold

You can change a resource-monitor threshold for a managed system or a group.

Complete the following steps to change a resource-monitor threshold:

1. In the IBM Director Console window, click a managed system or group in the **Managed Objects** pane.
2. Click **Tasks** → **Resource Monitors**
3. In the **Available Entries** pane, click the attribute for which the threshold was set. The attribute and corresponding statistics are displayed in the Selected Resources pane. A check mark on the entry indicates that a threshold has been applied.
4. Right-click on the attribute.
5. Click **Individual Threshold** or **Group Threshold**, depending on the type of threshold applied.
6. In the System Threshold window, make the appropriate changes.
7. Click **OK** to save your data and exit the window.

Exporting resource-monitor statistics

This topic describes how to export a record of resource-monitor statistics to a file in IBM Director.

You can export a resource-monitor record to a file in text, comma separated values (CSV), HTML, or XML format for the purpose of archiving statistics.

Complete the following steps to export a resource-monitor record:

1. In the IBM Director Console window, click **Tasks** → **Resource Monitors**.
2. In the Resource Monitors window, click **View** → **All Available Recordings**.
3. In the All Available Recordings window, right-click the record that you want to export.

4. Click **Export**.
5. In the Export window, select the drive and directory, type a name for the file, select the file type, and click **OK**.

   Note: You can save the file only to a file system that is accessible from the IBM Director Server. If the IBM Director Console is located on a separate system, you cannot save the file on the system running the IBM Director Console.

Exporting threshold tasks

This topic describes how to export threshold tasks in IBM Director.

You can export a threshold task for use on another management console. Complete the following steps to export a threshold task:

1. In the IBM Director Console Tasks pane, expand the **Resource Monitors** task.
2. Double-click the **All Available Thresholds** icon.
3. In the All Available Thresholds window, tight-click the threshold that you want to export to a task, and click **Export to Property File**.
4. In the "Export threshold to property file" window, type a file name in the **File Name** field, specifying .thrshplan for the file extension.
5. Click **OK**.

Importing threshold tasks

This topic describes how to import threshold tasks in IBM Director.

You can import a threshold task from another management console. Complete the following steps to import a threshold task:

1. In the IBM Director Console Tasks pane, right-click the Resource Monitors task and click **Import Plan from File**.
2. In the Import Threshold Plan from File window, either type a file name in the **File Name** field or navigate to the file and click the file name.
3. Click **OK**.

Monitoring the same resource on multiple groups or managed systems

This topic describes how to monitor the same resource on multiple groups or managed systems in IBM Director.

You can apply a threshold task, which is a resource-monitor threshold that you have already created, to individual managed systems or groups to monitor the same resource for a set of conditions on multiple groups or managed systems. Create a threshold task by taking a resource monitor that is configured already and exporting it to a task.

Complete the following steps to create a threshold task:

1. Create an individual or group threshold.
2. In the IBM Director Console Tasks pane, expand the **Resource Monitors** task.
3. Double-click the **All Available Thresholds** icon.
4. In the All Available Thresholds window, right-click the threshold that you want to export to a task, and click **Export to Task**.
5. In the Export Task window, type a descriptive name for the task, and click **OK**.

The new task is displayed in IBM Director Console under the Resource Monitors task. You can drag this new task onto other managed systems or groups to set identical threshold alerts.

Recording resource-monitor statistics

You can record a resource monitor to capture statistics about a managed system.

Note: You can set and record resource-monitor statistics for only individual managed systems or devices. You cannot record resource-monitor statistics for a group.

1. In the IBM Director Console Tasks pane, drag the **Resource Monitors** task onto the managed system that has the resource that you want to record.
2. In the Available Resources pane of the Resource Monitors window, expand the tree; then, double-click the resource that you want to record to add it to the Selected Resources pane.
3. Right-click the attribute cell relating to the resource and the managed system that you want to monitor. Click **Record**.
4. In the Resource Monitor Recording window, click **File** → **New**.
5. In the New Record window, type a description and select the length of time to record the resource monitor.
6. Click **OK** to start recording.
7. Click **View** → **Refresh** to update the status of the recording.

Removing a resource-monitor threshold

You can remove a resource-monitor threshold for a managed system or a group.

Complete the following steps to remove a resource-monitor threshold:

1. In the IBM Director Console window, click a managed system or group in the **Managed Objects** pane.
2. Click **Tasks** → **Resource Monitors**
3. In the Resource Monitors window, click **View** → **View All Thresholds**.
4. In the All Available Thresholds window, right-click on the threshold that you want to remove.
5. Click **Edit** → **Delete**.

Removing a resource-monitor records

You can remove a resource-monitor record for a managed system or a group.

Complete the following steps to remove a resource-monitor record:

1. In the IBM Director Console window, click a managed system or group in the **Managed Objects** pane.
2. Click **Tasks** → **Resource Monitors**
3. In the Resource Monitors window, click **View** → **All Available Recordings**.
4. Select one or more records to delete.
5. Click **Edit** → **Delete**.
6. Click **Yes** to confirm the deletion.

Saving a resource-monitor view

You can save and reuse a view that contains a set resource-monitors assembled in the Selected Resources pane of the Resource Monitor window. This user-defined view appears as a child entry of the Resource Monitors task.

Complete the following steps to save a resource-monitor view:

1. In the IBM Director Console window, click a managed system or group in the Managed Objects pane.
2. Click **Tasks** → **Resource Monitors**
3. In the Resource Monitors window, in the **Available Resources** pane, contains a list of monitors. The types of monitoring categories that are available depends on the management agent that is installed. Expand on a root category to view the categories of monitors. If a system is either CIM-enabled or DMI-enabled, a subattribute of IBM Director Agent is presented as CIM Monitors or DMI Monitors, respectively.
4. Continue expanding categories until attributes are displayed for which a resource monitor can be started.

   Note: A category might not contain attributes for these reasons:
   - None of the systems you selected for monitoring are online.
   - None of the systems you selected for monitoring allow monitor access.
   - None of the systems you selected for monitoring provide attributes for the selected category.
   - The response from one or more of the systems contacted did not complete in the time allowed.
   - The folder was opened before all systems had responded to the opening of the folder's parent category.
5. Right-click the attribute that you want to monitor and click **Add to Selected Resource Table** to view the data gathered by the selected monitor.
6. Repeats the previous steps for other monitors that you want in the view.
7. Click **File** → **Save As**
8. In the Input window, enter the name you want to assign to the view and click **OK**.

Setting a resource-monitor threshold

This topic describes how to set a resource-monitor threshold in IBM Director.

If you set a resource-monitor threshold for an attribute on a managed system or device, an event is generated when the threshold is met or exceeded. Most resource-monitor thresholds are numeric values, although for some resource monitors you can set text-string thresholds. If you set a text-string threshold, the text string you specify is monitored, and an event is generated if the text changes.

Complete the following steps to set a resource-monitor threshold:

1. In the IBM Director Console Tasks pane, drag the **Resource Monitors** task onto the managed system, device, or group that you want to monitor.
2. In the Available Resources pane of the Resource Monitors window, expand the tree; then, double-click the resource that you want to monitor. The resource is displayed in the Selected Resources pane.
3. In the Selected Resources pane, right-click the resource attribute that you want to monitor. Then, do one of the following: If you dropped the Resource

Monitors task onto an individual managed system or device, click **Individual Threshold**, or, if you dropped the Resource Monitors task onto a group, click **Group Threshold** .

4. In the System Threshold window, type a name for the threshold and complete the applicable fields. The **Enabled to generate events** check box is selected by default, so if the threshold you set in this window is met or exceeded, an event is generated. If you want to be notified when an event is generated, you must set up an event action plan that uses a threshold event filter.

   If you select the **Generate events on value change** check box, you cannot specify a threshold value. An event is generated if the value changes for the specified attribute and the **Enabled to generate events** check box is selected.

   To monitor a text-string threshold, click **Add** in the **Threshold strings** group box. In the Add string threshold setting window, type the text that you want to monitor, and select an event type from the list; then, click **OK**. The text string and event type are displayed in the **Threshold strings** group box.
5. Click **OK**. The threshold is set immediately.

If you set an individual threshold in the Resource Monitors window, a threshold icon is displayed in the data cell of the applicable attribute in the Selected Resources pane. In IBM Director Console, an icon is displayed beside the managed system in the Group Contents pane if the threshold state changes from Normal to Met or Exceeded.

If you set a group threshold, a threshold icon is displayed beside the applicable attribute in the Selected Resources column in the Selected Resources pane. If a threshold is met or exceeded on a managed system or device in the selected group, the data cell for the managed system that meets the criteria displays an icon indicating that the threshold has been met.

The following table describes the Resource-monitor status icons.

Table 15. Resource-monitor status icons

Icon	Description
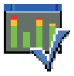	The threshold was set successfully and is in the Normal state.
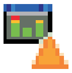	The threshold was met and has generated an event.
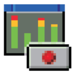	Statistics are being recorded.
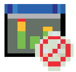	The monitor has been disabled.

Stopping the ticker-tape message display of data

This topic describes how to stop the ticker-tape message display of data in IBM Director.

To stop all resource-monitor data from being displayed in the ticker-tape area of the management console, in IBM Director Console, right-click the ticker-tape message, and then click **Remove All Monitors**.

To remove an individual monitor, select **Remove Monitor** and click the individual monitor you want to remove.

Viewing a resource-monitor threshold

You can view a resource-monitor threshold for a managed system or a group.

Complete the following steps to view a resource-monitor threshold:

1. In the IBM Director Console window, click a managed system or group in the **Managed Objects** pane.
2. Click **Tasks** → **Resource Monitors**
3. In the **Available Entries** pane, click the attribute for which the threshold was set. The attribute and corresponding statistics are displayed in the Selected Resources pane. A check mark on the entry indicates that a threshold has been applied.
4. Right-click on the attribute.
5. Click **Individual Threshold** or **Group Threshold**, depending on the type of threshold applied.

Viewing all resource-monitor thresholds

This topic describes how to view all resource-monitor thresholds in IBM Director.

To view all previously created resource-monitor thresholds, in the IBM Director Console Tasks pane, expand the **Resource Monitors** task; then, double-click the **All Available Thresholds** subtask to display all the thresholds that were created.

To view all the thresholds that are set on an individual managed system or group, drag the **All Available Thresholds** subtask onto a managed system or group to display all the thresholds that have been created for that system or group.

Viewing available resource monitors

This topic describes how to view the resource monitors that are available for a managed system, device, or group in IBM Director.

You can view the resource monitors that are available for a managed system, device, or group.

Complete the following steps to view resource monitors available for a managed system, device, or group:

1. In the IBM Director Console Tasks pane, drag the **Resource Monitors** task onto the managed system, device, or group that you want to monitor.
2. In the Available Resources pane in the Resource Monitors window, expand the tree to view which resource monitors are available.

Viewing a saved resource-monitor view

You can reuse a saved view that contains a set of resource monitors assembled in the Selected Resources pane of the Resource Monitor window. This user-defined view appears as a child entry of the Resource Monitors task.

Complete the following steps to save a resource-monitor view:

1. In the IBM Director Console window, click a managed system or group in the Managed Objects pane.
2. Click **Tasks** → **Resource Monitors**

3. Click the name of the saved view that you want to display in the Selected Resource pane.

Viewing resource-monitor statistics

You can view a resource-monitor statistics for a managed system or a group.

Complete the following steps to view resource-monitor statistics:

1. In the IBM Director Console window, click a managed system or group in the Managed Objects pane.
2. Click **Tasks** → **Resource Monitors**.
3. In the Resource Monitors window, in the **Available Resources** pane, contains a list of monitors. The types of monitoring categories that are available depends on the management agent that is installed. Expand on a root category to view the categories of monitors. If a system is either CIM-enabled or DMI-enabled, a subattribute of IBM Director Agent is presented as CIM Monitors or DMI Monitors, respectively.
4. Continue expanding categories until attributes are displayed for which a resource monitor can be started.

   Note: A category might not contain attributes for these reasons:
   - None of the systems you selected for monitoring are online.
   - None of the systems you selected for monitoring allow monitor access.
   - None of the systems you selected for monitoring provide attributes for the selected category.
   - The response from one or more of the systems contacted did not complete in the time allowed.
   - The folder was opened before all systems had responded to the opening of the folder's parent category.
5. Right-click the attribute that you want to monitor and click **Add to Selected Resource Table** to view the data gathered by the selected monitor.

Viewing a graph of a resource-monitor recording

This topic describes how to view a graph of a resource-monitor recording in IBM Director.

Complete the following steps to view a graph of a resource-monitor recording:

1. In the IBM Director Console Tasks pane, expand the **Resource Monitors** task.
2. Drag the **All Available Recordings** task onto the managed system or group for which you want to review the recordings.
3. In the All Available Recordings window, locate the recording that you want to review; then, right-click the cell and click **Graph** to display a graph of the recorded data.

Viewing resource-monitor statistics on the ticker tape

You can view the resource-monitor statistics for a managed system or group continually in IBM Director Console through the ticker-tape.

Complete the following steps to view resource-monitor statistics through the ticker tape:

1. In the IBM Director Console Tasks pane, drag the **Resource Monitors** task onto the managed system or group that has the resource monitor that you want to view using the ticker tape.
2. In the Available Resources pane of the Resource Monitors window, expand the tree and locate the resource monitor for which you want to display the data.
3. Right-click the resource monitor and click **Add to Ticker Tape on IBM Director Management Console**. The managed system name or group name and the resource-monitor data are displayed on the ticker tape.

Monitoring RAID devices

You can use the ServeRAID Manager task to monitor adapters or controllers that are installed locally or remotely on managed servers.

Starting the ServeRAID Manager task

This topic describes how to start ServeRAID Manager in IBM Director.

To start ServeRAID Manager, in the IBM Director Console Tasks pane, drag the **ServeRAID Manager** task onto a managed system that supports ServeRAID.

The left pane is the Enterprise view pane, and the right pane is the Physical and Logical devices pane. The bottom pane is the event viewer.

You can use ServeRAID Manager to view information about RAID controllers and the RAID subsystem, which includes arrays, logical drives, hot-spare drives, and physical drives.

Viewing system or device information

This topic describes how to view system or device information in IBM Director through ServeRAID Manager.

To view system or device information, expand the **Managed systems** tree in the Enterprise view pane; then, click the relevant tree object. Detailed information about the selected system or device is displayed in the right pane.

Viewing ServeRAID alerts

This topic describes how to view ServeRAID alerts in IBM Director.

You can view ServeRAID alerts in the event viewer. Three icons in the event viewer provide information about Error, Warning, and Information alerts.

Using ServeRAID Manager

Introducing the ServeRAID Manager

You might want to review the following information before using the ServeRAID Manager:

- Using the ServeRAID Manager interface
- Configuring the ServeRAID controller: the basic steps
- What's new in the ServeRAID Manager
- IBM ServeRAID publications
- Finding information on the World Wide Web

You can use the ServeRAID Manager to configure, administer, and monitor controllers that are installed locally or remotely in your IBM xSeries servers.

The first time you start the ServeRAID Manager, it will display only your local system. The local system is displayed in the Enterprise view ″tree″ and information about the system is displayed in the right pane, in the Physical and Logical device views.

Whenever you start the ServeRAID Manager **after** the first time, it will display an unknown system icon for any remote systems that you have added. When the ServeRAID Manager connects with the remote systems, it will update the status in the tree.

For additional information, see the IBM ServeRAID publications on the *IBM ServeRAID Support* CD.

More information

- Using the ServeRAID Manager in bootable-CD mode
- Working with systems in the ServeRAID Manager
- Monitoring systems over the network
- ServeRAID software features

What’s new in the ServeRAID TM Manager

This release of the ServeRAID Manager has the following new features:

- Ability to manage FlashCopy backups of logical drive data in external storage enclosures

About ServeRAID Manager

ServeRAID software features: You can use the ServeRAID software with the following controllers. Most advanced features are available only with ServeRAID controllers.

ServeRAID features	HostRAID controller	ServeRAID- 8i	ServeRAID- 7K	ServeRAID- 7t	ServeRAID- 6i/6i+	ServeRAID- 6M	ServeRAID- 5i
ROM Update Wizard	No	Yes	Yes	Yes	Yes	Yes	Yes
ServeRAID Manager	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BIOS Configuration program	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Comand-Line Tool	Yes	Yes	Yes	Yes	Yes	Yes	Yes
IPSSEND FlashCopy function (Windows XP Professional, Windows 2000, Windows Server 2003, and Windows NT only)	No	No	Yes	No	Yes	Yes	No
IPSMON (NetWare only)	No	No	Yes	No	Yes	Yes	Yes
Copy Back	No	Yes	Yes	No	Yes	Yes	No
Clustering (Windows 2000 and Windows NT only)	No	No	No	No	No	Yes	No
Failover (Windows 2000 and Windows NT only)	No	No	No	No	No	Yes (Windows 2000 only)	No
RAID level-1, RAID level-1E	Yes (1 only)	Yes	Yes	Yes (1 only)	Yes	Yes	Yes
RAID level-0, RAID level-5	Yes (0 only)	Yes	Yes	Yes	Yes	Yes	Yes
RAID level-5E	No	No	No	No	No	No	No
RAID level-5EE	No	Yes	Yes	No	Yes	Yes	No
RAID level-6	No	Yes	No	No	No	No	No
RAID level-x0	No	Yes	Yes	Yes (10 only)	Yes	Yes	Yes

ServeRAID features	Integrated RAID controller	ServeRAID-4	ServeRAID-3	ServeRAID-II
ServeRAID ROM Update Wizard	No	Yes	Yes	Yes
ServeRAID Manager	Yes	Yes	Yes	Yes
ServeRAID Mini-Configuration program	No	Yes	Yes	Yes
IPSSEND	No	Yes	Yes	Yes
IPSSEND FlashCopy function (Windows XP Professional, Windows 2000, Windows Server 2003, and Windows NT only)	No	Yes	Yes	Yes
IPSMON (NetWare only)	No	Yes	Yes	Yes
Copy Back	No	No	No	No
Clustering (Windows 2000 and Windows NT only)	No	Yes	Yes	Yes
Failover (Windows 2000 and Windows NT only)	No	Yes	Yes	Yes
RAID level-1, RAID level-1E	Yes (1 only)	Yes	Yes	Yes
RAID level-0, RAID level-5	No	Yes	Yes	Yes
RAID level-5E	No	Yes	Yes	No
RAID level-5EE	No	Yes (*except* 4H)	No	No
RAID level-6	No	No	No	No
RAID level-x0	No	Yes	No	No

IBM ServeRAID-3 hardware features: The following table lists the hardware features for the IBM ServeRAID-3 controllers:

Feature	ServeRAID-3HB	ServeRAID-3H	ServeRAID-3L
Arrays (max.)	8	8	8
Battery-backup cache	Yes	Optional	No
Cache memory	32 MB	32 MB	4 MB
Hard disk drives (max.)	45	45	15
Logical drives (max.)	8	8	8
Microprocessor	40 MHz	40 MHz	25 MHz
SCSI channels	3	3	1
SCSI transfer speed (max.)	80 MB per sec.	80 MB per sec.	80 MB per sec.
Supported RAID levels	0, 1, 5, Enhanced-1 (1E), and Enhanced-5 (5E)	0, 1, 5, Enhanced-1 (1E), and Enhanced-5 (5E)	0, 1, 5, Enhanced-1 (1E), and Enhanced-5 (5E)
System PCI data bus	64 bit	64 bit	32 bit

IBM ServeRAID-4 hardware features: The following tables list the hardware features for the IBM ServeRAID-4 controllers:

Feature	ServeRAID-4H	ServeRAID-4Mx	ServeRAID-4Lx
Arrays (max.)	8	8	8
Battery-backup cache	Yes	Yes	No
Cache memory	128 MB	64 MB	32 MB
Hard disk drives (max.)	60	30	15

Feature	ServeRAID-4H	ServeRAID-4Mx	ServeRAID-4Lx
Logical drives (max.)	8	8	8
Microprocessor	266 MHz	100 MHz	100 MHz
SCSI channels	4	2	1
SCSI transfer speed (max.)	160 MB per sec.	160 MB per sec.	160 MB per sec.
Supported RAID levels	0, 1, 5, Enhanced-1 (1E), and Enhanced-5 (5E), 00, 10, 1E0, 50	0, 1, 5, Enhanced-1 (1E), and Enhanced-5 (5E), 00, 10, 1E0, 50	0, 1, 5, Enhanced-1 (1E), and Enhanced-5 (5E), 00, 10, 1E0, 50
System PCI data bus	64 bit at 33 MHz	64 bit at 33 to 66 MHz	64 bit at 33 to 66 MHz

Feature	ServeRAID-4M	ServeRAID-4L
Arrays (max.)	8	8
Battery-backup cache	Yes	No
Cache memory	64 MB	16 MB
Hard disk drives (max.)	30	15
Logical drives (max.)	8	8
Microprocessor	100 MHz	100 MHz
SCSI channels	2	1
SCSI transfer speed (max.)	160 MB per sec.	160 MB per sec.
Supported RAID levels	0, 1, 5, Enhanced-1 (1E), and Enhanced-5 (5E), 00, 10, 1E0, 50	0, 1, 5, Enhanced-1 (1E), and Enhanced-5 (5E), 00, 10, 1E0, 50
System PCI data bus	64 bit at 33 MHz	64 bit at 33 MHz

IBM ServeRAID-5i hardware features: The following tables list the hardware features for the IBM ServeRAID-5i controller:

Feature	ServeRAID-5i
Arrays (max.)	8
Battery-backup cache	Yes
Cache memory	128 MB
Hard disk drives (max.)	30
Logical drives (max.)	8
Microprocessor	100 MHz
SCSI channels	0
SCSI transfer speed (max.)	320 MB per sec.
Supported RAID levels	0, 1, 5, Enhanced-1 (1E), 00, 10, 1E0, 50
System PCI data bus	64 bit at 66 MHz

Using the ServeRAID Manager interface: Before you begin using the ServeRAID Manager, take some time to review the layout of its user interface.

- Menubar
- Toolbar
- Enterprise view
- Physical and Logical device views
- Event viewer
- Status bar

You can work with the ServeRAID Manager interface by doing any of the following:

- Select actions from the menubar.
- Click an item on the toolbar.
- Right-click an object in the Enterprise view, the Physical devices view, or Logical devices view.
- Drag and drop objects in the Configuration wizard.
- Double-click objects and events.

More information

- Using ServeRAID Manager Assist, hints and tips
- Finding information in the ServeRAID Manager
- Hints and tips (action)
- Searching for information in ServeRAID Manager Assist
- ServeRAID Manager accessibility features
- Setting your preferences in the ServeRAID Manager

Using ServeRAID Manager Assist, hints, and tips: Use ServeRAID Manager Assist to tour and familiarize yourself with objects in the ServeRAID Manager Enterprise view and the Physical and Logical device views. Click on an object to view information about that object.

Tip: If you right-click on an object, you can select Hints and tips to view information about the object and the valid actions for that object.

Click ⬅ (Back) and ➡ (Forward) to review pages already presented during your current ServeRAID Manager Assist session.

If the ServeRAID Manager is installed on a server that is connected to a printer and its operating system supports printing in the ServeRAID Manager, the 🖨 (Print) is displayed on the ServeRAID Manager Assist toolbar. Click this icon to print the current page.

Note:

1. The ServeRAID Manager might print some pages with overlapping text.
2. You cannot print in bootable-CD mode.

To search for terms in the ServeRAID Manager Assist pages, type one or more words in the entry field and click 🔍 (Search). Search will list any page that contains all the words you typed (it is a Boolean AND search). To search for a specific phrase, surround the phrase with double quotes. For example, ″logical-drive migration.″

Note: The search is *not* case sensitive.

To exit the ServeRAID Manager Assist window, click ✕ (Exit) and you return to the ServeRAID Manager.

More information

- Finding information in the ServeRAID Manager
- Hints and tips (action)
- Searching for information in ServeRAID Manager Assist (action)

Finding information in ServeRAID Manager:

Tip: If you want to learn more about objects in the Enterprise view or the Physical and Logical devices views, right-click

If you want to learn more about the ServeRAID Manager interface, use the tool tips that are displayed when you hover the mouse over a window element.

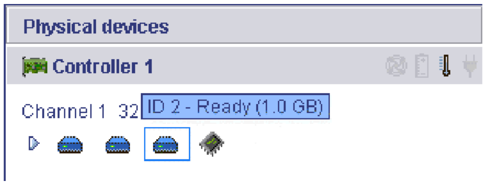

Hints and tips .

Note: If tool tips are not displayed when you hover over a window, verify that tool tips are enabled. Click **View** → **Tool tips** to enable or disable tool tips.

If you want context-sensitive information about the current window (for example, ″Creating logical drives″ in the Configuration wizard), click **Help** → **Information about this window** or the Help push-button.

The following items on the Help menu provide additional information resources:

- Search lets you search for terms in the ServeRAID Assist pages.
- ServeRAID publications lists IBM publications and where to find them.
- IBM online support lists IBM Web sites and available online support.
- What's new in the ServeRAID Manager lists the new features delivered in this version of the ServeRAID Manager program.
- About ServeRAID Manager reports the ServeRAID Manager version number, copyright, and legal information.

More information

- Using ServeRAID Assist, hints, and tips
- Hints and tips (action)
- Searching for information in ServeRAID Assist (action)

Using hints and tips: Use this action to view information about objects in the Enterprise view or the Physical and Logical device views.

1. In the Enterprise view, Physical devices view, or Logical devices view, click an object.
2. Right-click **Hints and tips**. The ″ServeRAID Manager Assist″ window opens.

To exit the ServeRAID Manager Assist window, click ✖ to return to the ServeRAID Manager.

More information

- Using ServeRAID Manager Assist, hints, and tips

Searching for information in ServeRAID Manager Assist: Use this action to search the ServeRAID Manager Assist pages.

1. Click **Help** → **Search**. The Search window opens.
2. Type one or more words in the entry field of the topic that you want to find. Search looks for any page that contains all the words you typed (that is, it is a Boolean AND search). To search for a specific phrase, surround the phrase with double quotes. For example, "logical-drive migration."

   Note: The search is *not* case sensitive.
3. Click **OK**. The Search window opens listing all the pages that contain all the words you typed.
4. Click on a topic to view its information.
5. Click ⬅ (Back) on the toolbar to view the Search listing again.

ServeRAID Manager accessibility features: Successful access to information and use of information technology by people who have disabilities is known as "accessibility."

The ServeRAID Manager provides keyboard accessibility for its supported operating systems.

On Microsoft Windows, the ServeRAID Manager works with the types of assistive technologies used by people who have disabilities and also supports high-contrast screen colors.

If you require accessibility, use the ServeRAID Manager on a Microsoft Windows operating system. If you need to work with systems that are installed with other operating systems, connect to these systems through the network using a Microsoft Windows system.

More information

- Accessing the ServeRAID Manager using the keyboard
- Accessing the ServeRAID Manager using screen-reading software
- Adding remote systems

Setting your preferences in the ServeRAID Manager: You can change some settings in the ServeRAID Manager to match your preferences.

The ServeRAID Manager uses an alarm to announce warning and error events. You can turn the repeating alarm on and off, change the amount of time between each repeating alarm, and change the duration of the alarm.

You can choose to hide the toolbar, the status bar, and the tool tips. To do so, click **View** in the menubar and click the item you want to turn off.

Also, you can refresh the ServeRAID Manager, including the Enterprise view. Click **View** → **Refresh**.

Additionally, you can adjust the size of the ServeRAID Manager viewing areas to make it easier to see information in which you are interested. Just drag the horizontal and vertical bars that separate the ServeRAID Manager panes.

More information

- Using the ServeRAID Manager interface
- Changing the alarm settings (action)

Using ServeRAID in a cluster environment:

Note: This action is not supported when using the following:

1. Integrated RAID controller
2. ServeRAID-5i controller
3. ServeRAID-6i/6i+ controller
4. ServeRAID-7k controller
5. ServeRAID-7t controller
6. ServeRAID-8i controller
7. HostRAID controller

A cluster is a group of independent computer systems that work together as a single logical system. A client interacts with a cluster as though it is a single server. In other words, a cluster is a group of computers linked together in such a way that they share and manage a set of resources that can support a number of users at the same time.

The ServeRAID high-availability clustering solution is based on a two-server cluster, where both servers can access the same storage devices, but only one server at a time controls the storage devices shared by both servers. If one of the two running servers fails, the remaining server automatically assumes control of the shared resources managed by the failed server while still controlling its own resources at the same time. The failed server can then be repaired offline without the loss of time or work efficiency because the shared data and applications earlier managed by the failed server are still online to the clients.

When the failed server is operational again, it can be placed back into the cluster; the resources are reallocated between the two servers and the cluster resumes normal operation.

You will need clustering hardware and software to configure a cluster. You can install IBM ServeRAID Cluster Solution by using one of the following software applications:

- Microsoft Cluster Server, which is part of Windows NT Server Enterprise Edition.
- Microsoft Cluster Server using Windows 2000, which is part of Windows 2000 Advanced Server Edition.

Important: Be sure to review *Installing the IBM ServeRAID Cluster Solution*.

Using Active PCI features:

Note: This action is not supported when using the following:

- Integrated RAID controller
- ServeRAID-5i controller
- ServeRAID-6i/6i+ controller
- ServeRAID-6M controller under Windows operating systems

Some IBM servers support Active PCI (also called hot-plug PCI) features. You can use these features to install or remove PCI controllers without turning off the server. The following table summarizes which operating systems support these features.

Feature	Windows NT 4.0	Windows 2000	Windows Server 2003	NetWare 4.x	NetWare 5.x
Hot add	Yes	Yes	Yes	No	Yes
Hot remove	No	Yes	Yes	Yes	Yes
Hot replace	Yes	No	No	No	No

You can use the hot-add feature to add a controller to a running server, thus expanding its capacity.

Note: If you are not using the hot-add feature, restore the controller to the factory-default settings before configuring arrays and logical drives.

Use the hot-remove feature to remove a controller from a running server. If a controller fails, use the hot-replace feature to replace a controller with an identical controller.

Attention: Do *not* attempt a hot-replace operation on Windows 2000, Windows Server 2003, or NetWare by hot-removing a failed controller and then hot-adding a new controller. Loss of data can occur. If a controller fails on these operating systems, you *must* shut down the server to replace the controller.

More information

- Using Windows NT 4.0 with Active PCI features
- Rebuilding a defunct drive
- Recovering from defunct drives
- Rebuilding a hot-swap drive
- Replacing a controller (action)
- Failing from the active to the passive controller (action)

Using Windows NT 4.0 with Active PCI features: To use Active PCI with Windows NT 4.0 and a ServeRAID-4 controller, you must install the following software components in this order:

1. DMI Service provider. A free version is included on the *IBM ServeRAID Support* CD in the following directory:

   e:\WINNT\DMISP\setup.exe

   where *e* is the CD-ROM drive.

   Note: This version is sufficient for most users' needs, but a Y2K compliant version is available at the following Web site:

   http://www.enablers.com
2. IBM Hot Plug for Windows NT 4.0 Package, version 4.2 or later. This package is available from the IBM Support Web site.

   Note: Be sure to read the instructions and restrictions for this software program.
3. ServeRAID Active PCI DMI component. This is component is installed automatically when you install the ServeRAID Manager program.

To perform a hot-replace operation, start the IBM ServeRAID Hot Replace Wizard. You can start this program from within either the IBM Hot Plug for Windows NT 4.0 program or the ServeRAID Manager program. You can use the ServeRAID Manager program to start the wizard either on the server with the failed controller, or across the network from a remote installation of the ServeRAID Manager.

Note: It is useful to start the IBM ServeRAID Hot Replace Wizard from a remote installation when the server with the failed controller does not have a monitor.

More information

- Using Active PCI features
- Replacing a ServeRAID controller (action)

Understanding unattended mode: Unattended mode is a setting that alters how the ServeRAID Manager startup code (that is, BIOS) handles failures during a system startup. Examples of possible failures during the startup are the following:

- The BIOS cannot communicate with the controller.
- A physical drive state has changed.
- A logical drive state has changed.

You can set unattended mode: either Disabled or Enabled. The default is Disabled.

When unattended mode is disabled and the BIOS detects a failure, the system remains at the recovery option screen and waits for the user to respond.

When unattended mode is enabled and the BIOS detects a failure, the system waits for 30 seconds for the user to respond to the recovery options screen. If no user responds, the BIOS automatically selects the first recovery option in the list and continues the system startup process. This mode is useful for remote systems where a user is not present for system startups.

Consider carefully whether to enable unattended mode. Depending on how you use the system, unattended mode might be undesirable.

For example, consider a system that includes an enclosure of multiple physical drives. If the enclosure is turned off when you start the system, the BIOS reports a drive failure. If unattended mode is enabled, the BIOS accepts a recovery option that changes the configuration. When the system restarts with the enclosure turned on, the configuration is no longer correct and the system no longer works as originally configured.

You can change unattended mode with the ServeRAID Manager in bootable-CD mode only.

Note: If you are configuring your system for clustering, you must enable unattended mode. If you use the ServeRAID Manager to enable your system for clustering, the ServeRAID Manager automatically enables unattended mode.

More information

- Enabling and disabling unattended mode (action)
- Configuring controllers for clustering (action)
- Using ServeRAID Manager in a cluster environment

Using the ServeRAID Manager interface

Before you begin using the ServeRAID Manager, take some time to review the layout of its user interface.

- Menubar
- Toolbar
- Enterprise view
- Physical and Logical device views
- Event viewer
- Status bar

You can work with the ServeRAID Manager interface by doing any of the following:

- Select actions from the menubar.
- Click an item on the toolbar.
- Right-click an object in the Enterprise view, the Physical devices view, or Logical devices view.
- Drag and drop objects in the Configuration wizard.
- Double-click objects and events.

More information

- Using ServeRAID Manager Assist, hints and tips
- Finding information in the ServeRAID Manager
- Hints and tips (action)
- Searching for information in ServeRAID Manager Assist
- ServeRAID Manager accessibility features
- Setting your preferences in the ServeRAID Manager

Finding information on the World Wide Web

IBM home page http://www.ibm.com/pc/

IBM support page http://www.ibm.com/pc/support/

Complete the following steps to access ServeRAID and Clustering support:

1. From the **Browse the support site** list box, select **Servers**. The IBM xSeries, Netfinity, and PC Server Support page opens.
2. From the **Family** list box, select **ServeRAID** or **Clustering**. The ServeRAID or Clustering Support page opens.
3. Click one of the following options that appear in the left margin:

   Downloadable files
   Download the latest versions of Clustering software, the ServeRAID Manager program, BIOS and firmware code, device-driver updates, and other important information.

   Hints and tips
   Obtain useful information about IBM Clustering and ServeRAID products, as well as tips for troubleshooting potential problems.

   Online publications
   Download the installation and user's guides, references, white papers, and other IBM publications.

Note: All IBM ServeRAID publications are available on the *IBM ServeRAID Support* CD.

4. From the **. . . by Category** list box, select **RAID**; a list of downloadable files appears below the list box.

IBM ServerProven compatibility page http://www.ibm.com/PC/us/compat/
Detailed information about server compatibility issues.

IBM ServeRAID publications

The following publications are available in Portable Data Format (PDF) on the *IBM ServeRAID Support* CD in the /BOOKS directory:

- *IBM ServeRAID User's Reference* (SRAID.PDF)
- *IBM ServeRAID Installation Guide (series 4, 5, 6, and 7 controllers)* (INSTALL.PDF)
- *IBM ServeRAID Device Driver Installation Instructions* (DEVDRV.PDF)

Note: Use Adobe Acrobat Reader to view these publications. The *IBM ServeRAID Support* CD contains the Acrobat Readers for Microsoft Windows, IBM OS/2, and Linux in the /BOOKS/READERS directory.

If you are installing the IBM ServeRAID Cluster Solution, you might need to refer to the *IBM Shared Disk Clustering Hardware Reference.* This publication provides general information about planning and configuring a shared-disk cluster using IBM server products. It contains illustrations, descriptions, and parts listings for various high-availability, shared-disk cluster examples.

Note: You can obtain this publication from the IBM Support Web site.

In addition, the following IBM Redbooks might be of interest:

- *Implementing Netfinity Disk Subsystems: ServeRAID SCSI, Fibre Channel, and SSA*
- *Tuning Netfinity Server for Performance: Getting the most out of Windows 2000 and Windows NT 4.0*
- *Netfinity Director: Integration and Tools*
- *Netfinity Clustering Planning Guide*

You can download these publications from the IBM Web site:
www.ibm.com/redbooks

Understanding RAID technology

Redundant array of independent disk (RAID) is the technology of grouping several **physical** drives in a computer into an array that you can define as one or more logical drives. Each **logical drive** appears to the operating system as a single drive. This grouping technique greatly enhances logical-drive capacity and performance beyond the physical limitations of a single physical drive.

When you group multiple physical drives into a logical drive, the ServeRAID controller can transfer data in parallel from the multiple drives in the array. This parallel transfer yields data-transfer rates that are many times higher than with nonarrayed drives. This increased speed makes the system better able to meet the **throughput** (the amount of data processed in a given amount of time) or productivity needs of the multiple-user network environment.

The ability to respond to multiple data requests provides not only an increase in throughput, but also a decrease in response time. The combination of parallel transfers and simultaneous responses to multiple requests enables disk arrays to provide a high level of performance in network environments.

More information

- Things to consider when changing the RAID level
- Understanding stripe-unit size
- RAID level-0
- RAID level-1
- RAID level-1E
- RAID level-5
- RAID level-5 Enhanced
- RAID level-5EE
- RAID level-6
- RAID level-x0
- RAID volumes

Understanding stripe-unit size

With RAID technology, data is **striped** across an array of physical drives. This data-distribution scheme complements the way the operating system requests data.

The granularity at which data is stored on one drive of the array before subsequent data is stored on the next drive of the array is called the **stripe-unit size** .

You can set the stripe-unit size to 8 KB, 16 KB, 32 KB, or 64 KB. You can maximize the performance of your ServeRAID controller by setting the stripe-unit size to a value that is close to the size of the system I/O requests. For example, performance in transaction-based environments, which typically involve large blocks of data, might be optimal when the stripe-unit size is set to 32 KB or 64 KB. However, performance in file and print environments, which typically involve multiple small blocks of data, might be optimal when the stripe-unit size is set to 8 KB or 16 KB.

SR-8i SR-7t HostRAID The ServeRAID-7t, ServeRAID-8i, and HostRAID controllers do not support an 8 KB stripe-unit size. The ServeRAID-7t and ServeRAID-8i controllers support these additional stripe-unit sizes: 128 KB, 256 KB, 512 KB, and 1024 KB.

The collection of stripe units, from the first drive of the array to the last drive of the array, is called a *stripe*.

SR-SCSI HostRAID After you configure an array and store data on the logical drives, you cannot change the stripe-unit size without destroying data in the logical drives.

You can set the stripe-unit size to 8 KB, 16 KB, 32 KB, or 64 KB. The default setting is 8 KB data bytes.

- When the stripe-unit size is 8 KB or 16 KB, the maximum number of physical drives in an array is 16.
- If you have a ServeRAID-3H or ServeRAID-3HB controller using ServeRAID firmware (version 4.0, or later) and the stripe-unit size is 32 KB or 64 KB, the

maximum number of physical drives in an array is 16. Otherwise, when the stripe-unit size is 32 KB or 64 KB, the maximum number of physical drives in an array is 8.

- If you have a ServeRAID-4 controller and the stripe-unit size is set to 32 KB or 64 KB, the maximum number of physical drives in an array is 16.

More information

- Changing the stripe-unit size (action)
- Application environment groups reference

Understanding RAID level-0

RAID level-0 stripes the data across all the drives in the array. This offers substantial speed enhancement, but provides no data redundancy. RAID level-0 provides the largest storage capacity of the RAID levels that are offered, because no room is taken for redundant data or data-parity storage.

RAID level-0 requires a minimum of one drive and, depending upon the level of firmware and the stripe-unit size, supports a maximum of 8 or 16 drives.

SR-8i SR-7t HostRAID For ServeRAID-8i, ServeRAID-7t, and HostRAID controllers, RAID level-0 requires a minimum of two drives.

The following illustration shows an example of a RAID level-0 logical drive.

RAID level-0 example

Start with two physical drives.

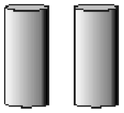

Create an array using the two physical drives.

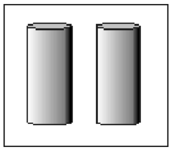

Then create a logical drive within that array.

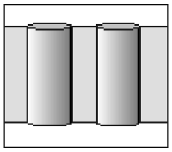

The data is striped across the drives, creating blocks. Notice that the data is striped across all the drives in the array, but no redundant data is stored.

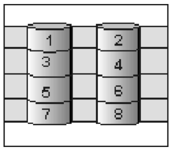

A physical drive failure within the array results in loss of data in the logical drive assigned RAID level-0, but only in that logical drive. If you have logical drives assigned RAID level-1, 1E, 5, or 5E in the same array, they will not lose data.

Note: If you have an array that contains only one physical drive, you can assign only RAID level-0 to the logical drive in that array.

When you replace a failed drive, the controller can rebuild all the RAID level-1, 1E, 5, 5E, and 5EE logical drives automatically onto the replacement physical drive. However, any data stored in a failed RAID level-0 logical drive is lost.

Although the risk of data loss is present, you might want to assign RAID level-0 to one of the logical drives to take advantage of the speed this RAID level offers. You can use this logical drive to store data that you back up each day and can re-create easily. You also might want to use a RAID level-0 logical drive when you require maximum capacity.

Advantages and disadvantages

RAID level-0 offers the following advantages and disadvantages.

Advantages	Disadvantages
• Substantial speed enhancement • Maximum utilization of physical drive storage capacity, because no room is taken for redundant data or data-parity storage	No data redundancy, resulting in data loss in the event that a physical drive fails

More information

- Understanding RAID technology
- Understanding stripe-unit size
- Selecting a RAID level
- RAID level-1
- RAID level-1 Enhanced
- RAID level-5
- RAID level-5 Enhanced
- RAID level-5EE
- RAID level-6
- RAID level-x0
- RAID volumes
- Software and hardware support of RAID levels

Understanding RAID level-1

RAID level-1 uses data mirroring. Two physical drives are combined into an array, and data is striped across the array. The first half of a stripe is the original data; the second half of a stripe is a **mirror** (that is, a copy) of the data, but it is written to the other drive in the RAID level-1 array.

RAID level-1 provides data redundancy and high levels of performance, but the storage capacity is diminished. Because the data is mirrored, the capacity of the logical drive when assigned RAID level-1 is 50% of the array capacity.

RAID level-1 requires two physical drives.

The following illustration shows an example of a RAID level-1 logical drive.

RAID level-1 example

Start with two physical drives.

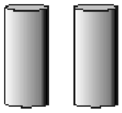

Create an array using the two physical drives.

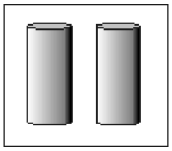

Then create a logical drive within that array.

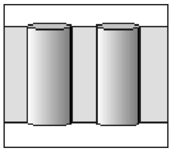

The data is striped across the drives, creating blocks. Notice that the data on the drive on the right is a copy of the data on the drive on the left.

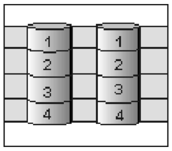

With RAID level-1, if one of the physical drives fails, the controller switches read and write requests to the remaining functional drive in the RAID level-1 array.

Advantages and disadvantages

RAID level-1 offers the following advantages and disadvantages.

Advantages	Disadvantages
• 100% data redundancy • High performance	Allows only 50% of the physical drive storage capacity to be used

More information

- Understanding RAID technology
- Understanding stripe-unit size
- Selecting a RAID level
- RAID level-0
- RAID level-1 Enhanced
- RAID level-5
- RAID level-5 Enhanced
- RAID level-5EE
- RAID level-6
- RAID level-x0
- RAID volumes
- Software and hardware support of RAID levels

Understanding RAID level-1 Enhanced

RAID level-1 Enhanced (RAID level-1E) combines mirroring and data striping. This RAID level stripes data and copies of the data across all of the drives in the array. As with the standard RAID level-1, the data is mirrored, and the capacity of the logical drive is 50% of the array capacity.

RAID level-1E has a similar profile to RAID level-1; it provides data redundancy and high levels of performance, but the storage capacity is diminished. However, RAID level-1E allows a larger number of physical drives to be used.

RAID level-1E requires a minimum of three drives and, depending upon the level of firmware and the stripe-unit size, supports a maximum of 8 or 16 drives.

The following illustration is an example of a RAID level-1E logical drive.

RAID level-1 Enhanced example

Start with three physical drives.

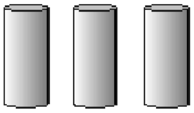

Create an array using the physical drives.

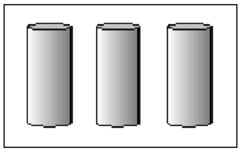

Then create a logical drive within that array.

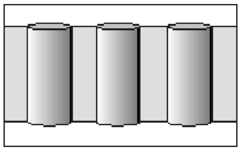

The data is striped across the drives, creating blocks. Notice that the stripe labeled * is the data stripe and the stripe labeled ** is the copy of the preceding data stripe. Also, notice that each block on the mirror stripe is shifted one drive.

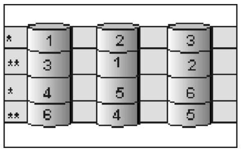

With RAID level-1E, if one of the physical drives fails, the controller switches read and write requests to the remaining functional drives in the RAID level-1E array.

Advantages and disadvantages

RAID level-1E offers the following advantages and disadvantages:

Advantages	Disadvantages
• 100% data redundancy • High performance	Allows only 50% of the physical drive storage capacity to be used

More information

- Understanding RAID technology
- Understanding stripe-unit size
- Selecting a RAID level
- RAID level-0
- RAID level-1
- RAID level-5
- RAID level-5 Enhanced
- RAID level-5EE
- RAID level-6
- RAID level-x0
- RAID volumes

Understanding RAID level-5

RAID level-5 stripes data and parity across all drives in the array.

RAID level-5 offers both data protection and increased throughput. When you assign RAID level-5 to an array, the capacity of the array is reduced by the capacity of one drive (for data-parity storage). RAID level-5 gives you higher capacity than RAID level-1, but RAID level-1 offers better performance.

RAID level-5 requires a minimum of three drives and, depending upon the level of firmware and the stripe-unit size, supports a maximum of 8 or 16 drives.

The following illustration is an example of a RAID level-5 logical drive.

RAID level-5 example

Start with four physical drives.

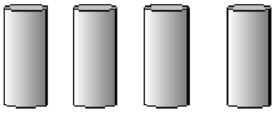

Create an array using three of the physical drives, leaving the fourth as a hot-spare drive.

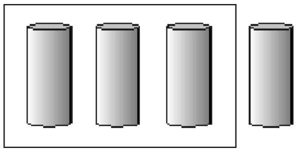

Then create a logical drive within that array.

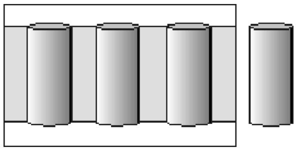

The data is striped across the drives, creating blocks.

Notice that the storage of the data parity (denoted by *) also is striped, and it shifts from drive to drive.

A parity block (*) contains a representation of the data from the other blocks in the same stripe. For example, the parity block in the first stripe contains data representation of blocks 1 and 2.

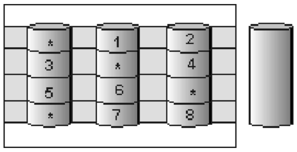

If a physical drive fails in the array, the data from the failed physical drive is reconstructed onto the hot-spare drive.

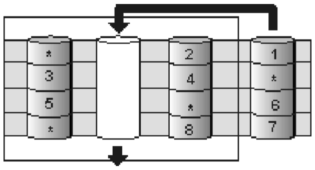

Advantages and disadvantages

RAID level-5 offers the following advantages and disadvantages.

Advantages	Disadvantages
• 100% data protection • Offers more physical drive storage capacity than RAID level-1 or level-1E	Lower performance than RAID level-1 and level-1E

More information

- Understanding RAID technology
- Understanding stripe-unit size
- Selecting a RAID level
- RAID level-0

- RAID level-1
- RAID level-1 Enhanced
- RAID level-5 Enhanced
- RAID level-5EE
- RAID level-6
- RAID level-x0
- RAID volumes

Understanding RAID level-5 Enhanced

Note: This RAID level is not available on all controllers.

RAID level-5E is the same as RAID level-5 with a built-in spare drive. Like RAID level-5, this RAID level stripes data and parity across all of the drives in the array.

RAID level-5E offers both data protection and increased throughput. When an array is assigned RAID level-5E, the capacity of the logical drive is reduced by the capacity of two physical drives in the array (one for parity and one for the spare).

Reading from and writing to four physical drives is more efficient than reading from and writing to three physical drives and an idle hot spare. Therefore, RAID level-5E provides a higher level of performance than RAID level-5.

The spare drive is actually part of the RAID level-5E array, as shown in the following example. With such a configuration, you cannot share the spare drive with other arrays. If you want a spare drive for any other array, you must have another spare drive for those arrays.

RAID level-5E requires a minimum of four drives and, depending upon the level of firmware and the stripe-unit size, supports a maximum of 8 or 16 drives. RAID level-5E is also firmware-specific.

Note: For RAID level-5E, you can have only one logical drive in an array. When using RAID level-5E, you can have a maximum of seven logical drives on the controller.

The following illustration is an example of a RAID level-5E logical drive.

RAID level-5 Enhanced example

Start with four physical drives.

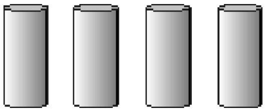

Create an array using all four physical drives.

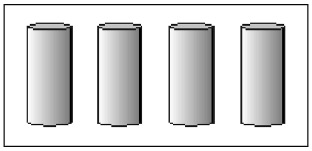

Then create a logical drive (labeled as 1) within the array. Notice that the distributed spare drive is the free space (labeled as 2) shown below the logical drive.

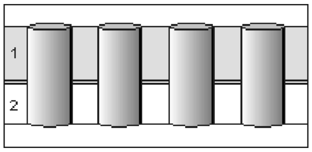

The data is striped across the drives, creating blocks in the logical drive. The storage of the data parity (denoted by *) is striped, and it shifts from drive to drive as it does in RAID level-5. Notice that the spare drive is *not* striped.

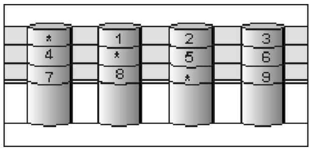

If a physical drive fails in the array, the data from the failed drive is reconstructed. The array undergoes compression, and the distributed spare drive becomes part of the array. The logical drive remains RAID level-5E.

When you replace the failed drive, the data for the logical drive decompresses and returns to the original striping scheme.

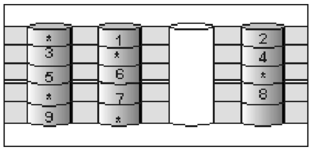

If you use a RAID level-5E logical drive in a failover or cluster configuration, the RAID level-5E logical drive will not failover while undergoing compression or decompression.

Advantages and disadvantages

RAID level-5E offers the following advantages and disadvantages.

Advantages	Disadvantages
• 100% data protection • Offers more physical drive storage capacity than RAID level-1 or level-1E • Higher performance than RAID level-5	• Lower performance than RAID level-1 and level-1E • Supports only one logical drive per array • Cannot share a hot-spare drive with other arrays • Not supported on all controllers

More information

- Understanding RAID technology
- Understanding stripe-unit size
- Selecting a RAID level
- RAID level-0
- RAID level-1
- RAID level-1 Enhanced

- RAID level-5
- RAID level-5EE
- RAID level-6
- RAID level-x0
- RAID volumes

Understanding RAID level-5EE

Note: This feature is not supported on all controllers.

RAID level-5EE is similar to RAID level-5E but with a more efficient distributed spare and faster rebuild times. Like RAID level-5E, this RAID level stripes data and parity across all of the drives in the array.

RAID level-5EE offers both data protection and increased throughput. When an array is assigned RAID level-5EE, the capacity of the logical drive is reduced by the capacity of two physical drives in the array: one for parity and one for the spare.

The spare drive is part of the RAID level-5EE array. However, unlike RAID level-5E, which uses contiguous free space for the spare, a RAID level-5EE spare is interleaved with the parity blocks, as shown in the following example. This allows data to be reconstructed more quickly if a physical drive in the array fails. With such a configuration, you cannot share the spare drive with other arrays. If you want a spare drive for any other array, you must have another spare drive for those arrays.

RAID level-5EE requires a minimum of four drives and, depending upon the level of firmware and the stripe-unit size, supports a maximum of 8 or 16 drives. RAID level-5EE is also firmware-specific.

Note: For RAID level-5EE, you can have only one logical drive in an array.

The following illustration is an example of a RAID level-5EE logical drive.

RAID level-5EE example

Start with four physical drives.

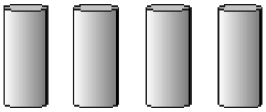

Create an array using all four physical drives.

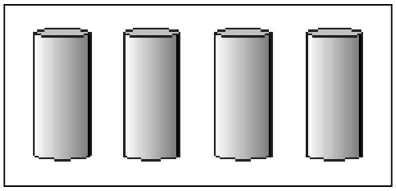

Then create a logical drive within the array.

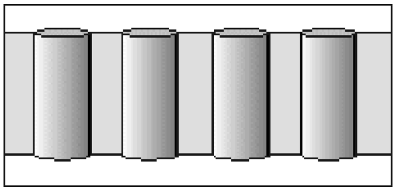

The data is striped across the drives, creating blocks in the logical drive. The storage of the data parity (denoted by *) is striped, and it shifts from drive to drive as it does in RAID level-5E. The spare drive (denoted by S) is interleaved with the parity blocks, and it also shifts from drive to drive.

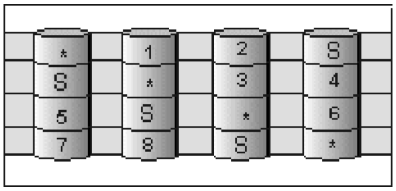

If a physical drive fails in the array, the data from the failed drive is reconstructed. The array undergoes *compaction*, and the distributed spare drive becomes part of the array. The logical drive remains RAID level-5EE.

When you replace the failed drive, the data for the logical drive undergoes *expansion* and returns to the original striping scheme.

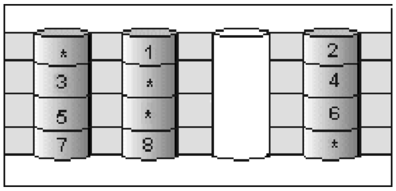

Advantages and disadvantages

RAID level-5EE offers the following advantages and disadvantages.

Advantages

- 100% data protection
- Offers more physical drive storage capacity than RAID level-1 or level-1E
- Higher performance than RAID level-5
- Faster rebuild than RAID level-5E

Disadvantages

- Lower performance than RAID level-1 and level-1E
- Supports only one logical drive per array
- Cannot share a hot-spare drive with other arrays
- Not supported on all controllers

More information

- Understanding RAID technology
- Understanding stripe-unit size
- Selecting a RAID level
- RAID level-0
- RAID level-1

- RAID level-1 Enhanced
- RAID level-5
- RAID level-6
- RAID level-x0
- RAID volumes

Understanding RAID level-6

RAID level-6 is similar to RAID level-5 but with two sets of parity information instead of one. RAID level-6 stripes blocks of data and parity across all drives in the array like RAID level-5, but adds a second set of parity information for each block of data.

When you assign RAID level-6 to an array, the capacity of the array is reduced for data-parity storage (the exact amount depends on the size of the drives in the array). The second set of parity information is added to improve fault tolerance. RAID level-6 can handle two simultaneous drive failures, where other single RAID levels can handle, at most, only one.

RAID level-6 requires a minimum of four drives and supports a maximum of 16 drives. The maximum stripe-unit size depends on the number of drives in the array.

The following illustration is an example of a RAID level-6 logical drive.

RAID level-6 example

Start with six physical drives.

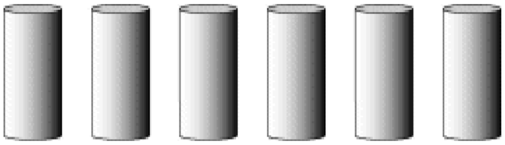

Create a logical drive using four physical drives, leaving two for hot spare drives.

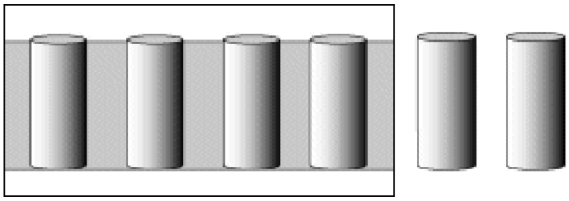

The data is striped across the drives, creating blocks in the logical drive. The storage of the data parity (denoted by * and **) is striped, and it shifts from drive to drive as it does in RAID level-5.

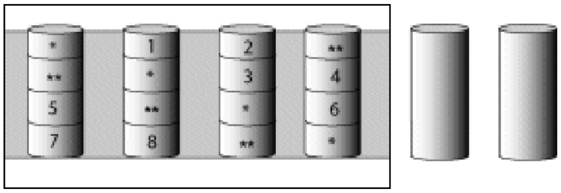

If a physical drive fails in the array, the logical drive is degraded but remains fault tolerant.

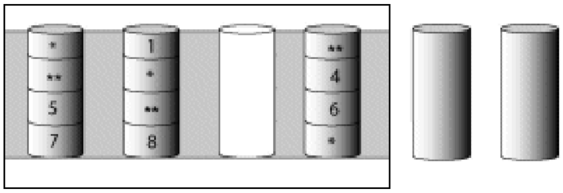

If a second physical drive fails in the array, the data from the failed drives are reconstructed onto the hot-spare drives, and the data for the logical drive return to the original striping scheme.

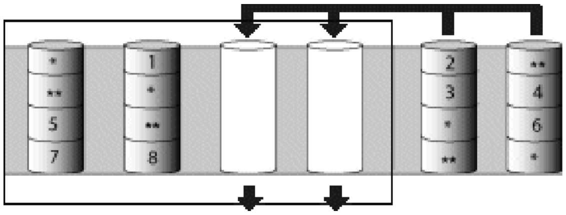

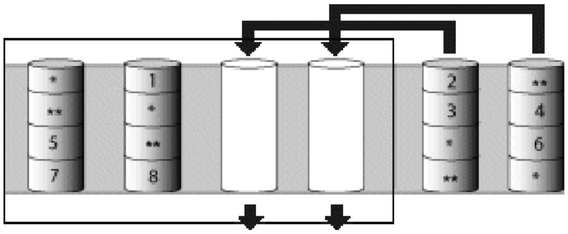

Advantages and disadvantages

RAID level-6 offers the following advantages and disadvantages.

Advantages	Disadvantages
• 100% data protection • Extremely high data fault tolerance • Can sustain two simultaneous drive failures • Good solution for mission critical applications	• Lower performance than RAID level-5 because of two parity drives • Not supported on all controllers

More information

- Understanding RAID technology
- Understanding stripe-unit size
- Selecting a RAID level
- RAID level-0
- RAID level-1
- RAID level-1 Enhanced
- RAID level-5
- RAID level-5E Enhanced
- RAID level-x0
- RAID volumes

Understanding RAID level-x0

Note: RAID level-x0 is not available on all controllers.

RAID level-x0 refers to RAID level-00, 10, 1E0, 50 and 60. RAID level-x0 uses an array of arrays, or a **spanned array**. The operating system uses the spanned array logical drive in the same way as a regular array logical drive.

RAID level-x0 allows more physical drives in an array. The benefits of doing so are larger logical drives, increased performance, and increased reliability. RAID level-0, 10, 1E, 5, 5E, and 6 cannot use more than 16 physical drives in an array; however, RAID level-1E0, 50, and 60 support 60 to 128 drives.

RAID level-x0 requires a minimum of two drives and supports a maximum of 60 to 128 drives, depending on the controller.

The following illustration is an example of a RAID level-10 logical drive.

RAID level-10 example

Start with six physical drives.

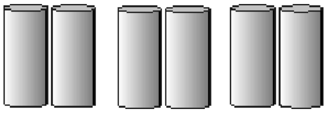

Create three arrays (labeled A, B, and C), each array using two physical drives.

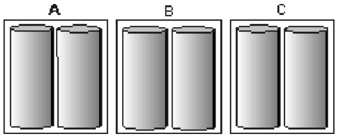

Then create a *spanned* array (labeled as *) that spans the three arrays.

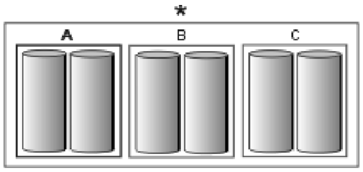

A sub-logical drive is created within *each* array (A, B, and C). Then the data is striped across the physical drives in the array, creating blocks.

Notice that, in each array, the data on the drive on the right is a copy of the data on the drive on the left. This is because the sub-logical drives (A, B, and C) are RAID level-1 in a RAID level-10 implementation (see the following table).

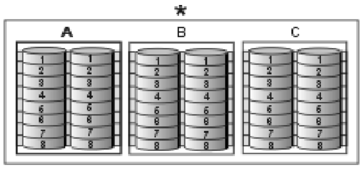

Then create a logical drive within the spanned array (*).

The data is striped across this logical drive, creating blocks (**1**- **12**). Notice that none of these blocks are redundant. This is because the logical drive is RAID level-0 in a RAID level-x0 implementation (see the following table).

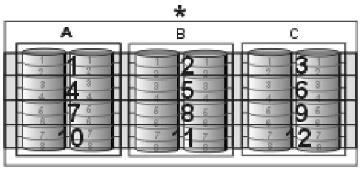

RAID level	Sub-logical drive	Spanned array logical drive
00	RAID level-0	RAID level-0
10	RAID level-1	RAID level-0
1E0	RAID level-1E	RAID level-0
50	RAID level-5	RAID level-0
60	RAID level-6	RAID level-0

With RAID level-10, 1E0, and 50, if one of the physical drives fails in a sub-logical drive, the ServeRAID controller switches read and write requests to the remaining functional drives in the sub-logical drive. With RAID level-60, if one or two of the physical drives fails in a sub-logical drive, the ServeRAID controller switches read and write requests to the remaining functional drives in the sub-logical drive. With RAID level-00, a physical drive failure within the sub-logical drive results in loss of data.

Advantages and disadvantages

RAID level-x0 offers the following advantages and disadvantages.

Advantages	Disadvantages
• Supports up to 60 physical drives on SCSI controllers • Supports up to 128 physical drives on SAS and SATA controllers • 100% data redundancy (except for RAID level-00)	• Not available on all controllers • No data redundancy for RAID level-00

More information

- Understanding RAID technology
- Understanding stripe-unit size
- Selecting a RAID level
- RAID level-0
- RAID level-1
- RAID level-1 Enhanced
- RAID level-5
- RAID level-5 Enhanced
- RAID level-5EE
- RAID level-6
- RAID volumes

Understanding RAID volumes

Note: HostRAID controllers do *not* support RAID volumes.

For the ServeRAID-8i SAS controller and ServeRAID-7t SATA controller, the ServeRAID Manager supports the following RAID volume types:

- **Simple Volume** - a single disk drive or segment; not redundant.
- **Spanned Volume** - two or more disk drives or segment with the same or different capacity, connected end-to-end. A spanned volume offers no redundancy or performance advantage over a single drive.
- **RAID Volume** - two or more logical drives with the same RAID level, connected end-to-end. The logical drives may have the same or different capacity and are *not* striped together; they may be redundant, depending on the RAID level.

  Notes:

  - You must create the individual logical drives before you create a RAID Volume. Use the Configuration wizard to create the logical drives; then, run the Configuration wizard again to create the RAID Volume.
  - When you create a RAID Volume, the ServeRAID Manager replaces the individual logical drives with a single RAID Volume icon in the Logical devices view; then, it reports the logical drives as deleted in the event viewer. The event message is normal and does not indicate a loss of data.

More information

- Understanding RAID technology
- Understanding stripe-unit size
- Selecting a RAID level
- RAID level-0
- RAID level-1
- RAID level-1 Enhanced
- RAID level-5
- RAID level-5 Enhanced
- RAID level-5EE
- RAID level-6
- RAID level-x0

Creating a simple volume: Use this action to create a simple volume. A simple volume is a single disk drive or segment used for storage in a RAID system; it is not redundant.

Complete the following steps to create a simple volume:

1. In the Physical devices view, click (online physical drive).
2. Right-click **Create simple volume**.
3. Click **Yes** to confirm.

More information

- Understanding RAID technology

Selecting a RAID level and tuning performance

Disk arrays are used to improve performance and reliability. The amount of improvement depends on the application programs that you run on the server and the RAID levels that you assign to the logical drive.

Each RAID level provides different levels of fault-tolerance (data redundancy), utilization of physical drive capacity, and read and write performance. In addition, the RAID levels differ in regard to the minimum and maximum number of physical drives that are supported.

When selecting a RAID level for your system, consider the following factors.

Note: Not all RAID levels are supported by all ServeRAID controllers.

RAID level	**Data redundancy**	**Physical drive capacity utili- zation**	**Read performance**	**Write performance**	**Built-in spare drive**	**Min. number of drives**	**Max. number of drives**
RAID level-0	No	100%	Superior	Superior	No	1	16
RAID level-1	Yes	50%	Very high	Very high	No	2	2
RAID level-1E	Yes	50%	Very high	Very high	No	3	16
RAID level-5	Yes	67% to 94%	Superior	High	No	3	16
RAID level-5E	Yes	50% to 88%	Superior	High	Yes	4	16
RAID level-5EE	Yes	50% to 88%	Superior	High	Yes	4	16
RAID level-6	Yes	50% to 88%	Very High	High	No	4	16
RAID level-00	No	100%	Superior	Superior	No	2	60
RAID level-10	Yes	50%	Very high	Very high	No	4	16
RAID level-1E0	Yes	50%	Very high	Very high	No	6	60
RAID level-50	Yes	67% to 94%	Superior	High	No	6	60 (SCSI) 128 (SAS, SATA)
RAID level-60	Yes	50% to 88%	Very High	High	No	8	128
Spanned Volume	No	100%	Superior	Superior	No	2	48
RAID Volume	No	50% to 100%	Superior	Superior	No	4	48

Physical drive utilization, read performance, and write performance depend on the number of drives in the array. Generally, the more drives in the array, the better the performance.

More information

- Understanding RAID technology
- Selecting the logical drive size
- Selecting the RAID level by array capacity
- Creating logical drives (action)
- Configuring RAID and creating arrays (action)
- Creating logical drives in the wizard
- Things to consider when changing the RAID level

Selecting the RAID level by array capacity

Note: Not all RAID levels are supported by all ServeRAID controllers.

If your array has...	**Consider this RAID level...**
One or more physical drives in an array. **Notes:** • This is the only choice if the array contains only one physical drive. • You can select RAID level-0 for any logical drive.	RAID level-0
Two physical drives. The default for two physical drives is RAID level-1.	RAID level-1
Three or more physical drives.	RAID level-1E
Three or more physical drives.	RAID level-5
Four or more physical drives.	RAID level-5E
Four or more physical drives.	RAID level-5EE
Four or more physical drives.	RAID level-6
Two or more physical drives.	RAID level-00
Four or more physical drives.	RAID level-10
Six or more physical drives.	RAID level-1E0
Six or more physical drives.	RAID level-50
Eight or more physical drives.	RAID level-60

More information

- Selecting a RAID level and performance tuning

Example: Total disk capacity: Physical drives in an array can be of different capacities (1 GB, or 2 GB, for example). The sum of the physical drives' capacity grouped in an array is the **total disk capacity**.

For example, if you group two 2 GB drives and one 3 GB drive into an array, the total disk capacity is the 7 GB physically available.

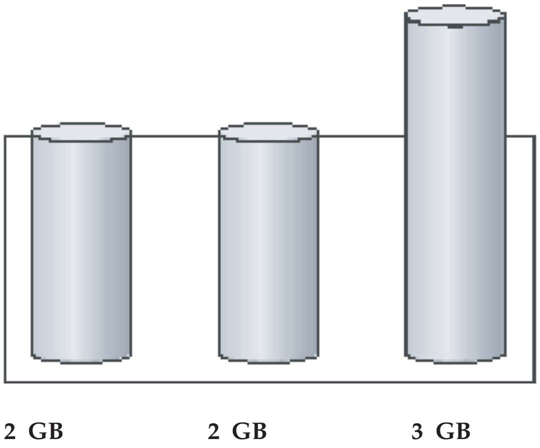

Similarly, if you group three 2 GB drives and 1 GB drive into an array, the total disk capacity is the 7 GB physically available.

More information

- Example: Usable and unusable capacity

Example: Usable and unusable capacity: Physical drive capacities influence the way you create arrays and logical drives. Drives in an array can be of different capacities (1 GB, or 2 GB, for example), but RAID controllers treat them as if they all have the capacity of the **smallest** physical drive.

For example, if you group two 2 GB drives and one 3 GB drive into an array, the usable capacity of the array is 2 GB times 3, or 6 GB, **not** the 7 GB physically available. The 7 GB is the total **disk** capacity. In the following diagram, usable capacity is labeled as 1 and unusable capacity is labeled as 2.

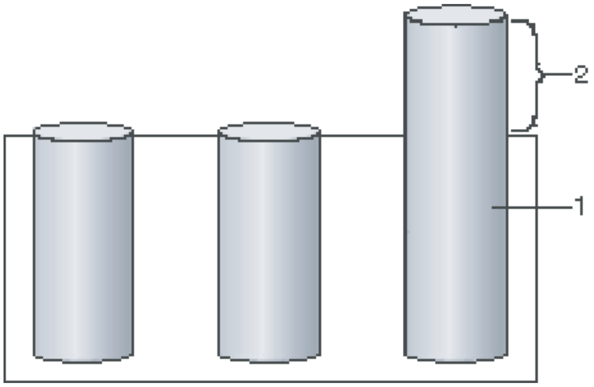

2 GB **2 GB** **3 GB**

Similarly, if you group three 2 GB drives and 1 GB drive into an array, the usable capacity of that array is 4 GB, **not** the 7 GB physically available. The 7 GB is the total **disk** capacity. The remaining capacity left on the three 2 GB drive is unusable capacity.

The optimal way to create arrays is to use physical drives that have the same capacity. Doing so avoids unusable capacity

SR-8i

SR-7t

For the ServeRAID-8i controller and ServeRAID-7t controller, usable capacity is the same as the total disk capacity. That is, you can use the remaining space to define another logical drive. The segments that make up the logical drive must be the same size on each physical drive. For example, you can group two 1 GB segments with the remaining space from the previous example to define a 3 GB logical drive. In the following diagram, the new logical drive is labeled 2; 4 GB of usable capacity remain, 2 GB on each drive, labeled 3.

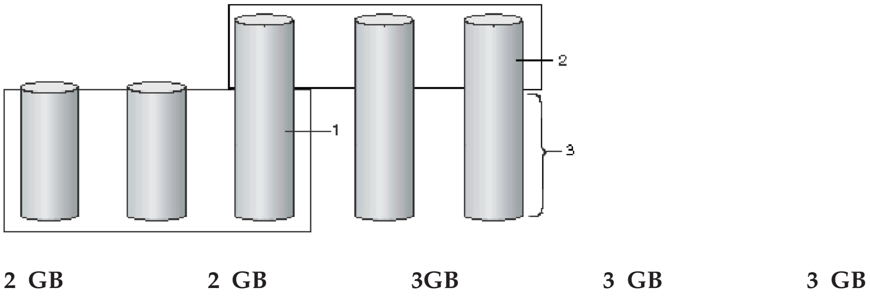

More information

- Example: Total disk capacity
- Understanding physical drive capacity and unusable capacity

Software and hardware support of RAID levels

RAID level	ServeRAID software* release that supports the RAID level	Hardware levels that support the RAID level
0	All	ServeRAID on the board, ServeRAID, ServeRAID-II, ServeRAID-3 family, ServeRAID-4 family, ServeRAID-5i, ServeRAID-6M, ServeRAID-6i/6i+, ServeRAID-7k, ServeRAID-7t, ServeRAID-8i
1	All	ServeRAID on the board, ServeRAID, ServeRAID-II, ServeRAID-3 family, ServeRAID-4 family, ServeRAID-5i, ServeRAID-6M, ServeRAID-6i/6i+, ServeRAID-7k, ServeRAID-7t, ServeRAID-8i, integrated RAID controller
1E	All	ServeRAID on the board, ServeRAID, ServeRAID-II, ServeRAID-3 family, ServeRAID-4 family, ServeRAID-5i, ServeRAID-6M, ServeRAID-6i/6i+, ServeRAID-7k, ServeRAID-8i
5	All	ServeRAID on the board, ServeRAID, ServeRAID-II, ServeRAID-3 family, ServeRAID-4 family, ServeRAID-5i, ServeRAID-6M, ServeRAID-6i/6i+, ServeRAID-7k, ServeRAID-7t, ServeRAID-8i

RAID level	ServeRAID software* release that supports the RAID level	Hardware levels that support the RAID level
5E	ServeRAID 3.50 or later	ServeRAID-3 and ServeRAID-4 families.
5EE	ServeRAID 6.00 or later	ServeRAID-4 family *except* ServeRAID-4H, ServeRAID-6M, ServeRAID-6i/6i+, ServeRAID-7k, ServeRAID-8i
6	ServeRAID 7.30 or later	ServeRAID-8i
x0	ServeRAID 4.00 or later	ServeRAID-4 family, ServeRAID-5i, ServeRAID-6M, ServeRAID-6i/6i+, ServeRAID-7k, ServeRAID-7t, ServeRAID-8i

* ServeRAID software refers to BIOS, firmware, device drivers, programs, and so forth.

RAID levels that can use a hot-spare or standby hot-spare drive

Hot-spare and standby hot-spare drives supply additional protection to a RAID configuration. When you select the RAID level for your configuration, consider the following:

RAID level	Can use a hot spare?	Can use a standby hot spare?
0	No	No
1	Yes	Yes
1E	Yes	Yes
5	Yes	Yes
5E *	Yes	Yes
5EE *	Yes	Yes
6	Yes	No
00	No	No
10	Yes	Yes
1E0	Yes	Yes
50	Yes	Yes
60	Yes	No

* RAID level-5E and RAID level-5EE integrate a distributed hot-spare drive, but also can use traditional hot-spare and standby hot-spare drives. If a physical drive fails in a RAID level-5E or level-5EE logical drive and the configuration includes a hot-spare or standby hot-spare drive, the data is rebuilt on the hot-spare or standby hot-spare drive. A RAID level-5E compression or RAID level-5EE compaction does not occur. If a second physical drive fails in the RAID level-5E or level-5EE logical drive, a data compression or compaction will take place on the distributed hot-spare drive.

More information

- Understanding RAID level-5E
- Understanding RAID level-5EE

Selecting the RAID level in the migration wizard

To select a new RAID level for the array:

1. Click the **RAID level** radio button. Only the most common choices appear in the list.

   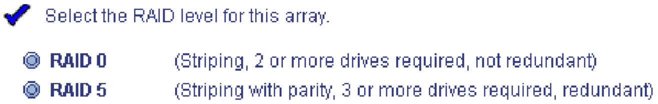

   ✓ Select the RAID level for this array.

   ○ **RAID 0** (Striping, 2 or more drives required, not redundant)
   ○ **RAID 5** (Striping with parity, 3 or more drives required, redundant)

2. To choose a different RAID level, click **Advanced settings**; then, choose a RAID level from the available options.

   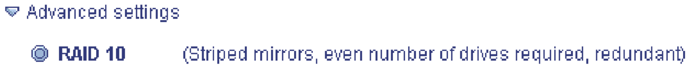

   ▽ Advanced settings

   ○ **RAID 10** (Striped mirrors, even number of drives required, redundant)

   Note: Only valid RAID level migrations appear in the list. Migration requirements for each RAID level are described here.

3. When you are ready to continue, click **Next**. The ″ Modify array″ window opens.

More information

- Things to consider when changing RAID levels
- Understanding RAID technology

Starting the ServeRAID Manager task

This topic describes how to start ServeRAID Manager in IBM Director.

To start ServeRAID Manager, in the IBM Director Console Tasks pane, drag the **ServeRAID Manager** task onto a managed system that supports ServeRAID.

The left pane is the Enterprise view pane, and the right pane is the Physical and Logical devices pane. The bottom pane is the event viewer.

You can use ServeRAID Manager to view information about RAID controllers and the RAID subsystem, which includes arrays, logical drives, hot-spare drives, and physical drives.

Logging in to the ServeRAID Manager

Use this action log into the ServeRAID Manager. You must log in when you first start the ServeRAID Manager, or if you are logged out.

You can log into the ServeRAID Manager as:

- **Administrator**: The Administrator log in allows you to view and modify the RAID configuration. You can create and delete logical drives, synchronize logical drives, perform migrations, add and remove hot-spare drives, and verify logical drives. To log in as Administrator, you must be a member of the Administrator group for your operating system. Use the Administrator user name and password or your own user name and password if you are a member of the Administrator group.

- **User**: The User log in allows you to view the RAID configuration, verify logical drives, and add and remove hot-spares. You cannot create or delete logical drives or perform migrations. To log in as User, use your normal network user name and password.

  Note: On Linux systems, the user name and password is defined in /etc/pam.d/storman.
- **Guest**: The Guest log in allows you to view the RAID configuration only. You cannot change or modify any information. To log in as Guest, simply click **Cancel** when the ServeRAID Manager ″Log in″ window opens.

Logging in at ServeRAID Manager startup

1. In the User Name field, type your user name.
2. In the Password field, type your password.

   Note: The password is case sensitive.
3. Click **Connect**. (To log in as ″Guest″, click **Cancel**.)

Logging in if you are currently logged out or logged in as Guest

1. In the Enterprise view, click 🖥 (system).
2. Click **Action** → **Log in**. The ServeRAID Manager ″Log in″ window opens.
3. In the User Name field, type your user name.
4. In the Password field, type your password.

   Note: The password is case sensitive.
5. Click **Connect**.

More information

- Logging out
- Configuring NetWare user authentication

Logging out of the ServeRAID Manager

Use this action to log out of the ServeRAID Manager.

1. In the Enterprise view, click 🖥 (system).
2. From Action menu, click **Action** → **Log out**.

   Note: After log out, you can still use the ServeRAID Manager Guest account. The Guest account allows you to view the RAID configuration but not change it.

More information

- Logging in to the ServeRAID Manager

Using the ServeRAID Manager from the command line

You can use the following command-line parameters when starting the ServeRAID Manager:

–h | –? | help
Displays basic help for the ServeRAID Manager.

–l
Starts the ServeRAID Manager in local-only mode.

The ServeRAID Manager does not obtain TCP/IP information and the remote actions are disabled. The only system appearing in the Enterprise view is the local system.

If you have remote systems defined in the Enterprise view when you run in nonlocal-only mode (that is, not using the **-l** flag), they do not appear in the Enterprise view when you run in local-only mode.

Note: Starting in local-only mode does not change any remote systems you have defined. The next time you start the ServeRAID Manager in nonlocal-only mode, your remote systems will appear in the Enterprise view again.

–r

Resets the Enterprise view. That is, it removes all remote systems from the Enterprise view tree.

Note: This command does not remove remote systems from the remote notification list.

Using the ServeRAID Manager with screen-reading software

To use the ServeRAID Manager with screen-reading software, follow these steps:</>

1. Download and install the Java 2 Platform for Microsoft Windows Operating System:
2.
   a. Go to http://java.sun.com/j2se/1.4.2/download.html and download the Java 2 Platform Standard Edition Java Runtime Environment Version 1.4.2. An executable is available for installation.
   b. Follow the instructions provided by the installation program, accepting all defaults.
3. Download and install the Java Access Bridge for Microsoft Windows Operating System:
4.
   a. Go to http://java.sun.com/developer/earlyAccess/accessbridge/ and download the compressed zip file. An older version is available at http://java.sun.com/products/accessbridge/.
   b. Extract the zip file and navigate to the Installer folder. Execute the file Install.exe.
   c. Make sure that the installation program detects the Java Virtual Machine you just installed.
   d. Follow the instructions provided by the installation program.
5. Install the screen reader software (ie, JAWS).
6. Double click the RaidManAcc.bat file to start the ServeRAID Manager program.
7.

   Note: Be sure that the java command is in your classpath. This should occur automatically when you install the Java Virtual Machine. If not, edit the batch file to explicitly point to the java.exe file found in the bin directory of the Java installation.

More information

- Accessing the ServeRAID Manager using the keyboard

Using the ServeRAID Manager in bootable-CD mode

The ServeRAID Manager program operates in two ways:

- bootable-CD mode
- As an installed software program.

When you run the ServeRAID Manager program from the *IBM ServeRAID Support* CD, you are using bootable-CD mode. The bootable-CD mode allows you to configure your ServeRAID controller *before* you install your operating system. After you have configured the ServeRAID controller and installed the operating system, you also can use bootable-CD mode to fine-tune specific ServeRAID controller settings.

To run the ServeRAID Manager program in bootable-CD mode, turn on the server; then, insert the *IBM ServeRAID Support* CD (or the CD that contains the ServeRAID Manager program that came with your server) into the CD-ROM drive.

If the ServeRAID Manager program detects unconfigured ServeRAID controllers and ready drives, the program automatically starts the Configuration wizard.

These are the actions that are available in bootable-CD mode only:

- Changing the rebuild rate
- Changing the stripe-unit size
- Changing BIOS-compatibility mapping
- Changing the SCSI-transfer speed
- Enabling and disabling read-ahead cache mode
- Enabling and disabling unattended mode
- Changing the write-cache mode on a physical drive
- Changing the write-cache mode on a logical drives
- Configuring for clustering

Accessing the ServeRAID Manager using the keyboard

Use the following key combinations to navigate in windows, tables, text areas, and the ServeRAID Manager tree.

Navigating in windows

Action	Key Combination
Navigate out forward	Tab
Navigate out backward	Shift + Tab
Activate	Enter Spacebar Alt + character accelerator key, if defined
Navigate within group	Arrow keys
Check / Uncheck	Spacebar
To beginning of list	Home
To end of list	End
Select all entries in list	Ctrl + A

Action	Key Combination
Select additional item in list	Ctrl + Spacebar
Navigate out of menu, toolbar and tool tip components	Esc
Open or post menu	Up arrow, Down arrow Enter Spacebar Alt + character accelerator key, if defined
Retract menu	Esc
Activate menu item	Enter Spacebar Alt + character accelerator key, if defined
Open submenu	Right arrow
Retract submenu	Left arrow Esc
Toggle drop-down list	Alt + Up arrow Alt + Down arrow
Move to next link in Web page	Ctrl + T
Activate link in Web page	Ctrl + Spacebar

Navigating in tables or text areas

Action	Key Combination
Navigate out forward in table or text area	Ctrl + Tab
Navigate out backward in table or text area	Ctrl + Shift + Tab
Move to next cell (wrap to next row)	Tab (Right arrow)
Move to previous cell (wrap to previous row)	Shift + Tab (Left arrow)
Move to cell above	Up
Move to cell below	Down

Accessing and navigating in the tree

Action	Key Combination
Accessing the tree pane	Ctrl + Tab
Expand entry	Right
Collapse entry	Left
Single select	Ctrl + Spacebar

Other keyboard navigation shortcuts

Action	Key Combination
Close active window	Alt + F4
Display context menu for selected object	Enter
Select controller, channel, or drive object	Spacebar
Display selected object properties	Alt + Enter
Open Help window	F1
Expand/collapse physical and logical devices	Up arrow Down arrow
Display event detail window (focus in event viewer)	Spacebar

More information

- ServeRAID Manager accessibility features
- Accessing the ServeRAID Manager using screen-reading software

Verifying that the ServeRAID Manager agent is running

Use the following procedure to verify that the ServeRAID Manager agent is running on your system.

Windows 2000, Windows XP, Windows Server 2003

1. In the lower-left corner of the desktop, click **Start** → **All Programs** → **Administrative Tools** → **Services**. The Services window opens.
2. In the Services table, verify that the ″ServeRAID Manager Agent″ is listed and started.

Windows NT

1. In the lower-left corner of the desktop, click **Start** → **Settings** → **Control Panel**. The Control Panel window opens.
2. Double-click **Services**. The Services window opens.
3. In the Services table, verify that the ″ServeRAID Manager Agent″ is listed and started.

OS/2

1. Press Ctrl + Esc. The OS/2 Window List opens.
2. Verify that the ″ServeRAID Manager Agent″ is listed.

NetWare

1. Press Alt + Esc to cycle through the NetWare consoles.
2. Verify that ″ServeRAID Manager″ is displayed on the first line of a console screen. If none of the console screens display ″ServeRAID Manager,″ the agent is not running.

Linux, OpenServer, UnixWare and Open UNIX

1. Open a shell window.
2. At the command prompt, type the following:

```
ps -ef | grep RaidAgnt.sh
```

3. If the agent is running, the process is listed as the following: sh RaidAgnt.sh.

More information

- Using the ServeRAID agent
- Starting the ServeRAID Manager agent
- Configuring the ServeRAID Manager agent

ServeRAID Manager and LED flash states

The amber and blue LEDs are controlled by the backplane (for internal physical drives) or enclosure (for external physical drives). Controllers send commands to the backplane or enclosure processor to set a logical state for the LED, and the subsystem translates that LED state into the following flash states:

Controller device state	Slot state	LED flash state
defunct	Device is faulty	On
Rebuilding	Device is rebuilding	Slow flash
Identify	Identify the device	Fast flash
Other	No error	Off

The green LED is an activity indicator for the physical drive and is controlled by the physical drive.

More information

- Identifying devices (action)
- Using Identify to work with your systems

Managing ServeRAID devices

Configuring ServeRAID controllers and enclosures in the Configuration wizard

Configuring an enclosure: the basic steps

The following steps are an overview of the enclosure configuration process:

1. In the Enterprise view, click (enclosure) that you want to configure.
2. In the Logical devices view, click ; or click on the toolbar. The Configuration wizard opens.
3. Select the physical drives that you want to include in your arrays or create as hot-spare drives.
4. Define the logical drives for your arrays.
5. Define the authentication method for each logical drive (iSCSI initiators only).
6. Define the initiators that can access the enclosure.
7. Assign logical drives to initiators.
8. Review the configuration summary.
9. Click **Apply**. The ServeRAID Manager will configure the controller and synchronize the logical drives.

More information

- Configuring RAID

Configuring ServeRAID SCSI controllers (ServeRAID series 3, 4, 5, 6, and 7K)

Choosing between Express and Custom configuration:

Note: If you intend to use your ServeRAID controllers in a Microsoft Windows failover or clustering environment, review the ServeRAID publications before configuring ServeRAID.

In the Configuration wizard, you can choose from two paths to configure your ServeRAID subsystem:

- Click **Express configuration** for a quick and easy path for you to automatically configure your ServeRAID controller. This path creates the most efficient ServeRAID configuration based on the number and capacity of the ready physical drives available in your system. If four or more ready drives of the same capacity are available, this choice also will define a hot-spare drive.

  Express configuration does the following:

  - Creates arrays by grouping together same-sized physical drives.
  - Creates one logical drive per array up to 2 terabytes (TB) maximum.
  - Assigns a RAID level based on the number of physical drives in an array:
    - An array with a single physical drive is assigned RAID level-0.
    - An array with two physical drives is assigned RAID level-1.
    - An array with three or more physical drives is assigned RAID level-5.

      Note: If the physical drive size exceeds 2TB, Express configuration removes drives until the size is less than 2 TB. It creates an additional array with the left-over drive(s).
  - Designates a hot-spare drive for the controller. If one or more arrays has four or more physical drives, the largest-sized drive from those arrays will be designated the hot-spare drive.

  Note: When you choose Express configuration, you have the opportunity to review and approve the configuration before the ServeRAID Manager applies the configuration. If you need to modify the configuration, you can do it from the configuration summary.
- Click **Custom configuration** for a path that lets you manually configure your ServeRAID controller. Choose this path to do any of the following:
  - Configure your controller manually.
  - Assign RAID level-1E, RAID level-5E, RAID level-5EE, or RAID level-x0 to a logical drive.
  - Create more than one logical drive per array.

  Using the Custom path, you can:

  - Select the physical drives that you want to include in each array and that you want to use as hot-spare drives.
  - Define the number, size, and RAID level of the logical drives for each array.
  - Review the configuration summary.

Click **Automatically initialize new logical drives**. Initializing a logical drive erases the first 1024 sectors on the drive and prevents access to any data previously stored on the drive. You can initialize the drive later if you do not choose automatic initialization.

When you are ready to continue, click **Next**.

More information

- Understanding stripe-unit size
- Initializing a logical drive
- Understanding physical drive capacity and unusable capacity

Creating arrays and hot-spare drives: Complete the following steps to create arrays:

1. Start the custom configuration wizard and complete all previous steps.
2. Click the appropriate **Array** tab in the right pane:

3. Then, from the list of ready drives, select the drives you want to add to the array:

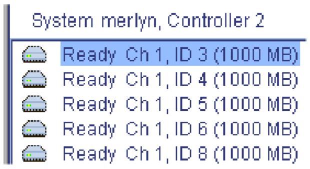

System merlyn, Controller 2
Ready Ch 1, ID 3 (1000 MB)
Ready Ch 1, ID 4 (1000 MB)
Ready Ch 1, ID 5 (1000 MB)
Ready Ch 1, ID 6 (1000 MB)
Ready Ch 1, ID 8 (1000 MB)

4. Click >> (Add selected drives) to add the drives to the array. You can click >> (Add all drives) to move **all** ready drives to an array.
5. Repeat steps 2 and 3 for each additional array or hot-spare drive that you want to configure.
6. If you do *not* want to create a spanned array, skip to step 6. Otherwise, select the **Span arrays** check box ☑ Span arrays . Then, click **Next**. If you created two arrays only, the ServeRAID Manager uses those arrays to create the spanned array; continue to step 6. Otherwise, the ″ Create spanned arrays″ window opens.

   Note: To create a spanned array, each array must have the same number of physical drives.
7. After you select the ready drives for your arrays and hot-spare drive, click **Next**. The ″ Create logical drives″ window opens.

To return to the ″ Express and Custom configuration″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Understanding creating basic arrays
- Example: Express configuration
- Creating spanned arrays
- Understanding creating spanned arrays
- Configuring the ServeRAID controller: the basic steps

Creating spanned arrays: If you want to assign RAID level-x0 to an array, you must create a spanned array.

Note: Spanned arrays are supported only by IBM ServeRAID-4 Ultra160, ServeRAID-5i Ultra320, ServeRAID-6M Ultra320, and ServeRAID-6i/6i+ Ultra320 SCSI controllers.

Complete the following steps to create one or more identical spanned arrays:

1. Start the custom configuration wizard and complete all previous steps.
2. In the list of arrays, click the arrays that you want to add to your spanned array.

   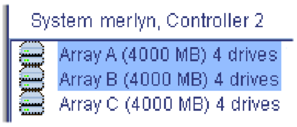
   System merlyn, Controller 2
   Array A (4000 MB) 4 drives
   Array B (4000 MB) 4 drives
   Array C (4000 MB) 4 drives

3. Click >> (Add selected arrays) to add the arrays to the spanned array. You can click >> (Add all arrays) to move **all** arrays to the spanned array.
4. To create additional spanned arrays, click the **New spanned array** tab in the right pane.

   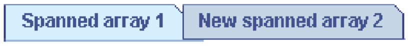

   Then, repeat steps 2 and 3.
5. Click **Next**; the ″ Create logical drives″ window opens.

To return to the ″ Create arrays″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Understanding RAID level-x0
- Understanding creating spanned arrays

Creating logical drives: Complete the following steps to create logical drives:

1. Start the custom configuration wizard and complete all previous steps.
2. Click the appropriate **Array** tab.

   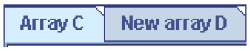

3. Select a RAID level from the drop-down list.

   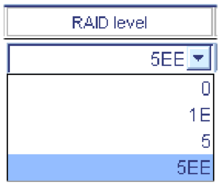

   Note:

   a. RAID level-5E and RAID level-5EE allow only one logical drive per array.
   b. If you are configuring a spanned array, you can set the RAID level only for the first logical drive you create.
   c. If you plan to use ″ Change RAID level,″ you must assign the same RAID level to all logical drives within a single array.

Attention: Before assigning a logical drive RAID level-5E or RAID level-5EE, consider the following. If a physical drive fails during a post-failover resynchronization, the logical drive will enter the blocked state. Data might be lost or damaged.

4. If you do not want to use the maximum size for the logical drive, type the size in the **Data (MB)** field.

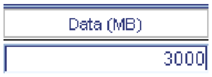

Data (MB)
3000

Note:

a. You can define up to eight logical drives per controller. There are two exceptions:
   - If an array contains a logical drive assigned RAID level-5E
   - If you want to use the logical-drive migration feature

   In these cases, one logical drive slot must be left free; therefore, you must define no more than seven logical drives.
b. Some operating systems have size limitations for logical drives. Before you save the configuration, verify that the size of the logical drive is appropriate for your operating system. For more detailed information, see your operating-system documentation.
c. A logical drive cannot exceed 2048 GB (2 terabytes); the minimum size is *n*MB, where n equals the number of drives in the array.
d. Typically, the first logical drive defined on the first ServeRAID controller found by system BIOS during startup will be your startup (boot) drive.
e. The actual logical-drive size might be slightly different from what you select. The RAID level and the number of physical drives determine the size of the logical drive. For example, an array consisting of three, 1 GB physical drives with a requested RAID level-0 logical drive of 1000 MB will actually contain only 999 MB because the data is striped across all three drives, with 333 MB on each drive.

5. If you have free space available and want to create additional logical drives, click .

6. Repeat steps 3 through 5 for each logical drive that you want to define in this array.
7. Repeat steps 2 through 6 for each additional array that you want to configure.
8. Click **Next**. The ″ Configuration summary″ window opens.

To return to the ″ Create arrays″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Selecting the RAID level by array capacity
- Selecting the array size
- Understanding physical drive capacity and unusable capacity

Confirming your express system configuration: Use the configuration summary to review all the changes that the ServeRAID Manager will apply to your configuration:

1. Review the information that is displayed in the ″Configuration summary″ window. To change the configuration, click Modify arrays or Modify logical drives .

   Notes:

   a. Additional notes and attentions regarding the configuration appear in the event viewer.
   b. Some operating systems have size limitations for logical drives. Before you save the configuration, verify that the size of the logical drive is appropriate for your operating system. For more detailed information, see your operating-system documentation.
2. Click **Apply**; then, click **Yes** when asked if you want to apply the new configuration. The configuration is saved in the ServeRAID controller and in the physical drives.

   Note: If you clicked Automatically initialize new logical drives, the ServeRAID Manager will initialize the logical drives automatically.
3. When you have completed configuring your controllers, you can change certain controller settings to fine-tune your configuration.

More information

- Configuration wizard notes and attentions
- Understanding logical-drive synchronization
- Initializing a logical drive
- Fine-tuning your system

Confirming your custom system configuration: Use the configuration summary to review all the changes that the ServeRAID Manager will apply to your configuration:

1. Review the information that is displayed in the ″Configuration summary″ window. To change the configuration, click **Back**.

   Note: Additional notes and attentions regarding the configuration appear in the event viewer.
2. Click **Apply**; then, click **Yes** when asked if you want to apply the new configuration. The configuration is saved in the ServeRAID controller and in the physical drives.

   Note: If you clicked Automatically initialize new logical drives, the ServeRAID Manager will initialize the logical drives automatically.
3. When you have completed configuring your controller, you can change certain controller settings to fine-tune your configuration.

More information

- Configuration wizard notes and attentions
- Understanding logical-drive synchronization
- Initializing a logical drive
- Fine-tuning your system

Configuring ServeRAID SAS, SATA, and HostRAID controllers (ServeRAID-7t, ServeRAID-8i, HostRAID)

Choosing between Express and Custom configuration: In the Configuration wizard, you can choose from two paths to configure your ServeRAID-8i, ServeRAID-7t, or HostRAID controller:

- Click **Express configuration** for a quick and easy path for you to automatically configure your controller. This path creates the most efficient configuration based on the number and capacity of the ready physical drives available in your system. If four or more ready drives of the same capacity are available, this choice also will define a hot-spare drive.

  Express configuration does the following:

  - Creates one logical drive, up to 2 terabytes (TB) maximum, by grouping together same-sized physical drives.

    Note: If the physical drive size exceeds 2TB, Express configuration removes drives until the size is less than 2 TB. It creates an additional logical drive with the left-over drive(s).
  - Assigns a RAID level based on the number of available physical drives:
    - A logical drive with a single physical drive is assigned RAID level-0.
    - A logical drive with two physical drives is assigned RAID level-1.
    - A logical drive with three physical drives is assigned RAID level-5.

    HostRAID For HostRAID controllers, Express configuration creates a RAID level-1 logical drive.
  - Designates a hot-spare drive for the controller. If four or more ready drives are available, the largest drive will be designated the hot-spare drive.

  Note: When you choose Express configuration, you have the opportunity to review and approve the configuration before the ServeRAID Manager applies the configuration. If you need to modify the configuration, you can do it from the configuration summary.
- Click **Custom configuration** for a path that lets you manually configure your controller. Choose this path to do any of the following:
  - Define the RAID level of the logical drives.
  - Select the physical drives that you want to include for each logical drive.
  - Define the number and size of the logical drives.
  - Fine tune your configuration for optimal performance
  - Review the configuration summary.

When you are ready to continue, click **Next**.

More information

- Understanding stripe-unit size
- Understanding logical drive synchronization
- Understanding physical drive capacity and unusable capacity

Choosing between Express and Custom configuration: In the Configuration wizard, you can choose from two paths to configure your integrated RAID controller subsystem:

- Click **Express configuration** for a quick and easy path for you to automatically configure your integrated RAID controller.

  Express configuration does the following:

  – Creates an array by grouping the first two physical drives that appear in the Physical devices view.

  – Creates a RAID level-1 logical drive.

  Note: If you click Express, you will have the opportunity to review and approve the configuration before the ServeRAID Manager applies the configuration. If you need to modify the configuration, you can do so from the configuration summary.

- Click **Custom configuration** for a path that lets you to manually configure your integrated RAID controller. Choose this path to do any of the following:

  – Configure your controller manually.

  – Select the two physical drives that you want to configure.

  – Create a hot-spare drive.

  – Preserve data on a mirror primary drive .

  Using this path, you can:

  – Select the physical drives that you want to include in the array and that you want to use as a hot-spare drive.

  – Review the configuration summary.

To initialize the new logical drives, click **Automatically initialize new logical drives**. Initializing a logical drive erases the first 1024 sectors on the drive and prevents access to any data previously stored on the drive. You can initialize the drive later if you do not choose automatic initialization.

When you are ready to continue, click **Next**.

More information

- Understanding stripe-unit size
- Initializing a logical drive

Example: Express configuration: If your server contains the following:

- One 1024 MB ready drive,
- Two 2150 MB ready drives,
- And four 4300 MB ready drives,

Express configuration will create three arrays and one hot-spare drive as follows:

Array A:
The total capacity of this array is 1024 MB (1 x 1024 MB) and it contains one, 1024 MB RAID level-0 logical drive.

Array B:
The total capacity of this array is 4300 MB (2 x 2150 MB) and it contains one, 2150 MB RAID level-1 logical drive.

Array C:
The total capacity of this array is 12900 MB (3 x 4300 MB) and it contains one, 8600 MB RAID level-5 logical drive.

Hot Spare:
Express configuration defines one of the four 4300 MB drives as a hot-spare drive.

Notes:

1. When there are four or more ready drives of the same capacity, Express configuration groups three of the drives into one array (as in Array C) and defines one of the drives as a hot spare.
2. A hot-spare drive must be of equal or greater capacity than the drive that it is intended to replace. In this configuration, the 4300 MB drive can replace any failed drives in Array B or Array C. Array A is not redundant, therefore a hot-spare drive is not used.

More information

- Choosing between Express and Custom configuration in the wizard (SCSI)
- Choosing between Express and Custom configuration in the wizard (SAS, SATA, HostRAID)
- Configuring RAID and creating arrays (action)

Choosing the RAID Level: Complete the following steps to choose the RAID level for the logical drive:

1. Click the **RAID level** radio button to choose the RAID level for the logical drive. Only the most common choices appear in the list.

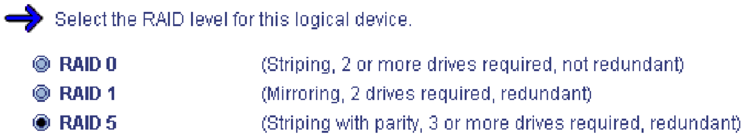

Select the RAID level for this logical device.

○ RAID 0 (Striping, 2 or more drives required, not redundant)
○ RAID 1 (Mirroring, 2 drives required, redundant)
● RAID 5 (Striping with parity, 3 or more drives required, redundant)

2. To choose a different RAID level, click **Advanced settings**; then, choose a RAID level from the available options.

Note: Not all RAID levels are available on all controllers.

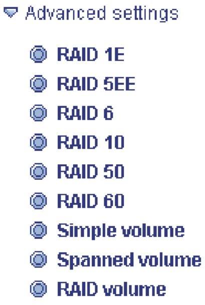

▽ Advanced settings

○ RAID 1E
○ RAID 5EE
○ RAID 6
○ RAID 10
○ RAID 50
○ RAID 60
○ Simple volume
○ Spanned volume
○ RAID volume

SR-8i SR-7t In addition to RAID Level-1E, 5EE, 6 and x0, the ServeRAID Manager supports the following RAID types for the ServeRAID-8i and ServeRAID-7t controllers:

3.

- **Simple Volume** - a single disk drive or segment; not redundant.
- **Spanned Volume** - two or more disk drives or segments with the same or different capacity, connected end-to-end. A spanned volume offers no redundancy or performance advantage over a single drive.

- **RAID Volume** - two or more logical drives with the same RAID level, connected end-to-end. The logical drives may have the same or different capacity and are **not** striped together; they may be redundant, depending on the RAID level.

  Notes:

  a. You must create the individual logical drives before you create a RAID Volume. Use the Configuration wizard to create the logical drives; then, run the Configuration wizard again to create the RAID Volume.
  b. When you create a RAID Volume, the ServeRAID Manager replaces the individual logical drives with a single RAID Volume icon in the Logical devices view; then, it reports the logical drives as deleted in the event viewer. The event message is normal and does not indicate a loss of data.

4. When you are ready to continue, click **Next**. The ″ Create logical drives″ window opens.

More information

- Understanding RAID technology
- Understanding stripe-unit size
- Understanding physical drive capacity and unusable capacity

Creating logical drives:

Note: The ServeRAID-7t and ServeRAID-8i firmware supports a maximum of 10 ″logical slices″ per physical drive.

Complete the following steps to create logical drives:

1. In the Physical devices view (on the right), click the drives you want the logical drive to use. To view available segments on each disk drive, switch to the full-size capacity view or relative-size capacity view.

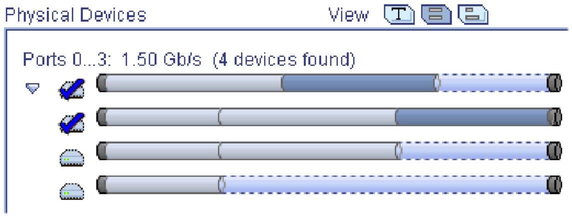

   Note: A segment shaded in light blue and outlined by a dashed line is not part of any logical drive.

2. To select a hot-spare for the logical drive, or to deselect a hot-spare, control-click the drive.
3. In the **Name** field, enter a name for the logical drive.
4. If you do **not** want to use the maximum size for the logical drive, click **Advanced settings**; then, type the size in the **Size (MB)** field:

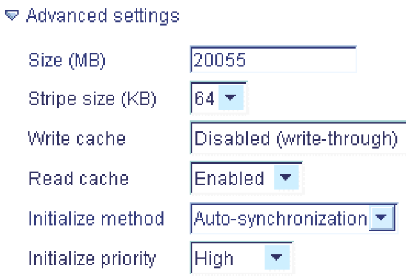

▽ Advanced settings

Size (MB) 20055

Stripe size (KB) 64

Write cache Disabled (write-through)

Read cache Enabled

Initialize method Auto-synchronization

Initialize priority High

HostRAID For HostRAID controllers, the logical drive uses all of the space on each disk drive; you cannot adjust the logical drive size.

5. To configure the logical drive for optimal performance, adjust the advanced settings as needed.

   Note: The default settings usually are adequate for most applications.
6. If you have free space available and want to create additional logical drives, click Add logical drive ; then, repeat steps 1-4.
7. When you are ready to continue, click **Next**. The ″ Configuration summary″ window opens.

More information

- Understanding physical drive capacity and unusable capacity

Creating logical drives > advanced settings:

Note: Not all options are available for all RAID levels.

Use the **Advanced settings** in the Configuration wizard to tune your logical drive configuration for optimal performance. You can set the following options:

- **Write cache** - Determines how data is stored in the controller's cache memory. You can choose from two available settings: write-through and write-back.
- **stripe-unit size** - Determines the amount of data, in kilobytes (KB), written to one segment of a logical drive before the next segment is used for subsequent data. Depending on the controller, you can choose a stripe-unit size from 16 KB to 1024 KB.

  Notes:

  1. The default stripe size usually provides the best performance for typical applications.
  2. For RAID Level-6 and Level-60 logical drives, the number of choices for stripe-unit size decreases as the number of disks in the logical drive increases.
- **Read cache** - Enables and disables read caching. When read-caching is enabled, the controller monitors read access to the logical drive. If it detects a pattern, it pre-loads the cache with the data that seems most likely to be read next. The default is Enabled.
- **Initialize method** - Determines the method used to initialize the logical drive. You can choose from three available methods:

- **Auto-synchronization** - For RAID Level-1 logical drives, copies the data from the primary drive to the mirror drive. For RAID Level-5 logical drives, computes and writes the correct parity for the entire logical drive.

  Note: Auto-synchronization can be a lengthy operation. The ServeRAID Manager performs this task in the background. However, you cannot use the logical drive until the task is complete.
- **Clear** - removes pre-existing data by overwriting every block in the logical drive. This method is faster than auto-synchronization but the logical drive is not available immediately. HostRAID For HostRAID controllers, the ServeRAID Manager supports the Clear initialization method only.
- **Quick** - makes the logical drive available immediately. It is the fastest method but should be used only for *new* physical disks.

- **Initialize priority** - Adjusts the priority of the initialization task. The default is High: initialize the logical drive as fast as possible.

When you are ready to continue, click **Next**. The ″ Configuration summary″ window opens.

More information

- Understanding stripe-unit size
- Understanding write-cache mode for logical drives
- Understanding logical-drive synchronization

Confirming your integrated RAID controller subsystem configuration: Use the configuration summary to review all the changes that the ServeRAID Manager will apply to your configuration:

1.

   Review the information that is displayed in the ″Configuration summary″ window. To change the configuration, click **Back**.
2. Click **Apply**; then, click **Yes** when asked if you want to apply the new configuration. The configuration is saved in the integrated RAID controller.

   Note: If you clicked Automatically initialize new logical drives, the ServeRAID Manager will initialize the logical drives automatically.

More information

- Configuration wizard notes and attentions
- Understanding logical-drive synchronization

Confirming your integrated RAID controller subsystem configuration: Use the configuration summary to review all the changes that the ServeRAID Manager will apply to your configuration:

1. Review the information that is displayed in the ″Configuration summary″ window. To change the configuration, click Modify arrays .

   Notes:

   a. Additional notes and attentions regarding the configuration appear in the event viewer.
   b. Some operating systems have size limitations for logical drives. Before you save the configuration, verify that the size of the logical

drive is appropriate for your operating system. For more detailed information, see your operating-system documentation.

2. Click **Apply**; then, click **Yes** when asked if you want to apply the new configuration. The configuration is saved in the integrated RAID controller.

   Note: If you clicked Automatically initialize new logical drives, the ServeRAID Manager will initialize the logical drives automatically.

More information

- Configuration wizard notes and attentions
- Understanding logical-drive synchronization

Confirming your system configuration: Use the configuration summary to review all the changes that the ServeRAID Manager will apply to your configuration:

1. Review the information that is displayed in the ″Configuration summary″ window. To change your Express configuration, click Modify logical drives ; to change your Custom configuration, click **Back**.

   Note: Additional notes and attentions regarding the configuration appear in the event viewer.
2. Click **Apply**; then, click **Yes** when asked if you want to apply the new configuration. The configuration is saved in the ServeRAID controller and in the physical drives.

   Note: The ServeRAID Manager will initialize the logical drives automatically. The initialization method depends on the RAID level and controller type.
3. After you configure your controller, you can change certain controller settings to fine-tune your configuration for optimal performance.

   Note: During Custom configuration, you can fine tune your configuration using the Advanced settings in the Configuration wizard.

Attention: After you apply the configuration, the logical drives will display as physical drives under your operating system. Before you can use these drives to store data, you must partition and format the drives using the disk management tools provided with your operating system. Each operating system provides its own disk management tools. For example, under Microsoft Windows, use the Computer Management tool to partition and format drives; then, assign each logical drive to a drive letter. For more information, see your operating system administrator’s guide.

More information

- Configuration wizard notes and attentions
- Understanding synchronizing logical drives

Selecting controllers to update

To select the ServeRAID controllers or enclosures to update:

1. In the tree, click the controller(s) or enclosure(s) you want to update. You can also choose the controllers or enclosures you want to *omit* from the update process.

For direct-attached storage devices, all controllers of the same type are selected by default. If the controller or enclosure software is already up-to-date, it is grayed out. To remove a controller or enclosure from the selection, just click it.

Note: To force a controller or enclosure to be updated, for example, if you want to down-level the software, press the Control key, then click the controller or enclosure.

2. When you are ready to continue, click **Next**; the ″ Update summary″ window opens.

Configuring ServeRAID enclosures

Choosing the configuration path: The Configuration wizard guides you through the configuration of your external storage enclosure. You can use the Configuation wizard to create arrays and logical drives for each controller in the enclosure. The Configuration wizard allows you to add drives to the array; choose the RAID level for each logical drive; choose an authentication method; and assign logical drives to users and machines (access control).

Note:

- The maximum number of arrays depends on the number of physical disks on the controller. You can configure one array per disk. *Example*: 4 physical disks=4 arrays maximum.
- You can create up to 64 logical drives per array.
- You can create up to 512 logical drives for each controller.

1. In the Enterprise view, click (enclosure).
2. Right-click **Configure storage**, or click in the Logical devices view. The Configuration wizard opens.
3. In the Select a configuration option field, choose a configuration path. Click either:
   - Create a logical drive in a new array
   - Create a logical drive in existing array
4. If you choose to create a logical drive in an existing array, select the array name from the drop-down list.
5. Click **Next**. For new arrays, the ″ Create arrays″ window opens. For existing arrays, the ″ Create logical drives″ window opens.

More information

- Configuring an enclosure: the basic steps
- Increasing logical drive capacity

Creating arrays and hot-spare drives: Complete the following steps to create arrays:

1. Click the appropriate **Array** tab in the right pane:

2. Then, from the list of ready drives, select the drives you want to add to the array:

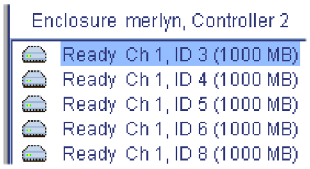
Enclosure merlyn, Controller 2

Ready Ch 1, ID 3 (1000 MB)
Ready Ch 1, ID 4 (1000 MB)
Ready Ch 1, ID 5 (1000 MB)
Ready Ch 1, ID 6 (1000 MB)
Ready Ch 1, ID 8 (1000 MB)

3. Click >> (Add selected drives) to add the drives to the array. You can click >> (Add all drives) to move **all** ready drives to an array.
4. Repeat steps 1 and 2 for each additional array or hot-spare drive that you want to configure.
5. If you do **not** want to create a spanned array, skip to step 5. Otherwise, select the **Span arrays** check box ☑ Span arrays . Then, click **Next**. If you created two arrays only, the ServeRAID Manager uses those arrays to create the spanned array; continue to step 5. Otherwise, the ″ Create spanned arrays″ window opens.

   Note: To create a spanned array, each array must have the same number of physical drives.
6. After you select the ready drives for your arrays and hot-spare drive, click **Next**. The ″ Create logical drives″ window opens.

To return to the ″ Configuration path″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Creating spanned arrays
- Understanding creating spanned arrays

Creating spanned arrays: If you want to assign RAID level-x0 to an array, you must create a spanned array.

Complete the following steps to create one or more identical spanned arrays:

1. In the list of arrays, click the arrays that you want to add to your spanned array.

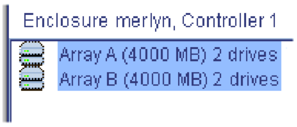
Enclosure merlyn, Controller 1

Array A (4000 MB) 2 drives
Array B (4000 MB) 2 drives

2. Click >> (Add selected arrays) to add the arrays to the spanned array. You can click >> (Add all arrays) to move **all** arrays to the spanned arrays.
3. To create additional spanned arrays, click the **New spanned array** tab in the right pane.

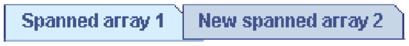
Spanned array 1 | New spanned array 2

   Then, repeat steps 1 and 2.
4. Click **Next**; the ″ Create logical drives″ window opens.

To return to the ″ Create arrays″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Understanding RAID level-x0
- Understanding creating spanned arrays

Creating logical drives: Complete the following steps to create logical drives:

1. Click the appropriate **Array** tab.

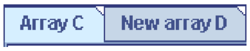

2. Select a RAID level from the drop-down list.

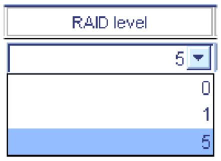

RAID level
5

   - If you are configuring a spanned array, you can set the RAID level only for the first logical drive you create.
   - In addtion to the standard RAID levels (0, 1, 5, x0), you can also create a **simple volume**: a single disk drive, non-redundant.

3. If you do not want to use the maximum size for the logical drive, type the size in the **Data (MB)** field.

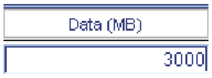

Data (MB)
3000

   Notes:

   a. Some operating systems have size limitations for logical drives. Before you save the configuration, verify that the size of the logical drive is appropriate for your operating system. For more detailed information, see your operating-system documentation.
   b. A logical drive cannot exceed 2048 GB (2 terabytes); the minimum size is 10MB.
   c. Typically, the first logical drive defined on the first controller found by system BIOS during startup will be your startup (boot) drive.
   d. The actual logical-drive size might be slightly different from what you select. The RAID level and the number of physical drives determine the size of the logical drive. For example, an array consisting of three, 1 GB physical drives with a requested RAID level-0 logical drive of 1000 MB will actually contain only 999 MB because the data is striped across all three drives, with 333 MB on each drive.

4. From the **Stripe Size** drop-down list, select the stripe-unit for the logical drive.

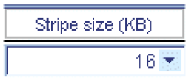

Stripe size (KB)
16

   Note: The default stripe-unit size is usually adequate for most applications.

5. From the **Controller** drop-down list, select the preferred owner for the logical drive, Controller A or Controller B. (For single controller configurations, this field is grayed out.)

Note: You can only choose the controller for the first logical drive in the array.

6. If you have free space available and want to create additional logical drives, click Define new logical drive . To create a new logical drive with the same settings as the current logical drive, click Replicate logical drive ; then, choose the number of replicas from the drop-down list.

   Attention: Use care when allocating space for new a new logical drive. You must leave enough space for FlashCopy backups to grow to their virtual size. For example, if the virtual size of a FlashCopy backup is 20 GB, you must leave at least 20 GB of free space in the array. Note that the Configuration wizard displays the *actual size* of the first FlashCopy backup (not the virtual size), and 0 for all other FlashCopy backups.
7. Repeat steps 2 through 6 for each logical drive that you want to define in this array.
8. Repeat steps 1 through 7 for each additional array you want to configure.
9. Click **Next**. Depending on the initiator type, the ″ Define authentication method″ window opens or the ″ Define initiators″ window opens.

To return to the ″ Create arrays″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

Defining the authentication method:

Note: The ServeRAID Manager supports authentication for iSCSI initiators only.

Use the Authentication information window to define the authentication method (if any) for each user permitted to use the logical drives in the array. Optionally, you can enable Radius authentication (an external authentication service) and define the SLP (Service Location Protocol) Scope name.

1. Select a logical drive from the list on the left.
2. From the **Authentication type** list, select the method used to authenticate users of this logical drive. You can choose:
   - **None** - Do not authenticate users
   - **CHAP** - Challenge Handshake Authentication Protocol
   - **SRP** - Secure Remote Password
3. If you enabled authentication (by choosing CHAP or SRP), click Add User ; the ″Global user name and password management″ window opens. Then, add users to the user list.
4. To enable Radius authentication for the CHAP authentication method, click **Radius authentication**; then, enter the following:
   - In the Primary server field, enter the host name or TCP/IP address of the authentication service; then, in the Port field, enter the server's startup port.
   - In the Secondary server field, enter the host name or TCP/IP address of the authentication service; then, in the Port field, enter the server's startup port.

   Note: Radius authentication is supported by the CHAP authentication method only.
5. To modify the default SLP Scope name, click **Advanced settings**; then, in the SLP Scope Name field, type a unique scope name or choose an existing name from the drop-down list.
6. Repeat steps 1 through 5 to enable authentication for additional logical drives. To replicate the settings for all logical drives, click Replicate settings .

7. When you are ready to continue, click **Next**. The ″ Define initiators″ window opens.

To return to the ″ Create logical drives″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Adding users to the user list

Defining initiator information: Use the Initiator information window to define an alias, or ″friendly″ name, for each initiator that can access the enclosure, including the initiator IQN (iSCSI Qualified Name) or fibre channel port name. An initiator represents a machine on the network that can access the enclosure.

The ServeRAID Manager displays previously defined initiators in a list. The alias is shown on the left. Depending on the initiator type (iSCSI or fibre channel), the IQN or port name is shown on the right. Initiator types are not mixed. You can add a new initiator, modify an initiator, or delete an initiator.

To add an initiator:

1. Click

   Add

   . Depending on the initiator type, the ″Add iSCSI initiator″ window opens or the ″Add Port Name″ window opens.
2. Follow the steps below to add an iSCSI initiator or add a fibre channel port name.
3. When you are ready to continue, click **Next**. The ″ Assign logical drives to initiators″ window opens.

To return to the previous Configuration wizard window, click **Back**.

To leave the Configuration wizard, click **Cancel**.

Adding an iSCSI initiator

1. In the **Initiator IQN** field, type a well-formed IQN in the format: `iqn.yyyy-mm.s1.s2[.|:]s3` where:
   - yyyy is the 4-digit year
   - mm is the 2-digit month (01-12)
   - s1 and s2 are 1-n character alphanumeric strings (for example, *mydomain.com*)
   - s3 is 5 or more alphanumeric characters, separated from s2 by ′.′ or ′:′
2. In the **Initiator alias** filed, type a name for the initiator. The initiator alias must start with a letter. It can contain only letters and numbers.
3. Optionally, click **Advanced settings**; then, click the check boxes to enable or disable the following features:

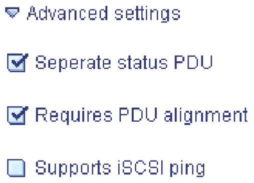

▿ Advanced settings	
☑ Seperate status PDU	Communications with initiator include separate status Protocol Data Unit (PDU).
☑ Requires PDU alignment	Communications with initiator include PDU alignment.
□ Supports iSCSI ping	Initiator supports iSCSI ping.

Note: Refer to your initiator documentation for more information about these features.

4. Click **OK**; then, repeat steps 1-3 to add additional initiators.
5. Click **Cancel** to close the ″Add iSCSI initiator″ window.

Adding a fibre channel port name

1. In the **Port Name** field, enter the 16-character Port Name. Each octet (8 characters) can be separated by a ′:′ or ′-′.

Port Name 5799B44907850444

2. In the **Initiator alias** field, type a name for the initiator. The initiator alias must start with a letter. It can contain only letters and numbers.
3. Click **OK**; then, repeat steps 1 and 2 to add additional initiators.

   Note: You can define a maximum of 255 ports (LUNs) for each fibre channel initiator.
4. To automatically detect initiators on your network, click

   . The ″Select a discovered initiator″ window opens.

   Note: The button will be grayed out if no new initiators are detected.

   Then, select a discovered initiator from the list.
5. Click **Cancel** to close the ″Add Port Name″ window.

More information

- Modifying and deleting initiator information
- Selecting discovered initiators

Selecting a discovered initiator: Use the ″Select a discovered initiator″ window to add an existing fibre channel initiator to the list of initiators that can access the enclosure.

1.

   From the Port Name list, select an initiator.
2. In the **Initiator alias** field, type a name for the initiator or accept the default. The initiator alias must start with a letter. It can contain only letters and numbers.
3. Click **OK**. The ″Select a discovered initiator″ window closes.
4. To select another discovered initiator, click . Then, repeat steps 1-3 to add additional initiators.
5.

Assigning logical drives to initiators:

Note: You can assign up to 32 initiators to each logical drive.

Complete the following steps to assign logical drives to initiators:

1. Click the appropriate **Initiator** tab in the right pane. To allow anyone to access the logical drive, regardless of users in the user list, click the Unrestricted tab.

2. Then, from the list of logical drives, select the drives you want the initiator to access:

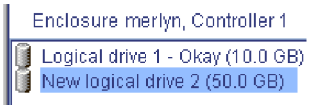

Enclosure merlyn, Controller 1

Logical drive 1 - Okay (10.0 GB)

New logical drive 2 (50.0 GB)

3. Click >> (Add selected drives) to assign the logical drives to the initiator. You can click >> (Add all drives) to assign **all** logical drives to the initiator.
4. Repeat steps 1 and 2 for each initiator.
5. After you assign logical drives to initiators, click **Next**. The ″ Configuration summary″ window opens.

To return to the ″ Define initiators″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

Confirming your system configuration: Use the configuration summary to review all the changes that the ServeRAID Manager will apply to your configuration:

1. Review the information that is displayed in the ″Configuration summary″ window. To change the configuration, click **Back**.

   Note: The Configuration wizard displays the *actual size* of the first FlashCopy backup (not its virtual size), and 0 for all other FlashCopy backups.
2. Click **Yes** when asked if you want to apply the new configuration. A status window opens, displaying the progress of the operation. When finished, the configuration is saved in the controller and physical drives in the enclosure.

   Note: The ServeRAID Manager initializes the logical drives automatically.

More information

- Understanding logical-drive synchronization
- Initializing a logical drive

Configuring NetWare user authentication

NetWare user authentication may be configured in one of two ways: out-of-box authentication and NDS authentication. The two methods are mutually exclusive.

Out-of-box authentication

Out-of-box authentication works as follows:

- Administrative access is exclusively granted to the ‘Admin’ user account, which must be defined within the NetWare Bindery context of the server running the ServeRAID Manager agent.
- User access is granted to all Bindery and NetWare Directory Services (NDS) user accounts. For NDS users, the user account must reside within the same tree as the server running the ServeRAID Manager agent.

When logging into the ServeRAID Manager, NDS users must provide the fully distinguished account name.

NDS authentication

NDS authentication is based on NDS group membership. To set up NDS group membership, create and populate a NDS group for storage administration. After you install the ServeRAID Manager, edit the ServeRAID Manager agent properties file to restrict authentication to the new group.

For example:

1. Choose the group name ″.CN=StorageAdmins.O=Acme″.

   Note: The group must reside within the same NDS tree as the server running the ServeRAID Manager agent.
2. Using the Novell Administration utility, add the appropriate users to the new group.
3. Add the following entry to the SYS:\RAIDMAN\RaidAgent.pps file:

   `agent.group.NetWareStorageAdmin=<` *`fully distinguished group name`*`>`The new entry should look like this :

   ```
   agent.group.NetWareStorageAdmin=.CN=StorageAdmins.O=Acme
   ```
4. Restart the server.

Limited user access is granted to all NDS user accounts.

More information

- Logging in to the ServeRAID Manager
- Adding a remote system
- Configuring the ServeRAID Manager agent

Novell NetWare considerations: (NetWare 5.x only) If you are preparing to install Novell NetWare 5.x from the startable Novell NetWare 5.x CD, set the write-cache mode *only* to write-through mode.

Attention: Do not attempt a hot-replace operation on Windows 2000, Windows Server 2003, or NetWare by hot-removing a failed controller and then hot-adding a new controller. Loss of data can occur. If a controller fails on these operating systems, you must shut down the server to replace the controller.

Copying the configuration from drives

Note:

1. This action is not supported on the integrated RAID controller.
2. (Cluster environment only) Use this action to merge non-shared logical drives.

Use this action to copy the configuration from all physical drives into a ServeRAID controller. This action is useful when replacing a controller and you want to reestablish the original configuration on the new controller.

If you are replacing a controller in a cluster environment, you must first configure the cluster parameters by using the ″ Configure for clustering″ action. The cluster parameters are the following:

- The controller name
- The partner controller name
- The SCSI initiator IDs

After you have configured these parameters, you can use "Copy configuration from drives" to copy the configuration from only the non-shared drives into the new controller.

1. In the Enterprise view, click (controller).
2. Right-click **Copy configuration from drives**.
3. If clustering is not enabled, go to step 4. Otherwise, right-click one of the following from the menu:
   - **Non-shared 206**
   - **Non-shared 207**
4. A confirmation window opens; click **Yes**.

Restore the configuration to the factory-default settings

Note: This action is not supported on the following:

1. Integrated RAID controller
2. ServeRAID-7t controller
3. ServeRAID-8i controller
4. HostRAID controller
5. Network storage controllers

Use this action to restore all parameters in the configuration to the factory-default settings:

1. In the Enterprise view, click (controller).
2. Right-click **Restore to factory-default settings**. A confirmation window opens.
3. Click **Yes**.

   Note: In the following step, the ServeRAID Manager will destroy all data on these logical drives.
4. If there **are** logical drives defined, click **Yes** to delete all logical drives and restore all parameters in the configuration to the factory defaults. If there are **no** logical drives defined, the ServeRAID Manager automatically restores all parameters in the configuration to the factory-default settings.

Saving a printable configuration and event logs

Use this action to save a report of your configuration and event logs that you can print.

1. In the Enterprise view, click (server).
2. Right-click **ServeRAID Manager Actions** → **Save printable configuration and event logs**. The ServeRAID Manager creates one text file for each controller. The file name is Raid **x**.log where **x** is a number assigned to the file.

   Note: If you are in bootable-CD mode, you can save the files to a diskette in the A drive only.

The ServeRAID Manager copies to each text file all configuration information for the following:

- The controller in the selected system
- Any physical drive on the controller
- Any array on the controller
- Any logical drive on the controller
- Device event logs
- defunct drive event logs
- Software event logs

- Hardware event logs

The ServeRAID Manager saves all events currently in the event log, not only the events associated with the selected system.

More information

- Example: Printable configuration
- Saving a printable configuration

Saving a printable configuration: Use this action to save a report of your configuration that you can print.

1. In the Enterprise view, click 🖥 (system).
2. Right-click **Save printable configuration**. The "Save file" window opens.
3. Choose the directory and file name for your report. The default directory is the directory in which the ServeRAID Manager is installed. The default file name is RaidCfg.log.

   Note: If you are in bootable-CD mode, you can save to a diskette in the A drive only.

The ServeRAID Manager copies to a text file all configuration information for the following:

- Each controller in the selected system
- All physical drives on the controller
- All arrays on the controller
- All logical drives on the controller

More information

- Example: Printable configuration
- Saving a printable configuration and event logs

Example: Printable configuration:

```
July 13, 2001 7:30:58 AM EDT

   Configuration summary
   ---------------------------

   Server name.....................myserver
   ServeRAID Manager Version.......4.80.xx
   Number of controllers...........1
   Operating system................Windows 2000

   Information for controller 1
   --------------------------------------------------------
   Controller type.................ServeRAID-4Mx
   BIOS version....................4.80.xx
   Firmware version................4.80.xx
   Device driver version...........4.80.xx
   Physical slot...................4
   Battery-backup cache............Installed
   Read-ahead cache mode...........Adaptive
   stripe-unit size................8K
   Rebuild rate....................High
   Hot-swap rebuild................Enabled
   Data scrubbing..................Enabled
   Auto-synchronization............Enabled
   Clustering......................Disabled
   Unattended mode.................Disabled
   BIOS compatibility mapping......Extended
```

```
Number of arrays................2
Number of logical drives........2
Number of hot-spare drives......0
Number of ready drives..........0

Array A
--------------------
Array identifier................A
Array size in MB................4157
Free space in MB................0
Number of logical drives........1
Stripe order (Channel/Device)...1/4
Number of physical drives.......1

Logical drives in array A
--------------------------------
Logical drive...................1
Array letter....................A
State...........................Okay
RAID level......................0
Data space in MB................4157
Parity space in MB..............0
Date created....................7/13/01
Write-cache mode................Write back
Merge group number..............207
Merge group state...............Non-shared


Physical drives in array A
--------------------------------
Type............................Hard disk drive
Channel.........................1
SCSI ID.........................4
Vendor..........................IBM_PSG
Product or model number.........xxxxxxxx
Serial number...................xxxxxxxx
Firmware level..................1.00
Size in MB......................4157
State...........................Online
Array letter....................A
PFA error.......................No

Array B
--------------------
Array identifier................B
Array size in MB................6143
Free space in MB................2
Number of logical drives........1
Stripe order (Channel/Device)...1/0 1/1 1/2
Number of physical drives.......3

Logical drives in array B
--------------------------------
Logical drive...................2
Array letter....................B
State...........................Okay
RAID level......................5
Data space in MB................4094
Parity space in MB..............2047
Date created....................7/13/01
Write-cache mode................Write back
Merge group number..............207
Merge group state...............Non-shared


Physical drives in array B
--------------------------------
```

```
Type............................Hard disk drive
Channel.........................1
SCSI ID.........................0
Vendor..........................IBM_PSG
Product or model number.........xxxxxxx
Serial number...................xxxxxxx
Firmware level..................1.00
Size in MB......................2047
State...........................Online
Array letter....................B
PFA error.......................No

Type............................Hard disk drive
Channel.........................1
SCSI ID.........................1
Vendor..........................IBM_PSG
Product or model number.........xxxxxxx
Serial number...................xxxxxxx
Firmware level..................1.00
Size in MB......................2047
State...........................Online
Array letter....................B
PFA error.......................No

Type............................Hard disk drive
Channel.........................1
SCSI ID.........................2
Vendor..........................IBM_PSG
Product or model number.........xxxxxxx
Serial number...................xxxxxxx
Firmware level..................1.00
Size in MB......................2047
State...........................Online
Array letter....................B
PFA error.......................No


SCSI channel 1
-------------------
Type............................Hard disk drive
Channel.........................1
SCSI ID.........................0
Vendor..........................IBM_PSG
Product or model number.........xxxxxxx
Serial number...................xxxxxxx
Firmware level..................1.00
Size in MB......................2047
State...........................Online
Array letter....................B
PFA error.......................No

Type............................Hard disk drive
Channel.........................1
SCSI ID.........................1
Vendor..........................IBM_PSG
Product or model number.........xxxxxxx
Serial number...................xxxxxxx
Firmware level..................1.00
Size in MB......................2047
State...........................Online
Array letter....................B
PFA error.......................No

Type............................Hard disk drive
Channel.........................1
SCSI ID.........................2
Vendor..........................IBM_PSG
```

```
Product or model number.........xxxxxxx
Serial number...................xxxxxxx
Firmware level..................1.00
Size in MB......................2047
State...........................Online
Array letter....................B
PFA error.......................No

Type............................Hard disk drive
Channel.........................1
SCSI ID.........................4
Vendor..........................IBM_PSG
Product or model number.........xxxxxxx
Serial number...................xxxxxxx
Firmware level..................1.00
Size in MB......................4157
State...........................Online
Array letter....................A
PFA error.......................No

Type............................Enclosure
Channel.........................1
SCSI ID.........................14
Vendor..........................SDR
Product or model number.........GEM200
Serial number...................0
Firmware level..................2

SCSI channel 2
-------------------

End of the configuration information for controller 1
--------------------------------------------------------
```

More information

- Saving a printable configuration (action)
- Saving a printable configuration and event logs (action)

Working with systems in the ServeRAID Manager

You can use the ServeRAID Manager to view information about managed systems, ServeRAID controllers, and the ServeRAID Manager subsystem (such as arrays, logical drives, hot-spare drives, and physical drives).

To view information, click an object in the Enterprise view or the Physical or Logical device views; then, click ☰ (Properties) on the toolbar. Using preferences, you can sort tree objects alphabetically or chronologically.

Note: The local system always appears first when you sort the tree alphabetically.

To display available actions for an item, right-click the item in the Enterprise view, Physical devices view, or Logical devices view.

You can use the following actions and applications in the ServeRAID Manager to modify or monitor other ServeRAID Manager systems in a network:

- Add or remove a remote system
- Notification Manager
- Email Notification Manager
- Task Manager
- ServeRAID Manager agent

To configure a new ServeRAID controller with the Configuration wizard, click (Create) on the toolbar or (Create) in the Logical devices view.

More information

- The ServeRAID Manager menubar
- The ServeRAID Manager toolbar
- Using the ServeRAID Manager interface
- Changing the ServeRAID Manager interface
- Finding information in ServeRAID Manager
- Using ServeRAID Manager Assist, hints and tips

Using the ServeRAID Manager interface > Menubar: In addition to using the mouse, you can use keyboard shortcuts to access the menubar. Selections from the Actions menu also are available when you right-click an object in the Enterprise view, Physical devices view, and Logical devices view.

Note: No actions from the Actions menu are available during any of the following operations:

- Rebuild
- Foreground synchronization
- Verification
- Logical-drive migration
- RAID level-5E compression or decompression
- RAID level-5EE compaction or expansion
- FlashCopy with backup option
- Copy back

More information

- The ServeRAID Manager menubar
- The toolbar
- The Enterprise view
- The Physical and Logical device views
- The event viewer
- The status bar

The ServeRAID Manager menu bar: The ServeRAID Manager menu bar has the following functions:

File >

Clear configuration event log
Clears the current contents of the event viewer. This option does not clear or delete the event logging file.

Preferences
Opens a window that you can use to specify the following settings:

- Remote access
- Alarm
- Display options

Close help
Closes the Help window

Exit Exits the ServeRAID Manager.

View >

Toolbar
Turns the toolbar on and off. The default is on (that is, selected).

Status bar
Turns the status bar on and off. The default is on (that is, selected).

Tool tips
Turns the tool tips on and off. The default is on (that is, selected).

Refresh
Refreshes the display with the latest configuration information.

Remote >

Add remote system
Adds a remote system to the Enterprise view for monitoring and configuration. This action is disabled in bootable-CD mode and when the ServeRAID Manager is a plug-in to another program.

Remove remote system
Displays a menu of all the remote systems currently in your Enterprise view. You can select one remote system to remove at a time. This action is disabled in bootable-CD mode and when the ServeRAID Manager is a plug-in to another program.

Actions >
Display any action that is valid for the selected object in the Enterprise view, the Physical devices view, or Logical devices view.

Note: If there are no valid actions for the object you selected, the Actions menu is grayed out.

Help >

Information about this window
View context-sensitive information for the current window.

Search
Searches for one or more specified words in ServeRAID Assist and displays a list of topics that include the words.

Contents
Presents the ServeRAID Assist contents. You can use the contents to acquaint yourself with ServeRAID Assist topics.

General Concepts
Presents ServeRAID concepts that you use with the ServeRAID Manager.

Publications
Lists IBM publications and where to find them.

IBM online support
Lists IBM Web sites and available online support.

What's new
Lists the new features in this version of the ServeRAID Manager program.

About ServeRAID Manager
Reports the ServeRAID Manager version number, copyright, and legal information.

The ServeRAID Manager toolbar: The ServeRAID Manager toolbar has the following functions:

Add

Add a remote system to the Enterprise view for monitoring and configuration.

Note: This action is not supported when using the following:
- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Create

Create a logical drive with the Configuration wizard. Choose Express or Custom configuration, if there are ready drives and you have not assigned the maximum number of logical drives.

Silence

Silence the console alarm and agent alarm on all connected systems. All Warning and Error events produce an audible alarm every five minutes, notifying you of the event. You can adjust the alarm interval in the ″Preferences″ window.

Properties

When a system, controller, or device is selected, displays system, controller, or device properties.

Events

Open the event log in a stand-alone viewer. The stand-alone viewer shows all ServeRAID Manager events. Use the embedded event viewer to view events for the current ServeRAID Manager session.

Configure

List tools for configuring the local or remote ServeRAID Manager agent or management station agent. To use a tool, select it from the drop-down list.

Help

View context-sensitive information for the current window.

Using the ServeRAID Manager interface > Toolbar: The toolbar provides quick-access icons for common tasks.

More information

- The menubar
- The Enterprise view

- The Physical and Logical device views
- The event viewer
- The status bar

Using the ServeRAID Manager interface > Enterprise view: The Enterprise view provides an expandable "tree" view of the systems, controllers, and enclosures you are managing. It is divided into two parts:

- **Direct attached storage** - managed systems with directly attached controllers and storage devices.
- **Networked storage** - management stations with network-attached storage enclosures.

You can perform most configuration and management tasks by selecting a controller or enclosure from the tree and working with related objects in the Physical and Logical device views.

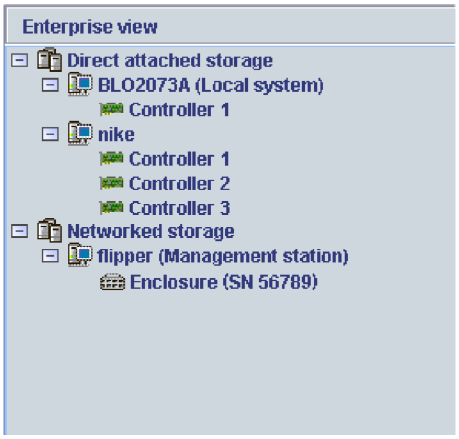

Tip: You can use display groups to group related systems under single tree object.

More information

- The menubar
- The toolbar
- The Physical and Logical device views
- The event viewer
- The status bar

Using the ServeRAID Manager interface > Physical and Logical device views: The Physical and Logical device views show the physical devices and logical devices connected to the controller or enclosure.

The **Physical devices** view (on the left) displays information about the drives, enclosures, and other devices attached to the controller. The devices are shown organized by the channel or port they are connected to and shown in numerical order. The display for each channel or port includes information on maximum speed capability, the number of the channel on the controller, and the number of devices attached to the controller.

The **Logical devices** view (on the right) displays information about the arrays and logical drives created using the physical devices. This information includes the number of arrays and logical drives, the RAID level of each device, and whether a logical drive is protected by a hot spare drive.

To view device information, click ⊞ in the Enterprise view to expand a managed system or management station. Then, select a controller or enclosure from the tree.

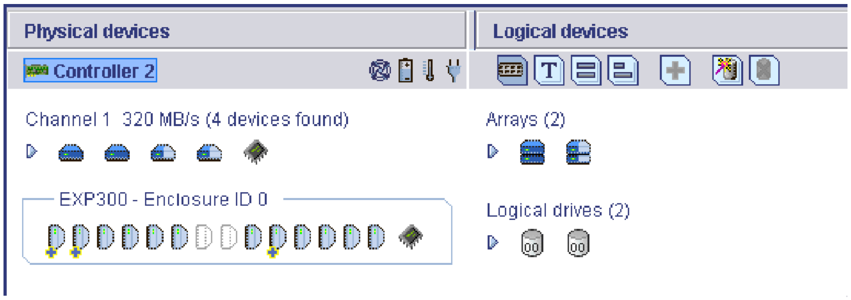

For network storage, the Physical devices view also shows the controllers in the enclosure. All other information is the same.

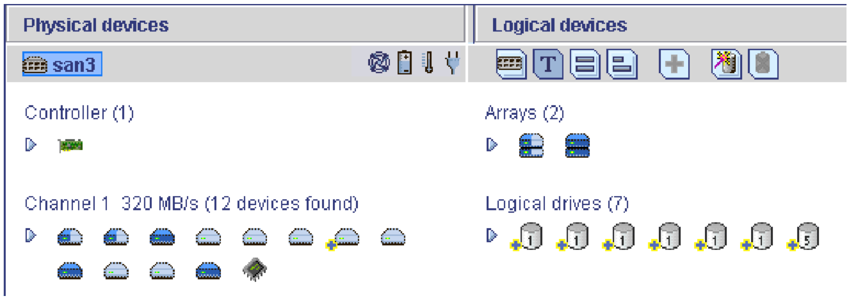

Tip: You can drag the bar between the Physical devices view and Logical devices view to adjust the size of each pane.

In the Physical devices view, four indicators report status of the fan, battery, temperature, and power modules on SAF-TE (enclosure management) devices and other devices that monitor these conditions. The indicator is blue for normal, yellow for warning, and red for error; the indicator is grayed out if none of the devices on the controller monitor the condition. Example: the fan indicator changes to yellow when one fan fails; it changes to red when a second fan fails and cooling is no longer adequate.

In the Logical devices view, use the buttons to create and delete arrays and logical drives. The Create option ( ) opens a wizard that presents the steps necessary to create a new array.

Other buttons in the Logical devices view allow you to:
- Change the way configuration information is displayed
- Collapse and expand configuration information
- Identify components of an array or logical drive

You can also create and delete hot-spare drives.

More information

- The menubar
- The toolbar
- The Enterprise view
- The event viewer
- The status bar

Using the ServeRAID Manager interface > Event viewer: The event viewer provides advisory and progressive-status information and messages for your ServeRAID Manager system. Every event includes a severity icon, a date and time stamp, the host name where the event originated, and a description.

Note: All Warning and Error events cause an audible alarm to sound, notifying you of the event.

You can double-click an event in the event viewer and the ″Configuration event detail″ window opens. This window contains the same information as the event viewer, but in a larger, easier-to-read window. Some events have additional details listed in this window.

If the ServeRAID Manager is not in bootable-CD mode, it appends each event to a log file, RAIDEVT.LOG. If this file exceeds 200 KB, the ServeRAID Manager copies the file to RAIDEVT.OLD and creates a new RAIDEVT.LOG. If there is a RAIDEVT.OLD already, the ServeRAID Manager overwrites it.

More information

- The menubar
- The toolbar
- The Enterprise view
- The Physical and Logical device views
- The status bar

Using the ServeRAID Manager interface > Status bar: The status bar displays the following information from left to right. You can drag the vertical bars to the right and left to adjust the size of the viewing areas.

Managed systems status icon. This icon is the same as the Managed systems icon displayed in the Enterprise view. The text next to the icon states whether the ServeRAID Manager has detected problems on any of the systems. You can see the text by dragging the vertical bar to the right.

Tree path. Reports the tree path location of the selected object.

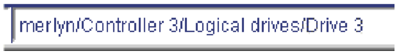

merlyn/Controller 3/Logical drives/Drive 3

Progress indicator (ServeRAID legacy controllers only). A labeled progress indicator is displayed in this pane if the currently selected system has one or more controllers containing a logical drive undergoing any of the following operations:

- Rebuild
- Migration
- Synchronization (either foreground or background)
- Compression or decompression (RAID level-5E only)
- Compaction or expansion (RAID level-5EE only)
- FlashCopy
- Copy back

SR-8i SR-7t HostRAID For ServeRAID-8i, ServeRAID-7t, and HostRAID controllers, the ServeRAID Manager displays (in animation) in the Logical devices view, instead of the progress indicator.

If no action is in progress, this pane is empty. The progress indicator is labeled with the controller number and logical drive number. If more than one controller is undergoing one of these operations, the progress of the selected controller is displayed in the pane.

More information

- The menubar
- The toolbar
- The Enterprise view
- The Physical and Logical device views
- The event viewer

Physical and Logical device views > Icons, buttons, and status indicators: The ServeRAID Manager displays the following icons, buttons, and status indicators in the Physical and Logical device views, representing the physical and logical devices in your system, their status, and the actions you can take to manage them.

Icons

Controller card
Array with no free space
Enclosure management device
Array with free space
Enclosure
Logical drive
Ready hard drive
Logical drive with hot spare
Hard drive with no free space
Creating/modifying array or logical drive
Hard drive with free space
Initialize logical drive
Hot spare drive protecting logical drive
FlashCopy backup
Hot spare drive **not** protecting logical drive
CD-ROM drive
defunct drive
Removable drive
Tape drive

Buttons

Use the following buttons to view physical and logical device information and to create and delete arrays and hot spares.

To view physical device information in enclosure view format (visible only if system has an attached storage enclosure)

T To view physical device information in text format

To view physical device information in full size capacity format

To view physical device information in relative size capacity format

⊕	To create or delete a hot spare drive
✎	To create an array or logical drive
✖	To delete an array or logical drive
⊞ ⊟	To expand and contract information in the Enterprise view
▷ ▽	To expand and contract physical and logical device information

Status indicators

Use the following indicators to monitor the status of controllers and SAF-TE enclosure management devices in your system.

⊗	To monitor status of the fan module
🔋	To monitor status of the battery module
🌡	To monitor temperature status
🔌	To monitor status of the power module

More information

- Physical device view options
- Component views
- Collapsed and expanded views

Physical and Logical device views > Collapsed and expanded views: In addition to changing the way physical device information is displayed in the Physical devices view (either textual or graphical), you can also view a collapsed or expanded view of the system configuration information. Initially, the ServeRAID Manager displays a collapsed textual view of the configuration information in both the Logical devices view and Physical devices view.

▷ ▽ In Logical devices view	Click to expand and collapse information about arrays and logical drives. The expanded display shows the following information about each logical device: • Array name and capacity (if available) • Logical drive size • Logical drive state • Build progress

▷ ▽

In Physical devices view

Click to expand and collapse information about physical drives. The expanded display shows the following information about each drive:

- Capacity of the drive
- Drive ID
- Drive state

More information

- Icons, buttons, and status indicators
- Physical device view options
- Component views

Physical and Logical device views > Physical device view options: You can choose the way information is displayed in the Physical devices view. To choose a view option, click the associated button in the Logical devices view.

T

Displays physical device information in text format. This is the default view for controllers with direct attached storage devices.

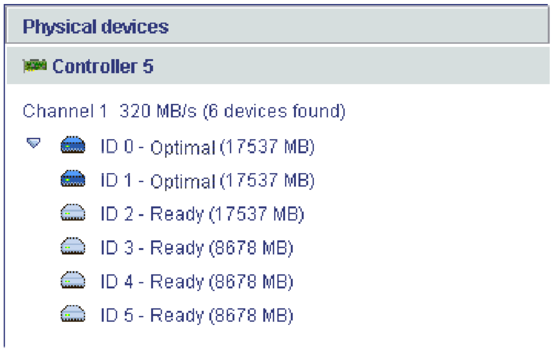

Physical devices

Controller 5

Channel 1 320 MB/s (6 devices found)

ID 0 - Optimal (17537 MB)
ID 1 - Optimal (17537 MB)
ID 2 - Ready (17537 MB)
ID 3 - Ready (8678 MB)
ID 4 - Ready (8678 MB)
ID 5 - Ready (8678 MB)

☰

Displays physical device information in full size capacity format. A full-length bar is displayed for each drive, regardless of capacity. A small segment on each drive is reserved for the RAID signature; this area is indicated by a gray ″cap″ at the end of each bar.

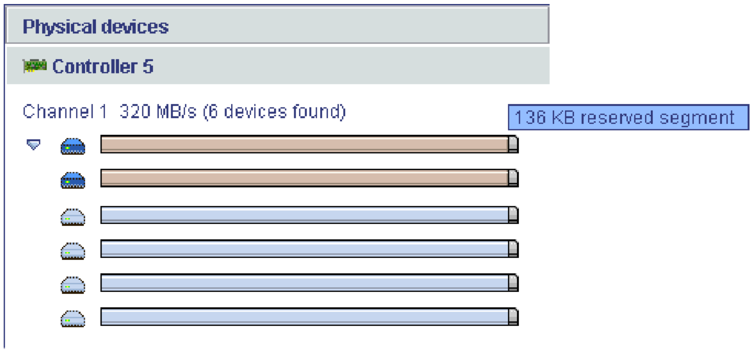

Physical devices

Controller 5

Channel 1 320 MB/s (6 devices found)

136 KB reserved segment

Note: A drive shaded in light blue is not part of any array.

Displays physical device information in relative size capacity format. A full length bar is displayed for the largest drive. Proportionally shorter bars are displayed for other drives.

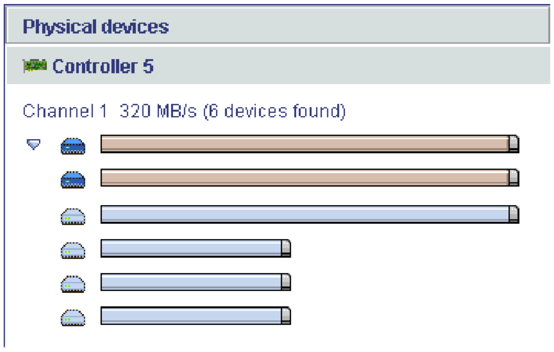

Displays physical device information in enclosure view format. Drives in the enclosure are shown in the physical slots they occupy with the proper vertical or horizontal orientation. Empty slots are shown as drive outlines. This is the default view for systems with an attached storage enclosure.

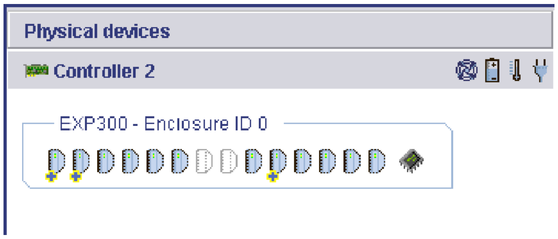

For network storage, enclosure view also shows the controllers in the enclosure.

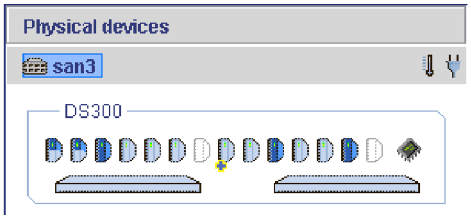

Attention: The enclosure view button is visible in the Logical devices view only if an enclosure is attached to the system. HostRAID controllers do *not* support enclosure view. Nor do some older enclosures.

More information

- Icons, buttons, and status indicators
- Component views
- Collapsed and expanded views

Physical and Logical device views > Component views: When you click either a physical or logical device in the device views, the related components are highlighted in the other view.

When you click an array, the associated logical drives are highlighted in the Logical devices view and the physical drives that are members of the array are highlighted in the Physical devices view.

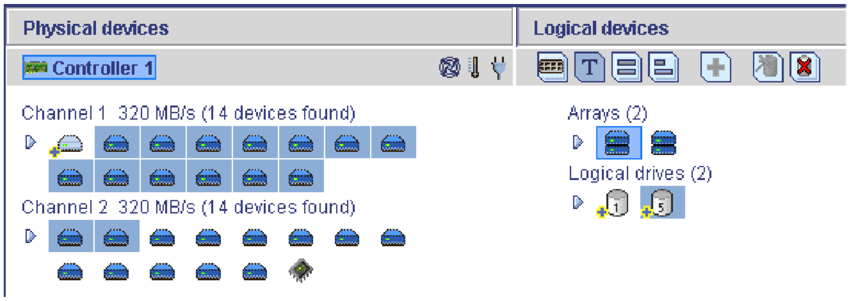
Physical devices
Controller 1
Logical devices
Channel 1 320 MB/s (14 devices found)
Channel 2 320 MB/s (14 devices found)
Arrays (2)
Logical drives (2)

When you click a hot spare drive, the logical drives protected by that spare are highlighted.

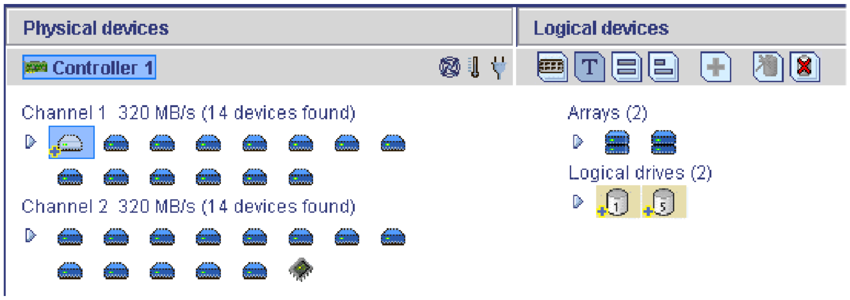
Physical devices
Controller 1
Logical devices
Channel 1 320 MB/s (14 devices found)
Channel 2 320 MB/s (14 devices found)
Arrays (2)
Logical drives (2)

In the graphical views, if the logical drive uses only part of the available storage, only those segments are highlighted, using the following color code:

- dark blue - storage used by selected logical drive
- brown - storage used by other logical drives
- light blue - free space (not used by any logical drive)
- gray - space reserved for RAID signature

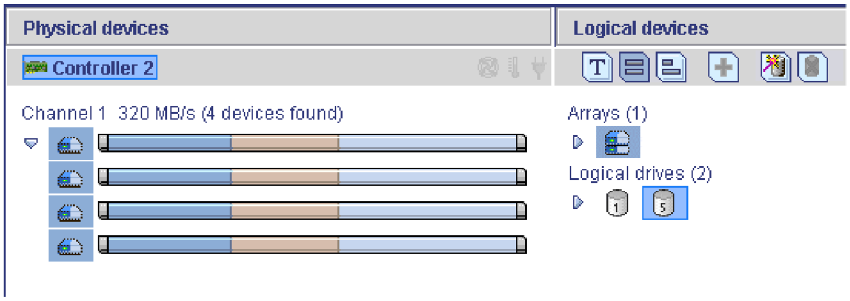

Note: Partially filled hard drive and array icons indicate that the device has free space.

More information

- Icons, buttons, and status indicators
- Physical device view options
- Collapsed and expanded views

Sending a test trap from the SNMP Trap Manager: Use this action to send a test trap to a remote system.

1. Click a system in the SNMP traps list.

   Note: This action is disabled if you do not select a system.
2. Click **Actions** → **Send test trap**.

Note: The SNMP Trap Manager cannot verify that the trap was received successfully. You must verify that the trap was received on the remote system. If the test fails, verify the following and try the test send again:

1. You typed the correct TCP/IP information when you added the system to the traps list.
2. The SNMP console is running on the remote system.

More information

- Using the SNMP Trap Manager
- Adding a system to the SNMP traps list
- Deleting a system from the SNMP traps list
- Modifying system properties in the SNMP traps list

Working with objects in the Enterprise view and device views: Use objects in the Enterprise view and the Physical and Logical device views to review information and to change or manage the ServeRAID configuration.

- Direct attached storage object
- System object
- Display group object
- Controller object
- Integrated RAID controller object
- Networked storage object
- Management station object
- Enclosure object
- arrays object
- array [letter] object
- logical drives object
- logical drive [number] object
- Channel [number] object
- Ports [numbers] object
- Physical drive [number] object
- Hot-spare drive object
- Enclosure management device object
- CD-ROM drive object
- Removable-media drive object
- Tape drive object

Direct attached storage object: Use this object, in the Enterprise view, to work with managed systems with direct attached storage devices or to display their properties. When you click this object, the ServeRAID Manager displays a summary of managed systems in the **Direct attached storage** view.

Possible subobject

- (System)
- (Display group)

Warning and error conditions

- ⚠ (warning) is displayed if any system has a warning or error condition.

Actions

- Add remote system
- Remove a remote system
- Update controller images

Adding a remote system:

Note: This action is not supported when using the following:
- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to connect to a remote system or management station and add it to the Enterprise view.

Tip: Before adding a remote system, verify that the system is running the required software. To add a remote system with direct-attached storage, the ServeRAID Manager must be running on that system. To add a management station with a network-attached storage enclosure, the ServeRAID Manager management station software must be running on that system.

1. From the Remote menu, select **Remote** → **Add** or click (Add) on the toolbar. The Add managed system window opens.
2. From the Type drop-down list, select:
   - **Managed System**, to add a remote system with direct-attached storage
   - **Management Station**, to add a remote management station with a network-attached storage enclosure
3. Type the remote host name or TCP/IP address.
4. (Managed system only) Type the remote system startup port number. The default port number is 34571.
5. Type your user name and password.

   Note: The password is case sensitive.
6. If you want to save the user name and password, select the **Save user name/Password** box. Once you successfully connect to the remote system, the ServeRAID Manager stores this information in a file along with the host name so you do not have to type your user name and password every time you run the ServeRAID Manager.
7. Click **Connect**.

More information

- Failing to add a remote system
- Successfully adding a remote system
- Specifying remote access settings
- Removing a remote system

Failing to add a remote system: If you fail to connect to the remote system, the ServeRAID Manager displays an error message and does not add the remote system to the Enterprise view. If the ServeRAID Manager fails to connect, verify the following:

- The ServeRAID Manager is running on the remote system.
- The remote system is using a compatible version of the ServeRAID Manager.
- The remote system has TCP/IP installed.

- You supplied the correct user name and password, if security is enabled on the remote system.
- The remote system is turned on.
- The host name is defined in the Domain Name Server or a Hosts file, if you are trying to connect using a host name.
- The remote session is running on the appropriate port.

More information

- Adding a remote system

Successfully adding a remote system: When you successfully add a remote system, the ServeRAID Manager updates the Enterprise view with the newly added system. The next time you start the ServeRAID Manager, it automatically loads the systems into the Enterprise view. When you click one of these systems, the ServeRAID Manager either:

- Opens the ″Add managed system″ window and prompts for the security information, if the remote ServeRAID Manager has security enabled and you did not save the user name and password.
- Connects to the system and retrieves the system information, if you saved the user name and password, or security is disabled.

More information

- Adding a remote system

Specifying remote access settings:

Note: This action is not supported when using the following:
- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Local-only mode

Complete the following steps to run the ServeRAID Manager in local-only mode:

1. In the ″User preferences″ window, click the **Remote access settings** tab.
2. Select the **Local only** box.
3. Click **OK**.
4. Restart the ServeRAID Manager for this setting to take effect.

   The next time you start the ServeRAID Manager, TCP/IP networking will be disabled. In local-only mode, you cannot monitor any remote systems and no remote systems can monitor your system.

Networking mode

Complete the following steps to run the ServeRAID Manager in networking mode:

1. In the ″User preferences″ window, click the **Remote access settings** tab.
2. If the **Local only** box is selected, deselect it.
3. Click **OK**.
4. Restart the ServeRAID Manager for this setting to take effect.

Note: The default port numbers for remote access are 34571-34574. If you have a conflict with these numbers, you must configure the agent to use different ports.

More information

- Adding a remote system
- Removing a remote system
- Configuring the ServeRAID Manager agent

Removing a remote system:

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Removing a remote system does not take a remote system offline, but only removes it from the ServeRAID Manager Enterprise view.

1. From the Remote menu, click either:
   - **Remote** → **Remove managed system** and then the system name, to remove a remote system with direct-attached storage
   -
   - **Remote** → **Remove management station** and then the system name, to remove a remote management station with network storage
2. If you are removing a remote system with direct-attached storage, the ″Remove managed system″ window opens; continue to step 2. Otherwise, skip to step 3.
3. If you want to continue receiving events from the remote system after having removed it, click ″Continue to receive events from the remote system″ from the drop-down list.
4. Click **OK**.

More information

- Receiving events from a removed system
- Adding a remote system
- Specifying remote access settings

System object: Use this object, in the Enterprise view, to work with a local or remote system, or to display its properties.

If a system object is gray, the ServeRAID agent is not available on that system.

Note: The ServeRAID software supports a maximum of 16 ServeRAID controllers.

Possible subobject

- (ServeRAID controller)
- (Integrated RAID controller)

Warning and error conditions

- ⚠ (warning) is displayed if any controller in the system has any problem.

Actions

Note: Not all actions are supported on all controllers.

- If the system is remote, remove remote system.
- Agent actions:

  - Configure (not available in bootable-CD mode):
  -
    - Notifications
    - Email Notifications
    - Task Manager
    - General settings
  - View event log
- Save printable configuration
- ServeRAID actions:
  - Save printable configuration and event logs
  - Clear all controller event logs
  - Validate cluster
- Change display group
- Save support archive

Display group object: Use this object, in the Enterprise view, to work with a display group.

Possible subobject

- 🖥 (System)
- 🖥 (Management station)

Warning and error conditions

- ⚠ (warning) is displayed if any system in the display group has any problem.

Actions

- Delete display group
- Rename display group

ServeRAID controller object: Use this object, in the Enterprise view, to work with ServeRAID controllers and related objects in the Physical and Logical device views, or to display controller properties.

Possible subobject

- Arrays (Arrays)
- Logical drives (Logical drives)
- Channel (Channel [number])
- Ports (Port [numbers])

Warning and error conditions

- ⚠ (warning) is displayed for the following conditions:
  - Any subobject of the controller has any warning or error conditions.
  - The controller has a battery-backup cache problem.
- ✖ (error) is displayed for the following conditions:
  - The controller has failed.
  - The controller has a bad configuration.

Actions

Note: Not all actions are supported on all controllers.

- Configure RAID
- Restore to the factory-default settings
- Copy configuration from drives
- Delete all arrays
- Create logical drive in array [number]
- Scan for new or removed ready drives
- Enable or disable copy back mode
- Identify all physical drives
- Clustering actions →
  - Configure for clustering
  - View shared drives
  - Fail from active to passive controller
- Replace a hot-swap controller

The following actions are available only on the ServeRAID-7t, ServeRAID-8i, and network-attached storage enclosures:

- Enable data scrubbing
- Change data scrubbing rate

The following actions are available only in bootable-CD mode:

- Enable or disable unattended mode
- Change rebuild rate
- Change stripe-unit size
- Enable or disable read-ahead cache mode
- Change BIOS-compatibility mapping

Integrated RAID controller object: Use this object, in the Enterprise view, to work with integrated RAID controllers and their subobjects or to display their properties.

The integrated RAID controller is displayed in the ServeRAID Manager Enterprise view as **two** controller objects. However, you can configure only one array and RAID level-1 logical drive on an integrated RAID controller. The controller information in the tree lists the controller type as LSI 1030.

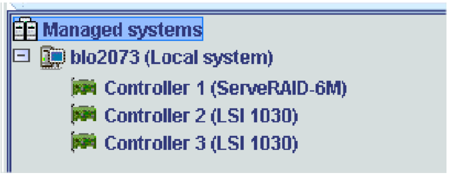

Managed systems
blo2073 (Local system)
Controller 1 (ServeRAID-6M)
Controller 2 (LSI 1030)
Controller 3 (LSI 1030)

Note:

1. The integrated RAID controller provides limited function compared to ServeRAID controllers.
2. The ServeRAID Manager can display a maximum of four LSI 1030 controller objects in a system.

3. Under Windows, LSI 1030 controller objects are not visible unless devices are attached to one of the LSI 1030 controller channels. If devices are attached to only one channel then only one controller object will be visible.

Possible subobject

- Arrays (Arrays)
- Logical drives (Logical drives)
- Channel (Channel [number])

Warning and error conditions

- ⚠ (warning) is displayed for the following conditions:
  - Any subobject of the controller has any warning or error conditions.
- ✖ (error) is displayed for the following conditions:
  - The controller has failed.
  - The controller has a bad configuration.

Actions

- Identify all physical drives

The following actions are available only in bootable-CD mode:
- Configure RAID
- Delete all arrays
- Enable or disable unattended mode

Enabling and disabling unattended mode:

Note: This action is supported in bootable-CD mode only.

Use this action to change the unattended-mode setting. This setting determines how the ServeRAID Manager startup code (BIOS) handles failures during a system startup. The default is Disabled.

1. In the Enterprise view, click (controller).
2. Right-click **Enable or Disable unattended mode.**
3. Click **Yes** to confirm.

More information

- Understanding unattended mode

Identifying devices: You can identify the devices in a server or enclosure, such as a physical drive, because the light associated with the device you selected is flashing.

1. Click any object in the Physical devices view or Logical devices view, **except** the following:

   (Controller)

   (CD-ROM drive)

   (Tape drive)

   (Removable-media drive)
2. Right-click the appropriate **Identify**.

The ServeRAID Manager sends a command to the firmware (microcode) to start flashing one or more physical drive or enclosure lights, depending on the following:

- The selected object
- Whether the item and its subsystems are all controlled by an enclosure device

Note: Identify works only on controllers with a SAF-TE processor and the appropriate firmware. If a device light or a series of device lights do not flash, it is because all or part of the subsystem is not controlled by a SAF-TE processor.

3. Click **OK** to stop the drive lights from flashing.

More information

- Using Identify to work with your systems
- ServeRAID Manager and LED flash states

Using Identify to work with your systems: When you click **Identify** for an object in the Physical devices view or Logical devices view, the ServeRAID Manager makes the hardware light associated with the selected object **flash**.

Depending on what you clicked, you can identify the following:

- A drive
- All drives
- A channel of drives
- A logical drive
- All logical drives
- An array
- A spanned array
- All arrays
- An enclosure

You can make only one selection at a time. The hardware will continue flashing until you click **OK** in the message window.

Identify works only on controllers with a SAF-TE processor and the appropriate firmware. The processor uses a SCSI ID; therefore, it is displayed in the Physical devices view as the following:

Enclosure management device

If a device light or a series of device lights do not flash, it is because all or part of the subsystem is not controlled by a SAF-TE processor.

For example, if SCSI channel 1 has a SAF-TE processor, but SCSI channel 3 does not, and you select to Identify all physical drives, only SCSI channel 1 will flash.

Identify is particularly useful for enclosures. You can use Identify to clarify which drives in the enclosure you need to work with, before you pull them out of the enclosure.

If a SAF-TE processor is present and the controller supports device identification, you can identify the components by doing the following:

Right-click:	To Identify...
Enclosure management device	The enclosure itself
Enclosure	All physical drives in the enclosure
Array [number]	All physical drives in the selected array or spanned array
Logical drive [number]	All physical drives in the selected logical drive
Physical drive [number]	The selected physical drive.
Arrays (Arrays object)	All physical drives in all arrays
Logical drives (Logical drives object)	All physical drives in all logical drives
Channel (Channel object)	All physical drives on the selected channel
Ports (Ports object)	All physical drives on the selected ports

More information

- Identifying devices (action)
- ServeRAID Manager and LED flash states

IBM ServeRAID SCSI controllers: You can use the IBM ServeRAID Manager with the following:

- IBM ServeRAID-7K Ultra320 SCSI controller
- IBM ServeRAID-6i/6i+ Ultra320 SCSI controller
- IBM ServeRAID-6M Ultra320 SCSI controller
- IBM ServeRAID-5i Ultra320 SCSI controller
- IBM ServeRAID-4H Ultra160 SCSI controller
- IBM ServeRAID-4Mx Ultra160 SCSI controller
- IBM ServeRAID-4Lx Ultra160 SCSI controller
- IBM ServeRAID-4M Ultra160 SCSI controller
- IBM ServeRAID-4L Ultra160 SCSI controller
- IBM ServeRAID-3HB Ultra2 SCSI controller
- IBM ServeRAID-3H Ultra2 SCSI controller
- IBM ServeRAID-3L Ultra2 SCSI controller
- IBM ServeRAID-II Ultra SCSI controller
- IBM ServeRAID SCSI controller

More information

- ServeRAID-5 hardware features
- ServeRAID-4 hardware features

- ServeRAID-3 hardware features
- ServeRAID software features

Integrated RAID controller: The integrated RAID controller (such as an LSI 1030) is a standard feature on some IBM xSeries servers. This controller has limited RAID capabilities. With an integrated RAID controller, you can use the ServeRAID Manager to:

- Configure two physical drives into a logical drive and assign it RAID level-1.
- Configure a ready physical drive as a hot-spare drive.
- Monitor the physical drives for PFA errors.
- Set drive status lights for a failed physical drive or PFA error.

The integrated RAID controller is displayed in the ServeRAID Manager Enterprise view as two controller objects. However, you can configure only one array and RAID level-1 logical drive on an integrated RAID controller. The controller information in the tree lists the controller type as LSI 1030.

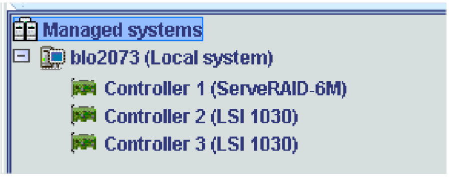

Managed systems
blo2073 (Local system)
Controller 1 (ServeRAID-6M)
Controller 2 (LSI 1030)
Controller 3 (LSI 1030)

Attention: If you choose to install a ServeRAID controller in the server, you cannot migrate the data from an array configured on the integrated RAID controller to an array configured on the ServeRAID controller.

For additional information about the integrated RAID controller, see the installation instructions and CDs provided with your IBM xSeries server.

More information

- Understanding RAID level-1
- Understanding RAID level-1 Enhanced
- Choosing to create a hot-spare or standby hot-spare drive
- ServeRAID software features

Configuring the ServeRAID controller for clustering:

Note: This action is not supported when using the following:

- Integrated RAID controller
- ServeRAID-5i controller
- ServeRAID-6i/6i+ controller
- ServeRAID-7k controller
- ServeRAID-7t controller
- ServeRAID-8i controller
- HostRAID controller

Use this action to configure and view ServeRAID settings for clustering computers or using high-availability.

The ServeRAID Manager configures the controller name, partner controller name, and SCSI initiator IDs in bootable-CD mode only; otherwise, you can view only this information .

You can configure merge-group numbers for a logical drive that you have created during the current session of the ServeRAID Manager. If you have exited from the ServeRAID Manager since you created the logical drive, you cannot configure merge-group numbers for it. You must use the ServeRAID Manager in startable-CD mode to change merge-group numbers for logical drives that are not new.

1. In the Enterprise view, click (controller).
2. Right-click **Clustering actions** → **Configure for clustering**.
3. Enter the controller name and partner controller name. You can specify a name up to 12 characters. When you enter the partner controller name, the merge-group information displays in the window.

   Note: Names are case sensitive.
4. Select the SCSI initiator IDs which can be 6 or 7.

   Note: The SCSI initiator IDs must be unique from the partner controller initiator IDs. The default is 7.
5. If you want a logical drive to merge to another controller if this controller fails, select the **Shared** box. Otherwise, if you do *not* want a logical drive to merge to another controller, select the **Non-shared** box.

   Note:

   a. Non-shared is the default for all logical drives. The merge-group number is defined as [200 + the SCSI initiator ID]. You cannot change the merge-group number for a non-shared logical drive.
   b. If a logical drive is defined as shared, all physical drives in the logical drive must be on a shared channel.
6. Enter a merge-group number for the shared logical drive. This merge-group number must be unique from the partner controller (that is, the merge-group number for any logical drive belonging to the partner controller cannot be the same merge-group number).

   Note: The default merge-group number is the logical-drive number.
7. Once you have configured both controllers for clustering, you can click View shared drives to view the physical drives owned by a partner controller.

More information

- Using ServeRAID in a cluster environment

Validating a cluster:

Note: This action is not supported when using the following:
- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application
- Integrated RAID controller
- ServeRAID-5i controller
- ServeRAID-6i/6i+ controller
- ServeRAID-7t controller
- ServeRAID-8i controller
- HostRAID controller

Note: You must have the ServeRAID Manager installed and running on all servers in the cluster for this feature to work. You must run the Validate cluster feature from one of two nodes in the cluster; you cannot run it remotely.

Use this action to determine if your cluster configuration is configured correctly.

1. In the Enterprise view, click 🖥 (local system only).
2. Right-click **ServeRAID actions** → **Validate cluster**. The ″Validate cluster″ window opens. The system you selected in step 1 is listed in the **Node A** box.

   Note: You cannot change the Node A system in the ″Validate cluster″ window. To select a different system, click **Cancel** to return to the tree.
3. Select your second system from the **Node B** drop-down list.

   Note: The Node B system defaults to the first system in the Enterprise view other than the Node A system. If there are no remote systems in the Enterprise view, click **Cancel** to return to the Enterprise view and add a remote system.
4. Click **Start**. The ServeRAID Manager scans both systems to verify the configuration.
5. If you want to stop the validation before it is completed, click **Cancel**. Otherwise, if the ServeRAID Manager found problems with the configuration, it reports the problems in the message field. If there are no problems with the configuration, the message field reports only the cluster pairings found.
6. To view the problem report events in an easy-to-read window, double-click an event in the message field.
7. To save the contents of the message field, click **Save**. A ″Save as″ window opens.
8. Specify the file name for the message field log and click **OK**.
9. Click **Done** to exit the ″Validate cluster″ window. If the configuration has problems, use ″ Configure for clustering″ to fix the problems; then, use ″Validate cluster″ again.

More information

- Configuring the ServeRAID controller for clustering
- *IBM ServeRAID User's Reference*

Clearing controller event logs: Use this action to clear (erase) all controller event logs. If you are having problems, do not clear the event logs; your service representative might need to reference the logs. If you must clear the logs, save the logs first.

Note: The event logs are automatically cleared when you update the firmware (microcode) on the controller.

1. In the Enterprise view, click 🖥 (system).
2. Right-click **ServeRAID Manager Actions** → **Clear all controller event logs**. The ServeRAID Manager clears all event logs belonging to all controllers in the selected server.

More information

- Saving a printable configuration
- Saving a printable configuration and event logs

Replacing a ServeRAID controller:

Note: This action is not supported when using the following:

- Integrated RAID controller
- ServeRAID-5i controller

- ServeRAID-6i/6i+ controller
- ServeRAID-6M controller under the Windows operating system

Use this action to perform a hot-swap replace of a controller.

1. In the Enterprise view, click (controller).
2. Right-click **Replace controller**. The IBM ServeRAID Hot-Swap Wizard opens.
3. Click **Next**. The wizard turns off the PCI slot.
4. Click **Next**. The ″Replace the controller″ window opens. This window lists the steps that you must follow to replace the controller hardware in the server.
5. When complete, click **OK**. The wizard does the following:
   a. Turns on the PCI slot.
   b. Verifies that the controller is working properly.
   c. Configures the controller by importing the configuration from the physical drives.

      Note: Configuration can take several minutes.

      .
6. Click **Next**.
7. Click **Finish**.

More information

- Steps for hot-replacing a ServeRAID controller
- Using Active PCI features

Configuring a ServeRAID controller: the basic steps: The following steps are an overview of the ServeRAID configuration process:

1. In the Enterprise view, click (controller) that you want to configure.
2. Right-click **Create logical drive**, or click in the Logical devices view. The Configuration wizard opens.
3. Click either **Express configuration** or **Custom configuration**: If you select Express configuration, the ServeRAID Manager automatically:
   a. Creates one or more arrays, based on the number and size of ready physical drives in your system.
   b. Defines a hot-spare drive, if four or more ready physical drives of the same size are available.
   c. Defines a logical drive for each array.

   If you select Custom configuration, the ServeRAID Manager walks you through:
   a. Selecting the physical drives that you want to include in your arrays or create as hot-spare drives.
   b. Defining the logical drives for your arrays.
4. Click **Automatically initialize new logical drives**. Initializing a logical drive erases the first 1024 sectors on the drive and prevents access to any data previously stored on the drive.
5. Click **Next**; then, review the configuration summary.
6. Click **Apply**. The ServeRAID Manager will configure the controller and initialize and synchronize the logical drives.

SR-8i SR-7t HostRAID For HostRAID controllers, Express configuration creates a single logical drive. Custom configuration allows you to select the physical drives and segments used to define your logical drives.

More information

- Example: Express configuration
- Choosing between Express and Custom configuration (SAS, SATA, HostRAID)
- Configuring RAID

Configuring two ServeRAID controllers in a failover environment:

Note: This action is not supported when using the following controllers:

- Integrated RAID controller
- ServeRAID-5i controller
- ServeRAID-6i/6i+ controller
- ServeRAID-7t controller
- ServeRAID-8i controller
- HostRAID controller

You can configure two ServeRAID controllers in a failover environment when using Microsoft Windows 2000 or Microsoft Windows NT 4.0. Recent versions of ServeRAID device drivers utilize fault-tolerant technology. With fault tolerance, you can pair two controllers and connect them to the same enclosure to ensure access to the physical drives, even after one controller fails.

Failover requires specific controller settings. You must set or change these settings by using the ServeRAID Manager in bootable-CD mode. The settings are the following:

- Controller name
- Partner controller name
- SCSI initiator IDs

Use Configure for Clustering to change these settings. You can use this action to view the settings from the installed ServeRAID Manager. To set or change these settings, you must use this action in bootable-CD mode.

Important: Be sure to review *Configuring IBM ServeRAID Controllers for Failover*.

More information

- Failing from active to passive controller (action)

Steps for hot-replacing a ServeRAID controller: If you are hot-replacing a ServeRAID controller in a server that you are monitoring remotely, print these steps so you can refer to them while working at the server.

1. Review the Safety Information in the ServeRAID book and the Safety Information book provided with your IBM server.
2. Detach the SCSI cables from the ServeRAID controller.
3. Open the locking handle and raise the latch.
4. Unseat the controller from the PCI slot and remove the controller from the server.
5. Insert the replacement controller in the PCI slot.
6. Lower the latch and close the locking handle.
7. Connect the SCSI cables to the replacement controller.

More information

- Replacing a ServeRAID controller (action)
- Using Active PCI features

Failing from the active to the passive controller:

Note: This action is not supported when using the following:

1. Integrated RAID controller
2. ServeRAID-5i controller
3. ServeRAID-6i/6i+ controller
4. ServeRAID-7t controller
5. ServeRAID-8i controller
6. HostRAID controller

Use this action to force a failover from the active to the passive controller for the following reasons:

- To test your failover configuration.
- If you are having problems accessing data through ServeRAID, attempt to failover to see if that fixes the problem.

Note:

1. This action is available only if clustering is enabled.
2. Failover detects **only** failed controllers.

1. In the Enterprise view, click (controller).
2. Right-click **Clustering actions** → **Fail from active to passive**.
3. Click **Yes** to confirm the failover.

Networked storage object: Use this object, in the Enterprise view, to work with the management stations and network-attached enclosures in your configuration or to display their properties. When you click this object, the ServeRAID Manager displays a summary of management stations in the **Networked storage** view.

Possible subobject

- (Management station)
- (Display group)

Warning and error conditions

- ⚠ (warning) is displayed if any management station has a warning or error condition.

Actions

- Add management station
- Remove management station
- Update controller images

Management station object: Use this object, in the Enterprise view, to work with a managment station or to display its properties.

Possible subobject

- (Enclosure)

Warning and error conditions

- ⚠ (warning) is displayed if any enclosure has a warning or error condition.

Actions

- Add agent
- Remove agent
- Configure management station
- View management station event log
- Change display group

Adding an agent: Use this action add a management station agent or FlashCopy agent:

- The management station agent is a monitoring agent for network storage. After you add a management station agent, you can monitor and configure the attached enclosures from a ServeRAID Manager console.
- The FlashCopy agent monitors FlashCopy backups for network storage. You must run at least one FlashCopy agent to manage FlashCopy backups of logical devices on external storage enclosures.

Adding a management station agent

To add a management station agent:

1. In the Enterprise view, click 🖥 (management station).
2. Right click **Add agent**. The ″Add agent″ window opens.
3. Enter the management port host name or TCP/IP address of one of the controllers in the enclosure. (Single controller configurations have only one controller in the enclosure.)
4. Enter the administrator’s password in the password field.

   Note: The administrator’s password is established when the enclosure is installed on the network. It is not the same as the management station password.
5. Click **Add**.

Adding a FlashCopy agent

To add a FlashCopy agent:

1. In the Enterprise view, click 🖥 (management station).
2. Right click **Add agent**. The ″Add agent″ window opens.
3. Enter the host name or TCP/IP address of the machine on which the agent is running.

   Note: For the FlashCopy agent, no password is required.
4. Click **Add**.

More information

- Using the management station agent
- Removing an agent

Removing an agent: Use this action to remove a management station agent or FlashCopy agent for networked storage. Removing the management station agent deletes the enclosure from the Enterprise view, but does not take the enclosure offline.

To remove an agent:

1. In the Enterprise view, click (management station).
2. Right click **Remove agent** and then the agent name. In the pull-right menu, management station agents are listed first, above the horizontal rule. FlashCopy agents are listed below the rule. The rule is visible (near the top of the panel) even if no management station agents are defined.
3. Click **Yes**.

More information

- Using the management station agent
- Adding an agent

Configuring a management station: Use this action to configure a management station. After you configure a management station, you can monitor the attached enclosures from a ServeRAID Manager console running on another system.

1. In the Enterprise view, click (management station).
2. Right click **Management station actions** → **Configure**. The ″Management Station Configuration″ window opens.
3. Configure the Security Manager, SNMP Traps Manager, and Email Notification Manager.

More information

- Using the management station agent
- Using the Security Manager
- Using the SNMP Traps Manager
- Using the Email Notification Manager

Viewing the management station event log: Use this action to view events in the management station event log.

1. In the Enterprise view, click (management station).
2. From the Actions menu, click **Management station actions** → **View event log**. The ″Management station event viewer″ window opens.
3. If you want to save the event log to a file, click **File** → **Save as**. The default is Events.txt.
4. Click **File** → **Close** to close event viewer window.

More information

- Using the management station agent

Using the management station agent: The management station agent is a monitoring agent for network storage. After you add an agent, you can configure the management station to:

- Notify users by email when events occur on the enclosures attached to the management station. Use the Email Notification Manager, from the Management Station Configuration console, to add recipients to the email notification list.
- Prevent unauthorized users from connecting to the management station using its built-in security feature. The management station checks the security list user

name and password to ensure that only authorized users can log in. Use the Security Manager, from the Management Station Configuration console, to add users to the agent security list.

- Send SNMP traps. You can receive SNMP traps using any available SNMP management program. Use the SNMP Trap Manager, from the Management Station Configuration console, to add systems to the SNMP traps destination list.

The agent monitors and generates events for critical or fatal problems in the enclosure configuration every 5 seconds. These changes include, but are not limited to:

- defunct drives
- PFA drives
- Failed battery
- Offline or critical logical drives
- Failed controllers
- Enclosure problems
- Non-warranted drives. An event is sent at startup, console connection, and every 30 days

More information

- Adding a management station agent
- Configuring a management station
- Viewing the management station event log

Enclosure object: Use this object to work with an enclosure or to display its properties.

Possible subobject

None.

Warning and error conditions

- (warning) is displayed if the enclosure has the following conditions:
  - A fan has failed.
  - A power supply has failed.
- (error) is displayed if the enclosure has the following conditions:
  - It is not responding.
  - It is overheating.
  - Multiple fans have failed.
  - Multiple power supplies have failed.
- (enclosure) is displayed if the enclosure is one of the following:
  - Working properly.
  - Not supported.
  - Supported, but not monitored for problems. For example, enclosures that are internal to a server do not have separate fans or power supplies.

Actions

- Identify all physical drives
- Fail back storage (dual controller configurations only)

- Shut down enclosure
- Restart enclosure
- Enter software key
- Change controller date and time
- Save support archive

Save support archive: Use this action to save the ServeRAID Manager configuration and status information in an archive file. Your ServeRAID Manager service representative might ask you to create this file to help diagnose a problem with your system.

1. In the Enterprise view, click (system) or (enclosure).
2. From Action menu, select **Save support archive**. The "Save as" window opens.
3. In the File name field, type a name for the archive file or accept the default name.
4. Click **Save**.

More information

- View event log
- Save printable configuration
- Save printable configuration and event logs

Arrays object: Use this object, in the Logical devices view, to work with basic arrays or spanned arrays or to display their properties:

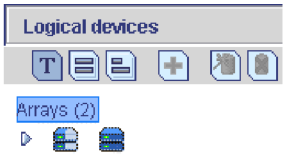

Note: A controller supports a maximum of eight arrays or four spanned arrays. An integrated RAID controller supports a maximum of one basic array.

Possible subobject

- (Array or spanned array)

Warning and error conditions

- None.

Actions

- Delete all arrays
- Identify all arrays
- If the array has free space, Create logical drive in array

Understanding creating basic arrays: When you group one or more physical drives together and configure the ServeRAID controller to access them in a particular pattern, you create a basic array. The maximum number of basic arrays depends on the controller and RAID level.

Arrays are used to improve security, performance, and reliability. The amount of improvement depends on the application programs that you run on the server, ServeRAID parameters, and the RAID levels that you assign to the logical drives in your arrays. The ServeRAID software supports RAID levels-0, 1, 1E, 5, 5E, and 6. (Not all RAID levels are supported by all controllers.)

Note: Each ServeRAID controller supports up to eight logical drives. If any of the existing arrays contain a RAID level-5 Enhanced logical drive, you can have only seven logical drives on a controller.

SR-8i SR-7t For the ServeRAID-8i controller you can configure a maximum of 128 logical drives. For the ServeRAID-7t controller, you can configure a maximum of 24 logical drives.

More information

- Understanding creating spanned arrays
- Creating arrays and hot-spare drives in the wizard
- Configuring RAID and creating arrays (action)
- Example: Express configuration
- Physical drive capacity and unusable capacity

Understanding creating spanned arrays:

Note: This feature is not supported on the integrated RAID controller.

When you group one or more arrays together and configure the ServeRAID controller to access them in a particular pattern, you create a spanned array.

Spanned arrays are used to create logical drives using as many as 60 physical drives. The ServeRAID controller supports RAID level-00, 10, 1E0, and 50.

More information

- Creating spanned arrays in the wizard
- Understanding RAID level-x0
- Understanding when RAID level-x0 logical drives change state
- Understanding creating basic arrays
- Physical drive capacity and unusable capacity

Configuring RAID and creating arrays:

Note: This action is supported on the integrated RAID controller in bootable-CD mode only.

You can use the Configuration wizard to create up to eight arrays and up to eight logical drives for each ServeRAID controller. The Configuration wizard provides two configuration options: Express and Custom. Express configuration automatically configures your ServeRAID controller. Custom configuration allows you to configure your controller manually. If you want to use RAID level-1E, RAID level-5E, RAID level-5EE, or RAID level-x0, you must use Custom configuration.

Note: If you intend to use your ServeRAID controllers in a Microsoft Windows failover or clustering environment, review the ServeRAID publications before configuring ServeRAID.

1. In the Enterprise view, click (controller).

Note: The integrated RAID controller is displayed in the ServeRAID Manager Enterprise view as two controller objects. However, you can configure only one array and RAID level-1 logical drive on an integrated RAID controller.

2. Right-click **Create logical drive**, or click on the toolbar or Logical devices view. The Configuration wizard opens.
3. Click either **Express configuration** or **Custom configuration**.

   Note: If you click **Express configuration**, you will have the opportunity to review and approve the configuration before the ServeRAID Manager applies the configuration.
4. Click **Automatically initialize new logical drives**. Initializing a logical drive erases the first 1024 sectors on the drive and prevents access to any data previously stored on the drive.
5. Create the arrays and hot-spare drives.
6. If you want spanned arrays, create the spanned arrays.
7. Create the logical drives.
8. Review the configuration summary and apply the configuration.

More information

- Configuring a ServeRAID controller: the basic steps
- Initializing a logical drive

Creating arrays and a hot-spare drive for your integrated RAID controller subsystem: Complete the following steps to create arrays:

1. Click the **Array** tab in the right pane:

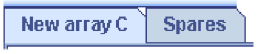

2. Then, from the list of ready drives, select the drives you want to move to the array:

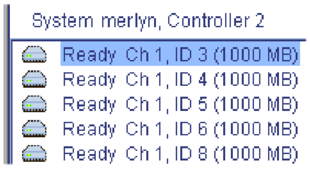
System merlyn, Controller 2

Ready Ch 1, ID 3 (1000 MB)
Ready Ch 1, ID 4 (1000 MB)
Ready Ch 1, ID 5 (1000 MB)
Ready Ch 1, ID 6 (1000 MB)
Ready Ch 1, ID 8 (1000 MB)

3. Click >> (Add selected drives) to add the drives to the array.
4. If you want to configure a hot-spare drive, complete the following steps:
   a. Click the **Spares** tab.
   b. Select the physical drive you want to designate as the hot-spare drive; then, click >> (Add selected drives).
5. Click **Next**. The ″ Configuration summary″ window opens.

To return to the ″ Express and Custom configuration″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Understanding creating basic arrays

Deleting all arrays:

Note:

1. This action is supported on the integrated RAID controller in bootable-CD mode only.
2. This action is not supported on the ServeRAID-8i and ServeRAID-7t controllers.

Use this action to delete all arrays on the specified controller.

For some operating systems, deleting an array and its associated logical drives might change the drive-letter assignments of the remaining drives when you restart the system.

Attention: If you delete all arrays associated with the controller, all data on the controller is destroyed.

1. In the Enterprise view, click (controller).

   Attention: Before proceeding, ensure that you have backed up any data and programs that you want to save.
2. Right-click **Delete all arrays**.
3. Click **Yes** to confirm the delete.

Replacing physical drives in an array: If you have an array built with one or more physical drives that are not the same size, or you simply want to upgrade the array with new, larger drives, you can replace the physical drives in the array.

Attention: During the rebuild operation, if another physical drive fails, there is the risk of losing data.

1. Replace the physical drive in the server or enclosure with a larger physical drive.
2. Wait for the rebuild to complete.
3. Repeat steps 1 and 2 until you have replaced all the smaller physical drives in the array.

More information

- Example: Usable and unusable capacity

Modifying an array in the migration wizard: To add or remove drives in an array:

1. In the Physical devices view (on the right), click the drives you want the array to use. To view available segments on each disk drive, switch to the full-size capacity view or relative-size capacity view:

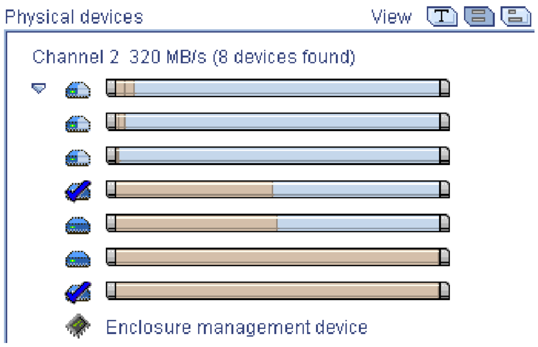
Physical devices
View
Channel 2 320 MB/s (8 devices found)
Enclosure management device

Note: A segment shaded in light blue is not part of any array or logical drive.

2. To remove or replace a physical drive, click the drive you want to remove (indicated by ✖). Then, click the drive you want to replace it with. To cancel your changes and start over, click Revert array .

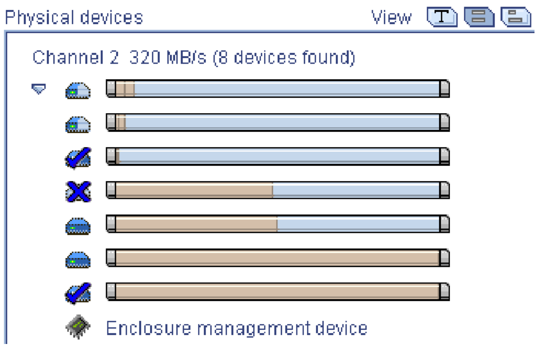
Physical devices
View
Channel 2 320 MB/s (8 devices found)
Enclosure management device

3. When you are ready to continue, click **Next**. The " Configuration summary" window opens.

More information

- Understanding physical drive capacity and unusable capacity

Choosing the primary physical drive of the array: When you create a RAID level-1 logical drive on an integrated RAID controller, you must specify the primary physical drive. Doing so ensures that you retain the data on the primary physical drive. In a RAID level-1 logical drive, the data on the primary physical drive is copied to the secondary physical drive (the mirror). Any data on the secondary physical drive is destroyed.

After the logical drive is created, the physical drive Properties panel reports the mirror role (primary or secondary) of each drive:

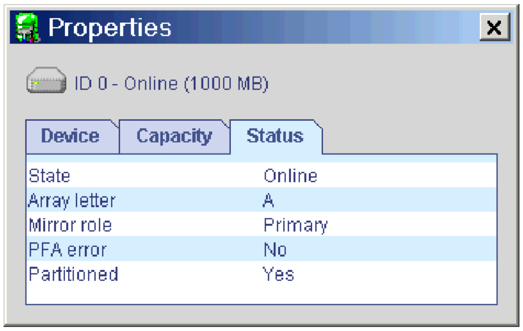

ID 0 - Online (1000 MB)

State	Online
Array letter	A
Mirror role	Primary
PFA error	No
Partitioned	Yes

Knowing the mirror role of each physical drive can be useful in the following situations:

- Creating a logical drive with a physical drive that already contains data. A server has one physical drive and the operating system is installed on that physical drive. You decide that you want a RAID level-1 logical drive to achieve redundancy. When you install a second physical drive and configure the logical drive, you must specify the original physical drive as the primary physical drive so you do not destroy the data on that drive.
- During a synchronization failure. After the logical drive is created, a synchronization begins to create the mirror of the primary physical drive. While the synchronization is in progress, the logical drive is displayed as critical, because the mirror is not complete and the logical drive has not achieved redundancy. During the synchronization, if the primary physical drive fails, all of the data on the drive is lost because the mirror is incomplete. If the secondary physical drive fails, the data on the primary physical drive is intact and you can create the mirror using a hot-spare drive.

More information

- Understanding RAID level-1
- Understanding logical-drive synchronization

Confirming your array migration configuration: Use the configuration summary to review all the changes that the ServeRAID Manager will apply to your modified array and logical drive configuration.

1. Review the information displayed in the ″Configuration summary″ window. It describes how the new configuration will affect the RAID level and free space. To change the configuration, click **Back**.
2. Click **Apply**.
3. Click **Yes** when asked if you want to apply the new configuration.

   Note: Logical-drive migration is a lengthy process. The ServeRAID Manager displays (in animation) in the Logical devices view while the operation is in progress. When the migration is complete, the configuration is saved in the controller and in the physical drives.

More information

- Understanding logical-drive migration

Exporting an array: Use this action to export the specified array. Exporting an array prepares the RAID configuration for transfer to another enclosure.

Note: When you export an array, its status changes to foreign. You cannot perform any action on a foreign array except view array components and import a RAID configuration.

To export an array:

1. In the Logical devices view, click (array).
2. Right-click **Export array**.
3. Click **Yes** to confirm the export.

More information

- Managing Foreign arrays and Alien arrays
- Importing an array

- Deleting an array

Deleting an array:

Note: This action is supported on the integrated RAID controller in bootable-CD mode only.

Use this action to delete a specified array.

For some operating systems, deleting an array and its associated logical drives might change the drive-letter assignments of the remaining drives when you restart the system.

Attention: If you delete an array, you delete all logical drives that are defined in that array. All data and programs on the logical drives in the array are lost during this procedure.

To delete an array:

1. In the Logical devices view, click (array).

   Attention: Before proceeding, ensure that you have backed up any data and programs that you want to save.

2. Right-click **Delete array**.
3. Click **Yes** to confirm the delete.

Importing an array: Use this action to import the RAID configuration into the specified foreign array. Importing the RAID configuration completes the transfer of an array from another enclosure.

To import an array:

1. In the Logical devices view, click (foreign array).
2. Right-click **Import array**.
3. Click **Yes** to confirm the import.

More information

- Managing Foreign arrays and Alien arrays
- Exporting an array
- Deleting an array

Moving an array to a different controller:

Note: This action is not supported on enclosures with a single controller configuration.

Use this action to move an array to a different controller in the enclosure. This action changes the preferred owner of the array and moves all associated logical drives to the new controller.

1. In the Logical devices view, click (array).
2. Right-click **Move to controller** and then the controller name (A or B).
3. Click **Yes** when prompted to confirm the move.

More information

- Recovering storage with fail back

Array [letter] object: Use this object, in the Logical devices view, to work with an array or to display its properties.

Notes:

1. Basic arrays are labeled with letters starting with A. Spanned arrays are labeled with numbers starting with 1. Spanned arrays are used in RAID level-x0 configurations only.
2. To view the array labels, switch to Text view; then, expand the array list.
3. A ServeRAID SCSI controller supports a maximum of eight basic arrays or four spanned arrays. An integrated RAID controller supports a maximum of one basic array.

Possible subobject

- Logical drives (Logical drives)
- (Logical drive)

Warning and error conditions

- ⚠ (warning) is displayed if the logical drives object has a warning condition.
- ✖ (error) is displayed if the logical drives object has an error condition.

Actions

- Identify array
- If the selected array is **not** a subarray in a RAID level-x0 configuration, the following actions are available:
  - Delete array
  - Change RAID level from...
    - RAID 0 to RAID 5
    - RAID 1 to RAID 5
    - RAID 5 to RAID 0
    - RAID 5E to RAID 5
  - Increase free space
  - Increase logical drive size
- If the array is on a network-attached storage enclosure:
-
  - Export array (if the array is online)
  - Import array (if the array is foreign)
  - Move array to controller (dual controller configurations only)
  - Expand or migrate array

Logical drives object: Use this object, in the Logical devices view, to work with logical drives or to display their properties:

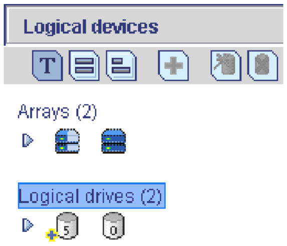

Notes:

1. ServeRAID SCSI controllers support a maximum of eight logical drives.
2. The ServeRAID-7t SATA controller supports a maximum of 24 logical drives.
3. The ServeRAID-8i SAS controller supports a maximum of 128 logical drives.
4. An integrated RAID controller supports a maximum of one logical drive.

Possible subobject

- (Logical drive)

Warning and error conditions

- None.

Actions

- Delete > logical drive number
- Identify all logical drives
- Configure access control list (networked storage only)
- Configure target information (networked storage only)

Understanding when RAID level-x0 logical drives change state: Start with four physical drives.

Use two physical drives to create array A and two physical drives to create array B.

Use arrays A and B to create spanned array 1. As part of a spanned array, arrays A and B are now called ″subarrays.″

A sub-logical drive is created in array A and another sub-logical drive is created in array B.

Then, create a logical drive in spanned array 1.

If the sub-logical drive in subarray A is in the okay state and the sub-logical drive in subarray B is in the okay state, then the logical drive in spanned array 1 is in the okay state.

If a physical drive fails in subarray A, then the sub-logical drive in subarray A changes to the critical state. Subarray B is still in the okay state. The logical drive in spanned array 1 changes to the critical state.

If the second physical drive fails in subarray A, then the sub-logical drive in subarray A changes to the offline state. Subarray B remains in the okay state. The logical drive in spanned array 1 changes to the offline state.

If the sub-logical drive in subarray A is in the critical state and a physical drive fails in subarray B, then the sub-logical drive in subarray B changes to the critical state. With both sub-logical drives in the critical state, the logical drive in spanned array 1 remains in the critical state.

More information

- Understanding creating spanned arrays
- Creating spanned arrays in the wizard
- Understanding RAID level-x0
- Physical drive capacity and unusable capacity
- Understanding creating basic arrays

Understanding logical-drive synchronization: The purpose of synchronizing logical drives is to compute and write the parity data on the selected drives. Synchronizing a logical drive verifies that the data redundancy for the logical drive is correct.

Note: You can determine if your controller firmware supports auto-synchronization and data scrubbing by viewing the controller properties. The properties pane reports if these features are enabled or disabled.

Using auto-synchronization

If the logical drive is attached to a controller that supports auto-synchronization, the ServeRAID Manager automatically synchronizes all new RAID level-5, 5E, 5EE, and 50 logical drives. These logical drives must be synchronized before storing data. Auto-synchronization ensures that the parity data is accurately computed. Accurate parity is critical for updating the parity when you begin to store data.

SR-8i SR-7t For the ServeRAID-8i controller and ServeRAID-7t controller, the ServeRAID Manager also supports auto-synchronization for RAID level-1 and 10 logical drives. During auto-synchronization, data from the primary drive is automatically copied to the mirror drive.

Depending on the controller, the ServeRAID Manager displays either a progress indicator in the status bar or a (in animation) in the Logical devices view, indicating that auto-synchronization is in progress. You can perform some actions on the controller during an auto-synchronization, such as deleting a logical drive; however, you **cannot** perform a logical drive migration. Also, you can turn off the server during an auto-synchronization and, when you turn on the server, the firmware will continue the auto-synchronization where it left off.

Using manual synchronization

When you create a logical drive that is attached to a controller that does **not** support auto-synchronization, you will receive a warning event in the event viewer stating that you must **manually** synchronize the logical drive before storing data. Use ″ Synchronize logical drives″ to manually synchronize the logical drive. You must synchronize **new** RAID level-5 logical drives after you create them.

Manual synchronization can be a lengthy operation. You **cannot** perform any other actions on the controller until the manual synchronization is completed. You can synchronize the logical drive immediately or schedule it for a later date and time. You can also schedule the synchronization as a recurring task. A recurring synchronization task starts at the initial start time, then runs periodically according to the given interval: the same time each day, each week, or each month. Do **not** turn off the server until the manual synchronization is complete.

Using manual synchronization for data scrubbing

Data scrubbing is an automatic background synchronization process. Data scrubbing keeps data ″fresh″ by doing the following:

- (For RAID level-5, 5E, 5EE, or 50) Reading data and rewriting the data parity.
- (For RAID level-1, 1E, 10, 1E0) Reading data and rewriting the mirror data.

If data scrubbing is disabled or your controller firmware does not support data scrubbing, consider manually synchronizing your logical drives weekly. This does not alter data on the drive.

Note: If data scrubbing is disabled on your controller, you **can** enable data scrubbing, but not through the ServeRAID Manager. Instead, use the IPSSEND command-line program. This program comes with the device option and is on the *IBM ServeRAID Support* CD. For more information, refer to the *IBM ServeRAID User′s Reference*.

More information

- Synchronizing logical drives (action)
- Understanding scheduled tasks
- Using the Task Manager

Creating logical drives: Complete the following steps to create logical drives:

1. Start the custom configuration wizard and complete all previous steps.
2. Click the appropriate **Array** tab.

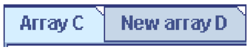

3. Select a RAID level from the drop-down list.

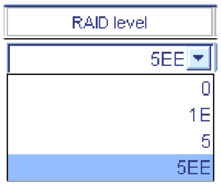

Note:

a. RAID level-5E and RAID level-5EE allow only one logical drive per array.

b. If you are configuring a spanned array, you can set the RAID level only for the first logical drive you create.

c. If you plan to use ″ Change RAID level,″ you must assign the same RAID level to all logical drives within a single array.

Attention: Before assigning a logical drive RAID level-5E or RAID level-5EE, consider the following. If a physical drive fails during a post-failover resynchronization, the logical drive will enter the blocked state. Data might be lost or damaged.

4. If you do not want to use the maximum size for the logical drive, type the size in the **Data (MB)** field.

   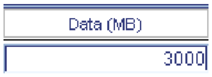

   Data (MB)
   3000

   Note:

   a. You can define up to eight logical drives per controller. There are two exceptions:
      - If an array contains a logical drive assigned RAID level-5E
      - If you want to use the logical-drive migration feature

      In these cases, one logical drive slot must be left free; therefore, you must define no more than seven logical drives.
   b. Some operating systems have size limitations for logical drives. Before you save the configuration, verify that the size of the logical drive is appropriate for your operating system. For more detailed information, see your operating-system documentation.
   c. A logical drive cannot exceed 2048 GB (2 terabytes); the minimum size is *n*MB, where n equals the number of drives in the array.
   d. Typically, the first logical drive defined on the first ServeRAID controller found by system BIOS during startup will be your startup (boot) drive.
   e. The actual logical-drive size might be slightly different from what you select. The RAID level and the number of physical drives determine the size of the logical drive. For example, an array consisting of three, 1 GB physical drives with a requested RAID level-0 logical drive of 1000 MB will actually contain only 999 MB because the data is striped across all three drives, with 333 MB on each drive.
5. If you have free space available and want to create additional logical drives, click Define new logical drive .
6. Repeat steps 3 through 5 for each logical drive that you want to define in this array.
7. Repeat steps 2 through 6 for each additional array that you want to configure.
8. Click **Next**. The ″ Configuration summary″ window opens.

To return to the ″ Create arrays″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Selecting the RAID level by array capacity
- Selecting the array size
- Understanding physical drive capacity and unusable capacity

Creating arrays and hot-spare drives: Complete the following steps to create arrays:

1. Start the custom configuration wizard and complete all previous steps.

2. Click the appropriate **Array** tab in the right pane:

3. Then, from the list of ready drives, select the drives you want to add to the array:

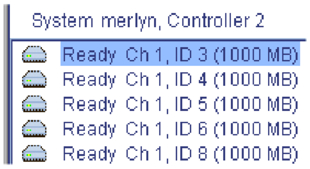

System merlyn, Controller 2

Ready Ch 1, ID 3 (1000 MB)
Ready Ch 1, ID 4 (1000 MB)
Ready Ch 1, ID 5 (1000 MB)
Ready Ch 1, ID 6 (1000 MB)
Ready Ch 1, ID 8 (1000 MB)

4. Click >> (Add selected drives) to add the drives to the array. You can click >> (Add all drives) to move **all** ready drives to an array.
5. Repeat steps 2 and 3 for each additional array or hot-spare drive that you want to configure.
6. If you do *not* want to create a spanned array, skip to step 6. Otherwise, select the **Span arrays** check box ☑ Span arrays . Then, click **Next**. If you created two arrays only, the ServeRAID Manager uses those arrays to create the spanned array; continue to step 6. Otherwise, the ″ Create spanned arrays″ window opens.

   Note: To create a spanned array, each array must have the same number of physical drives.
7. After you select the ready drives for your arrays and hot-spare drive, click **Next**. The ″ Create logical drives″ window opens.

To return to the ″ Express and Custom configuration″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Understanding creating basic arrays
- Example: Express configuration
- Creating spanned arrays
- Understanding creating spanned arrays
- Configuring the ServeRAID controller: the basic steps

Creating spanned arrays: If you want to assign RAID level-x0 to an array, you must create a spanned array.

Note: Spanned arrays are supported only by IBM ServeRAID-4 Ultra160, ServeRAID-5i Ultra320, ServeRAID-6M Ultra320, and ServeRAID-6i/6i+ Ultra320 SCSI controllers.

Complete the following steps to create one or more identical spanned arrays:

1. Start the custom configuration wizard and complete all previous steps.
2. In the list of arrays, click the arrays that you want to add to your spanned array.

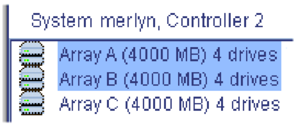

System merlyn, Controller 2

Array A (4000 MB) 4 drives
Array B (4000 MB) 4 drives
Array C (4000 MB) 4 drives

3. Click >> (Add selected arrays) to add the arrays to the spanned array. You can click >> (Add all arrays) to move **all** arrays to the spanned array.

4. To create additional spanned arrays, click the **New spanned array** tab in the right pane.

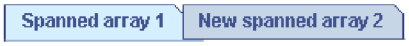

   Then, repeat steps 2 and 3.
5. Click **Next**; the ″ Create logical drives″ window opens.

To return to the ″ Create arrays″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Understanding RAID level-x0
- Understanding creating spanned arrays

Creating additional logical drives:

Note: This action is not supported on the integrated RAID controller.

Use this action to create additional logical drives in an existing configuration. If you configured your ServeRAID subsystem but left free space in an array, you can use this action to create additional logical drives using the free space.

1. In the Enterprise view, click (controller) that contains the array in which you want to create the logical drives.
2. In the Logical devices view, click Arrays (Arrays object).
3. Right-click **Create logical drive**. If the controller includes more than one array, right click **Create logical drive in array** and then the array name. The Configuration wizard opens with the ″Create logical drives″ window.
4. Create the logical drive.
5. Review the configuration summary and apply the configuration.

Changing the logical drive name: Use this action to change the name of a logical drive:

1. Press the Delete key to clear the drive's current name.
2. In the Type new name field, type the drive's new name.
3. Click **OK**.

Clearing a logical drive: Use this action to initialize a logical drive using the Clear initialization method. Initialization is usually automatic when you create a logical drive.

Attention: The Clear initialization method erases the entire logical drive and prevents access to any data previously stored on the drive.

1. In the Logical devices view, click (logical drive).
2. Right-click **Clear**.
3. Click **Yes**.

Deleting a logical drive:

Note:

1. This action is supported on the integrated RAID controller in bootable-CD mode only.

2. SR-SCSI For ServeRAID SCSI controllers, this action is available only if the selected logical drive is the last logical drive that was created in the array.

3. HostRAID To delete a partitioned logical drive on a HostRAID controller, you must remove the partition with your operating system disk management tools first. Then, you can delete the logical drive with the ServeRAID Manager.

4. Enclosure To delete a logical drive on a networked storage enclosure, all iSCSI initiators must be disconnected first. Then, you can delete the logical drive.

Attention: If you delete a logical drive, all data on that logical drive is destroyed. Before proceeding, ensure that you have backed up any data and programs that you want to save.

1. In the Logical devices view, click (logical drive).
2. Right-click **Delete**.
3. Click **Yes** to confirm the delete.

Initializing a logical drive: Use this action to initialize a logical drive. Initialization is usually automatic when you create a logical drive.

Attention: Initializing a logical drive erases the first 1024 sectors on the drive and prevents access to any data previously stored on the drive.

1. In the Logical devices view, click (logical drive).
2. Right-click **Initialize**.
3. Click **Yes**.

Synchronizing logical drives:

Note: This action is not supported on the integrated RAID controller.

If auto-synchronization is not supported on your controller, you must use this action to manually synchronize **new** RAID level-5, 5E, and 50 logical drives after you create them.

If data scrubbing is disabled or your controller firmware does not support data scrubbing, consider manually synchronizing your logical drives weekly using this action.

Manual synchronization can be a lengthy operation. You can synchronize the logical drive immediately or schedule it for a later date and time. You can also schedule the synchronization as a recurring task. A recurring task starts at the initial start time, then runs periodically according to the given interval: the same time each day, each week, or each month.

Note: You cannot perform any other actions on the controller until the manual synchronization is completed. For example, you cannot schedule a synchronization if another logical drive synchronization is in progress.

An example

You schedule a synchronization task with an initial start time of Sunday, October 13, 2004 at 2:00 AM. It recurs weekly. It will run again on Sunday October 20, 2004 at 2:00 AM.

Note: Do not turn off the server while manual synchronization is in progress.

To synchronize a logical drive:

1. In the Logical devices view, click (logical drive).
2. Right-click **Synchronize**.
3. Click **Yes** to synchronize the logical drive immediately (now). Depending on the controller, the ServeRAID Manager displays either a progress indicator in the status bar or (in animation) in the Logical devices view, indicating that the operation is in progress.
4. Or click **Schedule** to schedule the synchronization for a later date and time. The ServeRAID Manager displays the Synchronization Scheduler pane:
   - From the drop-down lists and calendar display, select the day, month, year, and time you want the task performed.
   - From the Recurring drop-down list, select the interval at which you want the task to recur: Never, Daily, Weekly, Monthly.
   - Click **OK** to schedule the task; click **Cancel** to cancel the scheduled task and return to the initial confirmation pane.

More information

- Understanding logical-drive synchronization
- Understanding scheduled tasks
- Using the Task Manager

Unblocking a logical drive:

Note: This action is not supported on the integrated RAID controller.

After the rebuild operation is completed on an array, if the array contained RAID level-0 logical drives, you can unblock the RAID level-0 logical drives and access them once again. But remember, the logical drive contains damaged data. You must re-create, install, or restore the data from the most recent backup disk or tape.

To unblock a blocked logical drive:

1. In the Logical devices view, click (logical drive).
2. Right-click **Unblock logical drive**.
3. Click **Yes** to confirm, understanding that the data on the logical drive is not valid. The ServeRAID Manager unblocks the blocked logical drive.
4. Initialize all unblocked logical drives before using them.
5. Restore the data to the drive from your most recent backup disk or tape.

More information

- Recovering from a defunct drive

Selecting the logical drive size: When you select the logical drive size, consider the following:

- If you change the RAID level of a logical drive, the total size (data plus parity) of the logical drive might change.
- If you do not use all available free space, you can create another logical drive. Click **Create new logical drive** to do so. You can assign the same or a different RAID level to additional logical drives.
- You cannot delete the last remaining logical drive in the array. You must have at least one logical drive for each array.
- Refer to the documentation provided with your operating system for information about the recommended logical-drive size.
- The default RAID configuration uses all available free space.

More information

- Selecting the RAID level by array capacity
- Creating logical drives (action)
- Configuring RAID and creating arrays (action)
- Creating logical drives in the wizard
- Example: Usable and unusable capacity

Scheduling the logical-drive migration: Use the calendar display to choose the date and time for the logical drive migration:

- From the calendar display, select the day, month, year, and time that you want the task performed.
- Click **Apply**.
- Click **Yes** when asked if you want to schedule the migration.

More information

- Understanding logical-drive migration
- Confirming your logical drive migration configuration
- Understanding scheduled tasks
- Using the Task Manager

Logical drive [number] object: Use this object, in the Logical devices view, to work with a logical drive or to display its properties.

Notes:

1. ServeRAID SCSI controllers support a maximum of eight logical drives.
2. The ServeRAID-7t SATA controller supports a maximum of 24 logical drives.
3. The ServeRAID-8i SAS controller supports a maximum of 128 logical drives.
4. HostRAID controllers support a maximum of 2 logical drives.
5. An integrated RAID controller supports a maximum of one logical drive.

Possible subobject

- None.

Warning and error conditions

- (warning) is displayed if the logical drive is in a critical, critical migrating, or critical system state. It is also displayed if a physical drive in the logical drive is a defunct drive or a defunct hot-spare drive.
- (error) is displayed if the logical drive is in a blocked or offline state.
- (initialize) is displayed when the logical drive is initialized (ServeRAID-8i, ServeRAID-7t, HostRAID controllers only).
- (build/modify) is displayed during a logical-drive migration (ServeRAID-8i, ServeRAID-7t, HostRAID controllers only).

Actions

Note: Not all actions are supported on all controllers.

- Delete logical drive
- Initialize logical drive
- Clear logical drive
- Synchronize logical drive
- Enable read-ahead cache
- Expand or change logical drive
- If a logical drive is blocked, unblock logical drive
- Identify logical drive [number]
- Create FlashCopy
- Remove FlashCopy

The following actions are supported for networked storage only (enclosures):

- Configure access control list
- Configure target information
- Increase logical drive capacity

The following action is supported in bootable-CD mode only, **except** for ServeRAID-8i and ServeRAID-7t controllers:

- Change write-cache mode to write back or write through

Channel [number] object: Use this object, in the Physical devices view, to work with a selected SCSI channel or to display its properties:

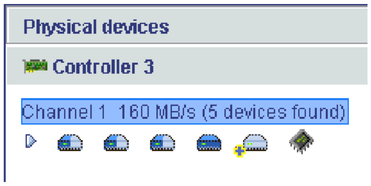

Physical devices

Controller 3

Channel 1 160 MB/s (5 devices found)

Notes:

1. A controller supports a maximum of four SCSI channels.
2. The ServeRAID-5i controller does not support tape drives and physical drives on the same channel.

Possible subobject

- (Physical drive)
- (Hot-spare drive)
- (Enclosure management device)
- (CD-ROM drive)
- (Tape drive)
- (Removable drive)

Warning and error conditions

None.

Actions

- Scan for new or removed ready drives
- Identify all physical drives on SCSI channel [number]

The following action is available only in bootable-CD mode:

- Change the SCSI transfer speed

Scanning for new or removed ready drives:

Note: This action is not supported on the integrated RAID controller.

Use this action to scan all SCSI channels or ports for either of the following:

- Physical drives that have been added.
- Ready physical drives that have been removed.

After physically removing a ready drive from a server or enclosure, the ServeRAID Manager still reports the drive as ready. The firmware (microcode) does not communicate with ready drives, therefore the firmware does not report the drive as removed. Use this action to remove the ready drive from the configuration.

1. In the Enterprise view, click (controller) or (enclosure).
2. Click (Scan for new or removed ready drives).
3. For ServeRAID SCSI controllers, the "Scan for new or removed ready drives on controller" window opens. The ServeRAID Manager scans all SCSI channels on the selected controller. Review the scan report; then, click **Done**. To stop the scan before it completes click **Cancel**.
4. For the ServeRAID-8i, ServeRAID-7t, and HostRAID controllers, the ServeRAID Manager scans all ports on the selected controller. It reports the results of the scan in the event viewer.

Changing the SCSI-transfer speed:

Note:

1. This action is supported in bootable-CD mode only.
2. This action is not supported on the integrated RAID controller.

3. This action is not supported on the ServeRAID-8i controller and ServeRAID-7t controller.

Use this action to change the SCSI-transfer speed.

1. In the Physical devices view, click the SCSI channel identifier:

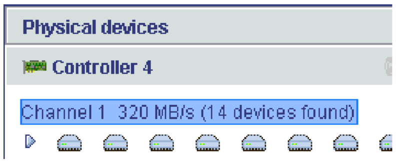

Physical devices

Controller 4

Channel 1 320 MB/s (14 devices found)

2. Right-click **Change the SCSI transfer speed** and then a transfer speed. The following choices are available:
   - Optimal

   This value is available if your controller supports it; otherwise, Optimal does not appear in the menu. Optimal is the default value. When the transfer speed is set to Optimal, the ServeRAID controller determines the best transfer speed, based on the types of SCSI drives and enclosures in use.
   - Ultra320 SCSI
   - Ultra160 SCSI
   - Ultra2 SCSI
   - UltraSCSI
   - Fast SCSI-2
   - SCSI-2
   - SCSI-1

   This value is available if your controller is set to SCSI-1; otherwise, SCSI-1 does not appear in the menu.

Ports [numbers] object: Use this object, in the Physical devices view, to work with selected ports or to display their properties:

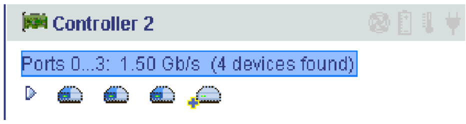

Controller 2

Ports 0...3: 1.50 Gb/s (4 devices found)

Possible subobject

- (Physical drive)
- (Hot-spare drive)

Warning and error conditions

- None.

Actions

- Scan for new or removed ready drives

Physical drive [number] object: Use this object, in the Physical devices view, to work with a physical drive or to display its properties.

Note: A SCSI channel supports a maximum 15 physical drives.

Possible subobject

None.

Warning and error conditions

- ⚠ (warning) is displayed if the selected physical drive is undergoing a rebuild operation, reports a PFA, or is a non-warranted physical drive.
- ✖ (error) is displayed if the selected physical drive is a defunct drive.

Actions

- Identify physical drive
- If the drive is ready:
  - Create hot-spare drive
  - Create standby hot-spare drive
  - Create assigned hot-spare drive
  - Clear
  - Verify
  - Initialize
- Rebuild drive
- If the drive is defunct, Remove the defunct drive
- Set drive state to online
- Change write-cache mode to write back or write through
- If the drive is rebuilding, online, or reports a PFA, set drive state to defunct
- Verify physical drive

Understanding physical drive capacity and unusable capacity: Physical drive capacities influence the way you create arrays. Drives in the array can be of different capacities (1 GB or 2 GB, for example), but the ServeRAID controller treats them as if they all have the capacity of the smallest physical drive.

For example, if you group three 1 GB drives and one 2 GB drive into an array, the total capacity of the array is 1 GB times 4, or 4 GB, not the 5 GB physically available. The remaining space on the 2 GB drive is unusable capacity.

SR-8i SR-7t For the ServeRAID-8i controller and ServeRAID-7t controller, the remaining space is usable capacity. That is, you can use the remaining space to define another logical drive; see Example: Usable and unusable capacity.

Conversely, if you add a smaller drive to an array of larger drives, such as a 1 GB drive to a group containing three 2 GB drives, the total capacity of that array is 4 GB, not the 7 GB physically available. Therefore, the optimal way to create arrays is to use physical drives that have the same capacity.

A hot-spare drive also must be at least as large as the smallest drive in the array.

More information

- Example: Usable and unusable capacity
- Example: Total disk capacity
- Understanding creating basic arrays

Clearing a disk drive: Use this action to remove all data from a disk drive:

1. In the Physical devices view, click (online physical drive).
2. Right-click **Clear**.
3. Click **Yes** to confirm.

More information

- Identify physical drive

Verifying physical drives: Use this action to check a physical drive for inconsistent or bad data. The Verify action examines each sector and block on the disk to ensure that it is readable. It does not repair the drive if bad data is found.

To verify a physical drive:

1. In the Physical devices view, click (physical drive).
2. Right-click **Verify**.
3. Click **Yes** to verify the drive.

   Note: To check the progress of the operation, switch to text view.

More information

- Understanding synchronizing logical drives

Initializing physical drives: Use this action to erase the metadata for a ready physical drive. The metadata includes all logical drive definition data.

To initialize a physical drive:

1. In the Physical devices view, click (physical drive).
2. Right-click **Initialize**.
3. Click **Yes** to initialize the drive.

   Note: To check the progress of the operation, switch to text view.

More information

- Verifying physical drives

Manually rebuilding a defunct drive: Use this action to rebuild a critical logical drive when a physical drive in the array is defunct. Normally, the controller rebuilds a logical drive when it detects the removal and reinsertion of a drive that is part of an array. Use this operation to force the rebuild to begin immediately, without physically removing and replacing the disk drive.

Note: The controllers can rebuild RAID level-1, level-1E, level-5, level-5E, level-5EE, level-10, level-1E0, and level-50 logical drives. They cannot, however, rebuild RAID level-0 logical drives because RAID level-0 is not redundant.

Complete the following steps to manually rebuild a defunct drive:

1. In the Physical devices view, click (defunct physical drive).
2. Right-click **Rebuild drive**.
3. Click **Yes** to confirm.

Note: During the rebuild operation, the defunct drive enters the rebuild state and the logical drive remains critical until the rebuild operation is completed.

More information

- Rebuilding a defunct drive (concepts)
- Recovering from defunct drives

Removing a defunct drive:

Note: This action is not supported on network storage controllers. It is supported on the integrated RAID controller in bootable-CD mode only.

You can use this action *only* on defunct physical drives that are **not** part of any array. Although the ServeRAID software implicitly removes defunct drives from the configuration, this action is useful if you want to remove a defunct drive from the Physical devices view so you can view an accurate configuration.

1. In the Physical devices view, click (defunct physical drive).
2. Right-click **Remove defunct drive**. The ServeRAID Manager removes the defunct drive from the configuration.

Setting a drive state to online:

Note: This action is not supported on the integrated RAID controller.

You can use this action *only* when the following is true:

- The defunct drive is part of an array.
- There are no critical logical drives in the array.

Attention: When you set a defunct drive to online, there is a high risk of data loss.

1. In the Physical devices view, click (defunct physical drive).
2. Right-click **Set drive state to online**.
3. Click **Yes** to confirm.

Setting a drive state to defunct:

Note: This action is not supported on the ServeRAID-8i controller and ServeRAID-7t controller.

You can use this action only on a physical drive that is part of an array. If a physical drive is part of an array, you must change the drive state to defunct before removing the drive from the server or enclosure. Doing so eliminates the risk of losing data. After you replace the physical drive, you can rebuild the affected array.

This action is useful in situations such as the following:

- You want to replace a physical drive that is marked with a PFA.
- You want to replace a physical drive that is unwarranted.
- You have an array built with one or more physical drives that are not the same size. You can replace the smaller physical drives in the array with larger physical drives so that the array no longer has unusable capacity and an inefficient configuration.

Attention:

1.

   If you do *not* change an online physical drive state to defunct before removing the drive from the server or enclosure, you risk losing data or damaging the physical drive.
2. You cannot use this action on a physical drive undergoing a rebuild operation.
3. (For non-RAID level-x0 logical drives) If you choose to use this action and a logical drive is:
   - In a critical state, the logical drive state *will* change to offline.
   - RAID level-0 and in an okay state, the logical drive *will* change to offline.
   - Not RAID level-0 and in an okay state, the logical drive state *will* change to critical.
4. (For RAID level-x0 logical drives) If you choose to use this action and a logical drive is in a critical state, the logical drive state *might* change to offline.

1. In the Physical devices view, click (physical drive) or (physical drive).
2. Right-click **Set drive state to defunct**. A confirmation window opens to warn that this action might corrupt data.

   Note: If the logical drive that contains this physical drive is already offline, the ServeRAID Manager does not display this warning because the data is already corrupt. Continue to step 4.
3. Click **Yes**. The ServeRAID Manager sets the drive status to defunct.
4. Replace the selected physical drive. If the logical drive is not offline, a rebuild automatically occurs.

More information

- Replacing physical drives in an array
- Understanding when RAID level-x0 logical drives change state

Rebuilding a defunct drive: A physical drive is marked defunct when there is a loss of communication between the controller and the physical drive. This can be caused by any of the following:

- An improperly connected cable, physical drive, or controller
- Loss of power to a drive
- An improperly assembled SCSI channel in an unsupported configuration
- A defective cable, backplane, physical drive, or controller
- Connecting unsupported SCSI devices (such as tape drives or CD-ROM drives) to the same SCSI channel used for an array

In each case, after the communication problem is resolved, a rebuild operation is required to reconstruct the data for the device in its disk array. The ServeRAID controllers can reconstruct RAID level-1, level-1E, level-5, level-5E, level-5EE, level-10, level-1E0, and level-50 logical drives. They cannot, however, reconstruct data stored in RAID level-0 logical drives because RAID level-0 is not redundant. If an array contains only RAID level-0 logical drives, the logical drives in the array are marked offline and the logical drives contain damaged data. You cannot rebuild the logical drives. You must correct the cause of the failure or replace the physical drives; then, you must restore your data.

To prevent data-integrity problems, the ServeRAID controllers set the RAID level-0 logical drives in the affected array to blocked during the rebuild operation for RAID level-1 or RAID level-5. After the rebuild operation is completed, you can unblock the RAID level-0 logical drives and access them once again. Remember, however, that the RAID level-0 logical drives contain damaged data.

Note: For logical drives in an IBM ServeRAID Cluster Solution:

- Because shared logical drives can have *only* one logical drive for each array, blocking a RAID level-0 logical drive during a rebuild operation does *not* apply to shared logical drives.
- Because non-shared logical drives can have *more* than one logical drive for each array, blocking a RAID level-0 logical drive during a rebuild operation *does* apply to non-shared logical drives.

More information

- Recovering from defunct drives
- Rebuilding a hot-swap drive
- Manually rebuilding a defunct drive
- Changing the rebuild rate (action)
- Unblocking a logical drive (action)
- RAID level-0
- RAID level-1
- RAID level-1E
- RAID level-5
- RAID level-5E
- RAID level-5EE
- RAID level-x0

Recovering from defunct drives: If the defunct drives are not part of an array, contact your service representative.

If a physical drive fails in an array or multiple physical drives fail in separate arrays (one physical drive per array), complete the following steps:

1. Replace each defunct physical drive. The controller starts the rebuild operation when it detects the removal and reinsertion of a drive that is part of an array.

   Note: (For a configuration that contains a hot-spare drive) If you replace a failed physical drive, it is not necessary to position the new physical drive on the same SCSI channel and SCSI ID as the original hot-spare drive. The replacement physical drive is automatically incorporated into the configuration as a hot-spare drive. Here is an example of how this works:

   a. The original configuration consists of a RAID level-5 logical drive composed of four physical drives. The physical drives are connected to SCSI channel 1; they are assigned SCSI IDs 0, 1, 2, and 3. SCSI ID 3 is a hot-spare drive.
   b. The physical drive at Channel 1, SCSI ID 2, fails; the logical drive enters the critical state.
   c. The hot-spare drive at Channel 1, SCSI ID 3, is rebuilt into the array.
   d. You remove the failed physical drive at Channel 1, SCSI ID 2, and replace it with a new physical drive. The new physical drive at Channel 1, SCSI ID 2, is automatically assigned to be a hot-spare drive.

2. If a rebuild operation is in progress, wait until the rebuild is complete. Otherwise, go to step 3.

   Note: If you are replacing multiple defunct drives, you must wait for each rebuild operation to complete before starting subsequent rebuild operations.
3. Verify the cables, physical drives, and controllers are installed properly.
4. Attempt to rebuild the defunct physical drive by performing a hot-swap rebuild.
5. If the hot-swap rebuild fails, contact your service representative.

More information

- Rebuilding a defunct drive
- Rebuilding a hot-swap drive
- Manually rebuilding a defunct drive
- Changing the rebuild rate (action)
- Unblocking a logical drive (action)
- RAID level-0
- RAID level-1
- RAID level-5

Viewing shared drives:

Note: This action is not supported when using the following:
- Integrated RAID controller
- ServeRAID-5i controller
- ServeRAID-6i/6i+ controller
- ServeRAID-7k controller
- ServeRAID-7t controller
- ServeRAID-8i controller
- HostRAID controller

Use this action to view the physical drives owned by a cluster partner system in the Enterprise view.

1. In the Enterprise view, click (controller).
2. Right-click **Clustering actions** → **View shared drives** to view the physical drives owned by a partner system.
3. Select the **Enable view shared drives** box. To disable viewing shared drives, deselect the **Enable view shared drives** box.
4. Define the location of the partner system online drive by specifying the SCSI channel and SCSI ID in the drop-down lists.
5. Click **OK**.

The ServeRAID Manager displays the physical drives owned by the partner system in the Enterprise view as (Reserved). You only can view these reserved drives; you cannot perform any actions upon them.

More information

- Configure controller for clustering

Hot-spare drive object: Use this object, in the Physical devices view, to work with a hot-spare drive in your configuration or to display its properties.

Possible subobject

None.

Warning and error conditions

- ⚠ (warning) is displayed if the hot-spare drive reports a PFA or is a non-warranted physical drive.
- ✖ (error) is displayed if the hot-spare drive is defunct.

Actions

- Create standby hot-spare drivce
- Delete hot-spare drive

Hot-spare configuration summary: If a hot-spare drive of the appropriate size is available, this column reports **Yes**.

If the logical drive is redundant (that is, RAID level-1, 1E, or 5), but a hot-spare drive of the appropriate size is not available, this column reports **No** (in red).

If the logical drive is RAID level-0, this column reports **No** (in black).

Creating a hot-spare drive:

Note: This action is supported on the integrated RAID controller in bootable-CD mode only.

Use this action to create a hot-spare drive.

1. In the Physical devices view, click (ready physical drive) or (standby hot-spare drive).
2. Click ⊞ in the Logical devices view. The ServeRAID Manager creates the hot-spare drive.

   Note: A hot-spare drive with a blue plus sign next to it indicates that the logical devices are protected by the hot spare. A hot-spare drive with a yellow plus sign next to it indicates that:
   - The drive is too small to protect any drives in the logical devices
   - None of the logical devices support hot spares (for example, a RAID Level-0 array)
   - No logical devices are defined

More information

- Choosing to create a hot-spare drive
- RAID levels that can use a hot-spare or standby hot-spare drive

Creating a standby hot-spare drive:

Note: This action is not supported on the ServeRAID-8i SAS controller, ServeRAID-7t SATA controller, and the integrated RAID controller.

Use this action to create a standby hot-spare drive.

1. In the Physical devices view, click (ready physical drive) or (hot-spare drive).

2. Right-click **Create standby hot-spare drive**. The ServeRAID Manager creates the standby hot-spare drive.

   Note: A standby hot-spare drive with a blue plus sign next to it indicates that the logical devices are protected by the hot spare. A standby hot-spare drive with a yellow plus sign next to it indicates that:

   - The drive is too small to protect any drives in the logical devices
   - None of the logical devices support hot spares (for example, a RAID Level-0 array)
   - No logical devices are defined

More information

- Choosing to create a standby hot-spare drive
- Deleting a hot-spare drive
- RAID levels that can use a hot-spare or standby hot-spare drive

Creating an assigned hot-spare drive:

Note: This action is supported only on the ServeRAID-7t, ServeRAID-8i, and external storage enclosures (networked storage).

Use this action to assign one or more hot-spare drives to a logical drive. An assigned hot-spare drive is a dedicated spare: it is used only to rebuild the logical drive to which it is assigned. An unassigned hot-spare drive is ″global″: it can be used to rebuild any logical drive on the controller.

Note: You cannot assign a hot-spare drive while a logical drive is being built or modified (such as a logical drive migration).

1. In the Physical devices view, click (ready physical drive).
2. Right-click **Create assigned hot-spare drive for** the logical device name. The ServeRAID Manager creates a dedicated hot-spare drive for the logical drive.

More information

- Choosing to create a hot-spare drive
- RAID levels that can use a hot-spare or standby hot-spare drive

Deleting a hot-spare drive:

Note: This action is supported on the integrated RAID controller in bootable-CD mode only.

Use this action to delete a hot-spare physical drive from your configuration.

1. In the Physical devices view, click (hot-spare drive).
2. Right-click **Delete hot-spare drive**. The ServeRAID Manager deletes the hot-spare drive from the configuration by setting the drive state to ready. If you removed the drive from the server, the ServeRAID Manager removes the drive from the configuration.

Deleting a standby hot-spare drive:

Note: This action is not supported on the ServeRAID-8i SAS controller, ServeRAID-7t SATA controller and the integrated RAID controller.

Use this action to delete a standby hot-spare physical drive from your configuration.

1. In the Physical devices view, click (standby hot-spare drive).
2. Right-click **Delete standby hot-spare drive**. The ServeRAID Manager deletes the standby hot-spare drive from the configuration by setting the drive state to ready. If you have physically removed the drive, the ServeRAID Manager removes the drive from the configuration.

Choosing to create a hot-spare or standby hot-spare drive: Hot-spare drives and standby hot-spare drives supply additional protection to a RAID configuration. If a physical drive fails in configurations other than RAID level-0 and RAID level-00, the hot-spare drive automatically replaces the failed physical drive. Subsequently, the logical drive can be rebuilt automatically.

Note: You cannot use hot-spare drives with RAID level-0 and RAID level-00.

You can create a hot-spare drive in the Configuration wizard when you first configure the controller. If you choose not to create a hot-spare drive at that time, you can do so later.

More information

- Creating a hot-spare drive (action)
- Creating a standby hot-spare drive (action)
- Creating an assigned hot-spare drive (action)
- Deleting a hot-spare drive (action)
- RAID levels that can use a hot-spare or standby hot-spare drive

Rebuilding a hot-swap drive: A hot-swap rebuild refers to a rebuild operation that is started by the controller when it detects that a drive that is part of an array and in the defunct state has been removed and reinserted on the SCSI cable or backplane. The reinsertion of the physical drive, whether it is the same drive or a new drive, will trigger the controller to start the rebuild operation. During the rebuild operation, the drive being rebuilt is in the rebuild state, and the logical drive remains critical until the rebuild operation has been successfully completed.

On most servers, when a hot-spare drive is available, the rebuild operation begins automatically without the need to replace the failed drive. If more than one drive fails within the same array, no rebuild takes place. If multiple drives fail in separate arrays (one physical drive per array), the controller initiates a rebuild operation for the logical drives within the array containing the first failed physical drive. This rebuild operation is performed on the *first* hot-spare drive of sufficient size to become a valid member of the array.

Complete the following steps to start a hot-swap rebuild:

1. Without removing the drive completely, gently remove the physical drive from the server, using the handle of the hot-swap tray. If necessary, refer to the documentation that comes with your server for information about removing a physical drive.

   Attention: When power is removed from a hot-swap drive, the drive immediately parks the heads, locks the actuator in the ″landing zone,″ and begins spinning down. However, the spinning down of the disk might require up to 20 seconds after power is removed. Do not move the drive while it is spinning down. Moving the drive while it is spinning down might damage the drive.
2. Wait 20 seconds to allow the physical drive to completely stop spinning.

3. If you are certain there is nothing wrong with the physical drive you removed, gently reinstall the drive into the server. Make sure the drive is completely installed in the backplane connector.

   Otherwise, replace the physical drive with a new drive that is the same size (or larger) and continue with the rebuild operation.

   Note:

   a. If multiple drives fail in separate arrays (one physical drive per array), replace each defunct physical drive. If multiple physical drives fail at the same time within the *same* array, contact your service representative.
   b. Although it is possible to rebuild a defunct physical drive to an online physical drive that is defective, avoid doing so.

More information

- Rebuilding a defunct drive
- Recovering from defunct drives
- Manually rebuilding a defunct drive
- Changing the rebuild rate (action)
- Unblocking a logical drive (action)

Enclosure management device object: Use this object, in the Physical devices view, to work with a storage enclosure attached to controller with a SAF-TE processor.

Note: The integrated RAID controller does not support enclosure monitoring. Any enclosure in an integrated RAID controller RAID subsystem is displayed in the ServeRAID Manager Enterprise view; however, the ServeRAID Manager and the ServeRAID Manager agent will not generate any events or enclosure status information.

Possible subobject

None.

Warning and error conditions

- ⚠ (warning) is displayed if the enclosure has the following conditions:
  - A fan has failed.
  - A power supply has failed.
- ✖ (error) is displayed if the enclosure has the following conditions:
  - It is not responding.
  - It is overheating.
  - Multiple fans have failed.
  - Multiple power supplies have failed.
- (enclosure) is displayed if the enclosure is one of the following:
  - Working properly.
  - Not supported.
  - Supported, but not monitored for problems. For example, enclosures that are internal to a server do not have separate fans or power supplies.

Actions

Identify enclosure

CD-ROM drive object: Use this object to display CD-ROM drive properties.

Note: The ServeRAID-5i, 6i/6i+, 6M, and 7t controllers do not support CD-ROM drives.

Possible subobject

None.

Warning and error conditions

CD-ROM drive objects cannot have warning or error conditions.

Actions

None.

Removable-media drive object: Use this object to display removable-media drive properties.

Possible subobject

None.

Warning and error conditions

Removable-media drive objects cannot have warning or error conditions.

Actions

None.

Tape drive object: Use this object to display tape drive properties.

Note: The ServeRAID-5i controller does not support tape drives and physical drives on the same channel.

Possible subobject

None.

Warning and error conditions

Tape drive objects cannot have warning or error conditions.

Actions

None.

Using display groups:

Note: Display groups are not available when using:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Display groups allow you organize related systems in the Enterprise view under an expandable tree object. Using actions on the System object, the Management station object, and the Display group object, you can:

- Create, delete, and rename display groups
- Move a system into a display group
- Remove a system from a display group
- Move a system from one display group to another

In this example, managed systems are organized by location, using three display groups.

Display groups are sorted alphabetically. They appear in the tree below all systems that are not part of any display group.

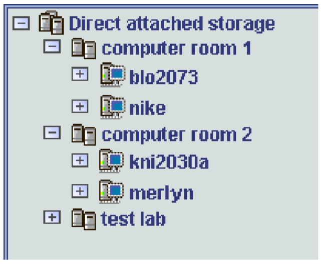

Note: A system can be part of only one display group at a time; it cannot appear in multiple display groups.

More information

- The Enterprise view

Creating a display group:

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to add a system or enclosure to a new display group. This action creates the display group and moves the system into it in a single step.

1. In the Enterprise view, click (system) or (enclosure).
2. Right-click **Change display group** → **new group**. The ″Group name″ window opens.
3. Enter a name for the group.
4. Click **OK**.

More information

- Using display groups
- Moving systems in and out of a display group (action)
- Deleting or renaming a display group (action)

Deleting or renaming a display group:

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to delete or rename a display group.

1. In the Enterprise view, click (display group).
2. To delete the display group, right-click **Delete display group.** All systems in the group move back to their original location in the tree.
3. To rename the display group:
   a. Right-click **Rename display group**. The ″Group name″ window opens.
   b. Enter a new name for the display group.
   c. Click **OK**.

More information

- Using display groups
- Creating a display group (action)
- Moving systems in and out of a display group (action)

Moving systems in and out of a display group:

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to move a system or enclosure into an existing display group; to move the system or enclosure from one display group to another; or to remove the system or enclosure from a display group.

1. In the Enterprise view, click (system) or (enclosure).
2. To move the system into a display group or from one display group to another, right-click **Change display group** and then the display group name.
3. To remove the system from the display group, right-click **Change display group** → **none**.

Note: When you remove the last system in a display group, the ServeRAID Manager deletes the display group.

More information

- Using display groups
- Creating a display group (action)
- Deleting or renaming a display group (action)

Changing the display options: You can sort objects in the ServeRAID Manager Enterprise view alphabetically or chronologically. You can also view the disk capacity unit of measure in megabytes (MB), gigabytes (GB), or terabytes (TB).

Note: The local system always appears first when you sort objects alphabetically.

Complete the following steps set the display options:

1. In the ″User Preferences″ window, click the **Display options** tab.
2. From the System tree sorting list, select chronological or alphabetical.

Note: This action is not supported in bootable-CD mode.

3. From the Capacity display units list, select one of the following:
   - MB (megabytes)
   - GB (gigabytes)
   - TB (terabytes)
   - Auto-select

   Choose Auto-select to let the ServeRAID Manager choose the unit of measure for you, based on the drive size. With this option, the ServeRAID Manager may display different drives with different units.
4. Click **OK**.

Monitoring systems over a network: You can monitor other systems over a network. To do so, the ServeRAID Manager must be installed and running on your local system and on the other systems. Then, you can add the remote systems to your ServeRAID Manager Enterprise view. Also, you can choose to use the ServeRAID Manager agent, instead of the ServeRAID Manager program.

Note: You cannot monitor other systems when using:
- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

In addition to the ServeRAID Manager, you can use the following:
- The Notification Manager is a tool that you can use to notify other systems running the ServeRAID Manager of all events that occur on this (that is, the local) system.
- Use the Email Notification Manager to notify users by email when events occur on the local system. You can send email notifications to anyone on the network; the recipient does not have to be running ServeRAID Manager.
- Use the Task Manager to manage scheduled tasks on local or remote systems, such as logical drive migration or synchronization. You can modify the task schedule, delete the scheduled task, or view task properties.

More information
- Adding a remote system (action)
- Using the ServeRAID Manager agent
- Starting the ServeRAID Manager agent
- Verifying that the ServeRAID Manager agent is running
- Using the Notification Manager
- Using the Email Notification Manager
- Using the Task Manager

Using the Notification Manager:

Note: You cannot use the Notification Manager when using:
- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

The Notification Manager is a tool that you can use to notify other systems running the ServeRAID Manager of all events that occur on this (that is, the local) system. The user of the system types system names in the notification list. Each system in the list is notified of all events that occur on the local system.

You can use the Notification Manager to:

- Add a system to the notification list
- Delete a system from the notification list
- Modify system properties in the notification list
- Send a test event to a system in the notification list
- Monitor events sent from the Notification Manager

The Notification Manager is enabled by default. To disable the Notification Manager, click **Actions** → **Disable notifications** . If you disable the Notification Manager, the events are generated, but not dispatched to remote systems.

An example

You install SystemA in a lab with a ServeRAID Manager subsystem. You run the ServeRAID Manager on SystemA to monitor for events and problems, but you want to monitor from your workstation and not from SystemA. You open the SystemA Notification Manager from the ServeRAID Manager and define your workstation in the notification list. When running the ServeRAID Manager from your workstation, you are notified of all problems and events that occur on SystemA.

Note: You can use the Email Notification Manager to notify users of events by email. This allows you to notify users who have not installed or are not currently running ServeRAID Manager.

The Notification Manager and its events

Events generated by the Notification Manager include the following:

- Progress information, such as rebuilds, synchronizations, and migrations.
- Problems, such as defunct physical drives and PFA errors.
- Changes to the local configuration, such as creating a hot-spare drive or defining a logical drive.

When an event is generated on a system, the Notification Manager connects with each system in the notification list and relays the event to these systems' ServeRAID Manager. If the Notification Manager successfully connects and sends the event, the notification list updates its ″ Last event sent″ column. If the Notification Manager does not successfully connect or send the event, the Notification Manager:

- Updates the ″Last event sent″ column with a ⊗ notifying you of a problem.
- Logs an event in the Notification Manager event viewer detailing why the event could not be sent.

The Notification Manager interface

The Notification Manager consists of the following:

Toolbar
Provides quick-path icons for common tasks.

Notification list
Displays the remote systems configured to receive event notification.

Notification event viewer
Displays status information for the Notification Manager.

In addition to displaying the events in the event viewer, the Notification Manager appends each event to a logging file, RAIDNOT.LOG. If this file exceeds 200 KB, the ServeRAID Manager copies the file to RAIDNOT.OLD and creates a new RAIDNOT.LOG. If there is a RAIDNOT.OLD already, the ServeRAID Manager overwrites it.

More information

- Adding a system to the notification list
- Deleting a system from the notification list
- Modifying system properties in the notification list
- Sending a test event to a system in the notification list
- Monitoring events sent from the Notification Manager
- The Notification Manager menubar
- The Notification Manager toolbar
- Using the Email Notification Manager
- Using the ServeRAID Manager agent

Notification Manager menubar: The Notification Manager menu bar has the following functions:

File >

Clear event log
: Click Notification Manager from the menu. Clears the current contents of the Notification Manager event viewer. This option does not clear or delete the event logging file.

Close
: Closes the Notification Manager.

View >

Toolbar
: Turns the toolbar on and off. The default is on (selected).

Actions >

Add system

Delete system

Modify system

Send test event

Help >

Information about this window
: View context-sensitive information for the current window.

Search
: Searches for one or more specified words in ServeRAID Manager Assist and displays a list of topics that include the words.

Contents
: Presents the ServeRAID Manager Assist contents. You can use the contents to acquaint yourself with ServeRAID Manager Assist topics.

About ServeRAID Manager
Reports the ServeRAID Manager version number, copyright, and legal information.

Notification Manager toolbar: The Notification Manager toolbar has the following functions:

Add system

Delete system

Modify system

Information about this window View context-sensitive information for the current window.

Adding a system in the Notification Manager:

Note: This action is not supported when using the ServeRAID Manager in bootable-CD mode.

Use this action to add a new system to the notification list. Every new system that you add to the list is notified of all events that occur on this system.

1. Click (Add system).
2. Type the host name or TCP/IP address of the system that you want to add.
3. Type the TCP/IP port for the remote system Startup port number. The default port is 34571. If you changed the system Startup port number in the ″User preferences″ window on the remote system, type that number here.
4. Click **Add**.
5. Continue to add systems, or click **Cancel** to close the window.

After you add the system to the notification list, the ServeRAID Manager will attempt to connect to the remote system and send an initial set of events. The Manager uses these events to update the ServeRAID Manager interface running on the remote system.

More information

- Using the Notification Manager
- Deleting a system from the notification list
- Modifying system properties in the notification list
- Sending a test event to a system in the notification list
- Monitoring events sent from the Notification Manager

Deleting a system in the Notification Manager:

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to delete a remote system from the notification list. After deleting the remote system, the Notification Manager no longer notifies the remote system of events occurring on the local system.

1. Click the system you want to delete from the notification list.
2. Click (Delete system).
3. Click **Yes** to confirm. The Notification Manager deletes the system from the notification list.

More information

- Using the Notification Manager
- Adding a system to the notification list
- Modifying system properties in the notification list
- Sending a test event to a system in the notification list
- Monitoring events sent from the Notification Manager

Monitoring events sent from the Notification Manager: For each system in the notification list, you can view details about the last event the Notification Manager sent (or tried to send) to that system.

In the notification list, double-click the **Last event sent** column for the system in which you are interested. The ″Last event detail″ window opens. This window contains information about the last event the Notification Manager sent (or tried to send) to the selected system. The event information includes the following: status, type, date, time, source, target, and description.

More information

- Using the Notification Manager
- Adding a system to the notification list
- Deleting a system from the notification list
- Modifying system properties in the notification list
- Sending a test event to a system in the notification list

Modifying a system in the Notification Manager:

Note: This action is not supported when using the following:
- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to modify a remote system host name, TCP/IP address, or port number:

1. Click the system that you want to modify from the notification list.
2. Click (Modify system). The ″System properties″ window opens.
3. Change the system properties in the appropriate entry fields.
4. Click **OK**.

More information

- Using the Notification Manager
- Adding a system to the notification list
- Deleting a system from the notification list
- Sending a test event to a system in the notification list
- Monitoring events sent from the Notification Manager

Sending a test event in the Notification Manager:

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to send a test event to a remote system:

1. Click the system in the notification list that you want to send a test event.

   Note: If you click no system or more than one system, this action is disabled.
2. Click **Actions** → **Send test event**. A window opens to report success or failure. The remote system will respond to the test event with a beep.
3. Click **OK**.

If the test fails, verify the following and try the test send again:

1. You typed the correct TCP/IP information.
2. The ServeRAID Manager is running on the remote system.

More information

- Using the Notification Manager
- Adding a system to the notification list
- Deleting a system from the notification list
- Modifying system properties in the notification list
- Monitoring events sent from the Notification Manager

Changing the alarm settings:

Note: This action is not supported in bootable-CD mode.

Warning and Error events cause an audible alarm to sound every five minutes, notifying you of the event. Any Error or Warning event can trigger the alarm if the event causes the system to enter the critical state or the event arrives while the system status is critical. Once triggered, the repeating alarm continues to sound as long as any system in the ServeRAID Manager Enterprise view is critical.

Some events, such as the non-warranted drive event, do not affect the system status in the Enterprise view; if the system status is not critical, these events will not sound the alarm.

You can adjust the alarm settings with the following steps:

1. In the User Preferences window, click the **Alarm settings** tab.
2. Select the check box to enable or disable the repeating alarm.

   Note: If you disable the alarm, you will not hear an audible alarm when you receive problem events.
3. If you have enabled the alarm, you can adjust the time interval (in seconds) that you want between each alarm. The default is 300 seconds (5 minutes).
4. If you have enabled the alarm, you can adjust the length of time (in beeps) that you want the alarm to continue. The default is 3 beeps.
5. Click **OK**.
6. Restart the ServeRAID Manager for these settings to take effect.

Using the Email Notification Manager:

Note: You cannot use the Email Notification Manager when using:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

The Email Notification Manager is a tool that you can use to notify users by e-mail of events that occur on this (that is, the local) system. The user of the system types recipient names and e-mail addresses in the e-mail notification list. Each recipient in the list is notified of all or selected events that occur on the local system.

You can use the Email Notification Manager to:

- Add a recipient to the email notification list
- Delete a recipient from the email notification list
- Modify recipient properties in the email notification list
- Send a test message to a recipient in the email notification list
- Monitor messages sent from the Email Notification Manager
- Configure the Email Notification Manager

The Email Notification Manager is enabled by default. To disable the Email Notification Manager, click **Actions** → **Disable Email Notifications** . If you disable the Email Notification Manager, events are generated but users do not receive e-mail notification when the events occur.

An example

You install SystemA in a lab with a ServeRAID Manager subsystem. You run the ServeRAID Manager on SystemA to monitor for events and problems. You want to monitor SystemA from your workstation but do not have ServeRAID Manager installed. You open the SystemA Email Notification Manager from the ServeRAID Manager and add your name and e-mail address to the e-mail notification list. You are notified by e-mail of all problems and events that occur on SystemA.

The Email Notification Manager and its events

Events monitored by the Email Notification Manager include the following:

- Progress information, such as rebuilds, synchronizations, and migrations.
- Problems, such as defunct physical drives and PFA errors.
- Changes to the local configuration, such as creating a hot-spare drive or defining a logical drive.

Each event has an **event type** indicating its severity: error, warning, or informational. You can configure a recipient to be notified of errors, errors and warnings, or errors, warnings, and informational events.

When an event is generated on a system, the Email Notification Manager attempts to send a message to each recipient in the email notification list. If the Email Notification Manager successfully sends the message, it updates the ″ Last message sent″ column in the email notification list. If the Email Notification Manager does not successfully connect or send the message, the Email Notification Manager:

- Updates the ″Last message sent″ column with a ⊗ notifying you of a problem.
- Logs an event in the Email Notification Manager event viewer detailing why the message could not be sent.

The Email Notification Manager interface

The Email Notification Manager consists of the following:

Toolbar
Provides quick-path icons for common tasks.

Email notification list
Displays the recipients configured to receive e-mail notifications.

Email notification event viewer
Displays status information for the Email Notification Manager.

In addition to displaying the events in the event viewer, the Email Notification Manager appends each event to a logging file, RAIDSMTP.LOG. If this file exceeds 200 KB, the ServeRAID Manager copies the file to RAIDSMTP.OLD and creates a new RAIDSMTP.LOG. If there is a RAIDSMTP.OLD already, the ServeRAID Manager overwrites it. The Email Notification Manager appends SMTP transport errors in a separate log file, SMTPERR.LOG.

More information

- Adding a recipient to the email notification list
- Deleting a recipient from the email notification list
- Modifying recipient properties in the email notification list
- Sending a test message to a recipient in the email notification list
- Monitoring messages sent from the Email Notification Manager
- Configuring the Email Notification Manager
- The Email Notification Manager menubar
- The Email Notification Manager toolbar
- The Email Notification Manager message format
- Using the ServeRAID Manager agent

About Email Notification Manager:

Email Notification Manager menu bar: The Email Notification Manager menu bar has the following functions:

File →

Clear event log
Click Email Notifications from the menu. Clears the current contents of the Email Notification Manager event viewer. This option does not clear or delete the event logging file.

Close Closes the Email Notification Manager.

View →

Toolbar
Turns the toolbar on and off. The default is on (selected).

Actions →

Add email recipient

Delete email recipient

Modify email recipient

Send test message

SMTP server settings

Help →

Information about this window
View context-sensitive information for the current window.

Search
Searches for one or more specified words in ServeRAID Manager Assist and displays a list of topics that include the words.

Contents
Presents the ServeRAID Manager Assist contents. You can use the contents to acquaint yourself with ServeRAID Manager Assist topics.

About ServeRAID Manager
Reports the ServeRAID Manager version number, copyright, and legal information.

Email Notification Manager menu bar: The Email Notification Manager menu bar has the following functions:

File >

Close Closes the Email Notification Manager.

View >

Toolbar
Turns the toolbar on and off. The default is on (selected).

Actions >

Add email recipient

Delete email recipient

Modify email recipient
Send test message
SMTP server settings

Help >

Information about this window
View context-sensitive information for the current window.

Search
Searches for one or more specified words in ServeRAID Manager Assist and displays a list of topics that include the words.

Contents
Presents the ServeRAID Manager Assist contents. You can use the contents to acquaint yourself with ServeRAID Manager Assist topics.

About ServeRAID Manager
Reports the ServeRAID Manager version number, copyright, and legal information.

Email Notification Manager toolbar: The Email Notification Manager toolbar has the following functions:

Add email recipient

Delete email recipient

Modify email recipient

Information about this window View context-sensitive information for the current window.

Email Notification Manager toolbar: The Email Notification Manager toolbar has the following functions:

Add email recipient

Delete email recipient

Modify email recipient

Information about this window View context-sensitive information for the current window.

The Email Notification Manager Message Format: When an event is generated on a system, the Email Notification Manager sends a message to each recipient in the email notification list. The message format is:

To: <recipient email address>

From: <administrator email adress>

Subject: ServeRAID Manager Event Notification - Event type <event type>

This message was generated by ServeRAID Manager Agent.

Please do not reply to this message.

Event Description: <event description>

Event Type: <event type>

Event Source: <agent domain name>

Date: <date>

Time: <time>

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Deleting a recipient from the email notification list
- Modifying recipient properties in the email notification list
- Sending a test message to a recipient in the email notification list

The Email Notification Manager Message Format: When an event is generated on a system, the Email Notification Manager sends a message to each recipient in the email notification list. The message format is:

```
From: <management station email adress>
Sent: <date>
To: <recipient email address>
Subject: <management station> - <Event message>

This message was generated by the ServeRAID Manager. Please
do not reply to this message.

Event Description: <event description>
Event Type: <event type>
Event Source: <management station> / <IP address>
Time: <date> <time>
```

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Deleting a recipient from the email notification list
- Modifying recipient properties in the email notification list
- Sending a test message to a recipient in the email notification list

Adding a recipient to the Email Notification Manager:

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to add a recipient to the email notification list. Every recipient that you add to the list is notified of all or selected events that occur on this system.

1. Click 👤 (Add recipient).
2. Type the name of the recipient that you want to add to the list.
3. Type the recipient’s e-mail address. Include the user name and domain, such as jack_smith@synteca.com.
4. Select the type of event you want the recipient to be notified of: error, error and warning, or error, warning, and informational.
5. Click **Add**.
6. Continue to add recipients, or click **Cancel** to close the window.

More information

- Using the Email Notification Manager
- Deleting a recipient from the email notification list
- Modifying recipient properties in the email notification list
- Sending a test message to a recipient in the email notification list
- Monitoring messages sent from the Email Notification Manager

Deleting a recipient from the Email Notification Manager:

Note: This action is not supported when using the following:
- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to delete a recipient from the e-mail notification list. After deleting a recipient, the Email Notification Manager no longer notifies the user of events occurring on the local system.

1. Click the recipient you want to delete from the e-mail notification list.
2. Click (Delete recipient).
3. Click **Yes** to confirm. The Email Notification Manager deletes the recipient from the e-mail notification list.

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Modifying recipient properties in the email notification list
- Sending a test message to a recipient in the email notification list
- Monitoring messages sent from the Email Notification Manager

Modifying a recipient in the Email Notification Manager:

Note: This action is not supported when using the following:
- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to modify an e-mail recipient name, address, or event type:

1. Click the recipient that you want to modify in the e-mail notification list.
2. Click (Modify email recipient). The ″Modify email recipient″ window opens.
3. Change the recipient properties in the appropriate fields.
4. Click **OK**.

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Deleting a recipient from the email notification list
- Sending a test message to a recipient in the email notification list
- Monitoring messages sent from the Email Notification Manager

Sending a test message in the Email Notification Manager:

Note: This action is not supported when using the following:
- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to send a test message to a recipient in the e-mail notification list:

1. Click the recipient in the e-mail notification list that you want to send a test message.

   Note: If you click no recipient or more than one recipient, this action is disabled.

2. Click **Actions** → **Send test message**. If the Email Notification Manager does not successfully connect or send the message, it:
   - Updates the ″Last message sent″ column with a ⊗ notifying you of a problem.
   - Logs an event in the Email Notification Manager event viewer detailing why the message could not be sent.
3. Click **OK**.

If the test fails, verify the following and try the test again:

1. You typed the correct e-mail address when you added the recipient to the Email Notification Manager.
2. You typed the correct SMTP server address when you configured the Email Notification Manager.

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Deleting a recipient from the email notification list
- Modifying recipient properties in the email notification list
- Monitoring messages sent from the Email Notification Manager

Monitoring messages sent from the Email Notification Manager: For each recipient in the e-mail notification list, you can view details about the last message the Email Notification Manager sent (or tried to send) to that recipient.

In the e-mail notification list, double-click the **Last message sent** column for the recipient in which you are interested. The ″Last message detail″ window opens. This window contains information about the last message the Email Notification Manager sent (or tried to send) to the recipient. The message information includes the following: status, type, date, time, source, target, and description.

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Deleting a recipient from the email notification list
- Modifying recipient properties in the email notification list
- Sending a test message to a recipient in the email notification list

Configuring the Email Notification Manager:

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to specify the SMTP server address and the administrator’s ″reply to″ address. The reply to address typically is the address of the ServeRAID Manager system administrator or IT or network administrator.

1. Click **Actions** → **SMTP server settings**. The SMTP server settings window appears.
2. Type the SMTP server address. You can enter the host name, including domain, or the TCP/IP address.
3. Type the administrator’s reply to address. Include the user name and domain, such as jack_smith@syntecha.com.

4. Click **OK**.

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Deleting a recipient from the email notification list
- Modifying recipient properties in the email notification list
- Sending a test message to a recipient in the email notification list
- Monitoring messages sent from the Email Notification Manager

Adding a recipient to the Email Notification Manager: Use this action to add a recipient to the e-mail notification list. You can add **one** recipient only, typically the ServeRAID Manager administrator, or system or network administrator. The recipient is notified of all or selected events that occur on the enclosure(s) connected to the management station.

1. Click (Add e-mail recipient). The ″Add e-mail recipient″ window opens.
2. In the To e-mail address field, type the name of the recipient that you want to add to the list. Include the user name and domain, such as jack_smith@synteca.com.
3. Select the type of event(s) you want the recipient to be notified of: error, warning, or informational.
4. Click **Add**.

More information

- Using the Email Notification Manager
- Deleting a recipient from the email notification list
- Modifying recipient properties in the email notification list
- Sending a test message to a recipient in the email notification list

Configuring the Email Notification Manager: Use this action to specify the SMTP server address and the ″from″ address of the e-mail sender. The default ″from″ address includes a generalized management station name and a partial e-mail address of the form `<name@%s>`. Replace the the partial address with the address you want to appear in the From field of the e-mail message.

1. Click **Actions** → **SMTP server settings**. The ″SMTP server settings″ window opens.
2. Type the SMTP server address. You can enter the host name, including domain, or the TCP/IP address.
3. Type the from address, replacing the general management station name and e-mail address with the name and address you want to appear in the From field of the e-mail message.
4. Click **Add**.

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Deleting a recipient from the email notification list
- Modifying recipient properties in the email notification list
- Sending a test message to a recipient in the email notification list

Deleting a recipient from the Email Notification Manager: Use this action to delete a recipient from the email notification list. After deleting a recipient, the Email Notification Manager no longer notifies the user of events occurring on the enclosure(s) connected to the management station.

1. In the e-mail notification list, click the recipient you want to delete.
2. Click (Delete e-mail recipient).
3. Click **Yes** to confirm. The Email Notification Manager deletes the recipient from the e-mail notification list.

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Modifying recipient properties in the email notification list
- Sending a test message to a recipient in the email notification list

Modifying a recipient in the Email Notification Manager: Use this action to modify the e-mail recipient's address or monitored event types:

1. In the e-mail notification list, click the recipient that you want to modify.
2. Click (Modify email recipient). The "Modify email recipient" window opens.
3. Change the recipient properties in the appropriate fields.
4. Click **Modify**.

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Deleting a recipient from the email notification list
- Sending a test message to a recipient in the email notification list

Sending a test message in the Email Notification Manager: Use this action to send a test message to a recipient in the email notification list:

1. Click a recipient in the e-mail notification list.

   Note: This action is disabled if you do not click a recipient.
2. Click **Actions** → **Send test message**.

If the test fails, verify the following and try the test again:

1. You typed the correct e-mail address when you added the recipient to the Email Notification Manager.
2. You typed the correct SMTP server address when you configured the Email Notification Manager.

More information

- Using the Email Notification Manager
- Adding a recipient to the email notification list
- Deleting a recipient from the email notification list
- Modifying recipient properties in the email notification list

Using the Task Manager:

Note: You cannot use the Task Manager in bootable-CD mode or when using the ServeRAID Manager as a plug-in application.

The Task Manager is a tool that you can use to manage scheduled tasks, such as a logical drive migration or synchronizing a logical drive. For each scheduled task, the Task Manager displays the task start time, status, and description.

You can use the Task Manager to:

- Modify a task’s schedule
- Delete a scheduled task
- View task properties

The Task Manager is enabled by default. To disable the Task Manager, click **Actions** → **Disable Task Scheduler** .

Note: If you disable the Task Manager, a task may miss its start time. To run the task, you must reschedule it. See Understanding scheduled tasks for more about missed start times.

An example

You schedule a synchronization task with an initial start time of Sunday, October 13, 2003 at 2:00 AM. Due to a power outage, the task misses its start time. The Task Manager flags the error. The next day, you modify the task schedule to run on Sunday October 20, 2003 at 2:00 AM.

The Task Manager status

If a scheduled task runs successfully, the Task Manager updates the ″Status″ column in the task list with Complete. For recurring tasks, an asterisk (*) indicates that the task has completed one cycle and is scheduled to run again at the given interval. If a scheduled task does not run successfully, the Task Manager:

- Updates the ″Status″ column in the task list with Error.
- Updates the first column of the task list with a ⚠ , notifying you of a recoverable error (a warning message); it displays a ⊗ if the error is non-recoverable (a critical error).

To see details about the error, use View task properties. See Understanding scheduled tasks for more about Task Manager status, start time validation, missed start times, and error conditions.

The Task Manager interface

The Task Manager consists of the following:

Toolbar
Provides quick-path icons for common tasks.

Task list
Displays the scheduled tasks.

Task event viewer
Displays status information for the Task Manager.

More information

- Understanding scheduled tasks
- Starting the Task Manager
- Deleting a scheduled task
- Modifying a scheduled task
- Viewing task properties
- The Task Manager menubar
- The Task Manager toolbar

- Using the ServeRAID Manager agent

Task Manager menu bar: The Task Manager menu bar has the following functions:

File →

Clear event log
Click Task from the menu. Clears the current contents of the Task Manager event viewer.

Close Closes the Task Manager.

View →

Toolbar
Turns the toolbar on and off. The default is on (selected).

Actions →

Delete task

Modify task

View task

"Sending a test message in the Email Notification Manager" on page 332
Disable Task Scheduler

Turns the Task Scheduler on and off. The default is on.

Help →

Information about this window
View context-sensitive information for the current window.

Search
Searches for one or more specified words in ServeRAID Manager Assist and displays a list of topics that include the words.

Contents
Presents the ServeRAID controller Assist contents. You can use the contents to acquaint yourself with ServeRAID Manager Assist topics.

About ServeRAID Manager
Reports the ServeRAID Manager version number, copyright, and legal information.

Task Manager toolbar: The Task Manager toolbar has the following functions:

Delete task

Modify task

View task

Information about this window View context-sensitive information for the current window.

Starting the Task Manager: Use this action to start the Task Manager. You must start the Task Manager to manage scheduled tasks.

1. In the Enterprise view, click (system).
2. From the "Actions" menu, select **Agent actions** → **Configure**. The ServeRAID Manager Agent window opens.
3. Click the **Task** note tab.

More information

- Using the Task Manager
- Deleting a scheduled task
- Modifying a scheduled task
- Viewing task properties
- The Task Manager menubar
- The Task Manager toolbar
- Using the ServeRAID Manager agent

Modifying a scheduled task in the Task Manager:

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to modify a scheduled task's start time and date, description, and recurrence interval (if the task can be recurring):

1. Click the task that you want to modify in the task list.
2. Click (Modify task). The "Modify task" window opens.
3. Change the task properties in the appropriate fields.
4. Click **OK**.

More information

- Using the Task Manager
- Starting the Task Manager
- Understanding scheduled tasks
- Deleting a scheduled task
- Viewing task properties

Deleting a scheduled task from the Task Manager:

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to delete a scheduled task from the task list. After deleting a task, the Task Manager no longer attempts to run the task at the scheduled time.

1. Click the task you want to delete from the task list.
2. Click (Delete task).
3. Click **Yes** to confirm. The Task Manager deletes the task from the task list.

More information

- Using the Task Manager
- Starting the Task Manager
- Understanding scheduled tasks

- Modifying a scheduled task
- Viewing task properties

Viewing task properties from the Task Manager: You can view details about each task in the task list.

In the task list, double-click the **task** in which you are interested. The "Task properties" window opens. This window contains information about the task, including the following: description, task ID, start time, recurrence interval, and last time a recurring task started. An error description appears if the Task Manager flagged an error for the task.

More information

- Using the Task Manager
- Starting the Task Manager
- Understanding scheduled tasks
- Deleting a scheduled task
- Modifying a scheduled task

Monitoring storage enclosures on a network: You can monitor network-attached storage enclosures from a ServeRAID Manager console running on the local system. The ServeRAID Manager management station software must be installed and running on the system to which the storage enclosure is attached. Then, you can add a management station to your ServeRAID Manager Enterprise view.

You can configure the management station to monitor and manage network-attached storage:

- Use the Security Manager to allow users to configure and view enclosure information from remote systems.
- Use the Email Notification Manager to notify users by email when events occur on the enclosures attached to the management station. You can send email notifications to any user on the network; the recipient does not have to be running ServeRAID Manager.
- Use the SNMP Trap Manager to notify other systems running an SNMP console of all SNMP traps that occur on the enclosures attached to the management station.

More information

- Adding a management station
- Using the management station agent
- Adding a management station agent

Using the Security Manager: Use the Security Manager to give users the ability to connect to a management station from remote systems and configure and view the attached enclosures. You must configure at least one user name and password before any user can connect to the management station.

Note: You cannot disable security for a management station.

Use the following actions in the Security Manager to manage users' ability to configure and view network storage from remote systems:

- Add a user
- Delete a user
- Modify a user

The first time you open the Security Manager, the ServeRAID Manager prompts you to type an Admin (administrator) password. Either type a password or click **Cancel** to not have an Admin user.

The Security Manager interface

The Security Manager consists of the following:

Toolbar
Provides quick-path icons for common tasks.

Security list
Displays the users with access to this management station.

The Security Manager logs events to a logging file, RAIDSEC.LOG. If this file exceeds 200 KB, the ServeRAID Manager copies the file to RAIDSEC.OLD and creates a new RAIDSEC.LOG. If there is a RAIDSEC.OLD already, the ServeRAID Manager overwrites it.

More information

- The Security Manager menubar
- The Security Manager toolbar
- Using the management station agent

Using the SNMP Trap Manager: Use the SNMP Trap Manager to notify other systems running an SNMP console of all SNMP traps that occur on the enclosures connected to a management station. You add system names to the SNMP traps list. Each system in the list is notified of all traps that occur on the enclosure.

You can use the SNMP Trap Manager to:

- Add a system to the SNMP traps list
- Delete a system from the SNMP traps list
- Modify system properties in the SNMP traps list
- Send a test trap to a system in the SNMP traps list

The SNMP Trap Manager is enabled by default. To disable the SNMP Trap Manager, click **Actions** → **Disable SNMP Traps** . If you disable the SNMP Trap Manager, the SNMP traps are generated, but not dispatched to remote systems.

An example

You install SystemA in a lab with a network-attached storage enclosure. You configure SystemA to send SNMP traps from the enclosure, but you want to monitor the traps from your workstation, not SystemA. You open the SystemA SNMP Trap Manager from the ServeRAID Manager and define your workstation in the SNMP traps list. When running the SNMP console from your workstation, you are notified of all SNMP traps that occur on the enclosures attached to SystemA.

The SNMP Trap Manager and traps

Traps generated by SNMP include the following:

- Progress information, such as rebuilds, synchronizing, and migrations.
- Problems, such as defunct physical drives and PFA errors.
- Changes to the local configuration, such as adding or replacing a controller.

When an SNMP trap is generated on an enclosure, the SNMP Trap Manager connects with each system in the SNMP traps list and relays the trap to these systems' SNMP consoles.

The SNMP Trap Manager also logs each trap to a logging file, RaidSNMP.log. If this file exceeds 200 KB, the Manager copies the file to RaidSNMP.old and creates a new RaidSNMP.log. If there is a RaidSNMP.old already, the Manager overwrites it.

The SNMP Trap Manager interface

The SNMP Trap Manager consists of the following:

Toolbar
Provides quick-path icons for common tasks.

SNMP traps destination list
Displays the remote systems configured to receive SNMP traps.

More information

- The SNMP Trap Manager menubar
- The SNMP Trap Manager toolbar
- Using the management station agent

Using the Email Notification Manager: Use the Email Notification Manager to notify users by e-mail of events that occur on enclosures connected to a management station. You type the recipient's name and e-mail address in the e-mail notification list. The recipient is notified of all or selected events that occur on the enclosure.

You can use the Email Notification Manager to:

- Add a recipient to the email notification list
- Delete a recipient from the email notification list
- Modify recipient properties in the email notification list
- Send a test message to a recipient in the email notification list
- Configure the Email Notification Manager

The Email Notification Manager is enabled by default. To disable the Email Notification Manager, click **Actions** → **Disable Email Notifications**. If you disable the Email Notification Manager, events are generated but users do not receive e-mail notification when the events occur.

An example

You install SystemA in a lab with a network-attached storage enclosure. You add a management station agent on SystemA to monitor for events and problems. You want to monitor SystemA from your workstation but do not have ServeRAID Manager installed. You open the SystemA Email Notification Manager from the ServeRAID Manager and add your name and e-mail address to the e-mail notification list. You are notified by e-mail of all problems and events that occur on the enclosure attached to SystemA.

The Email Notification Manager and its events

Events monitored by the Email Notification Manager include the following:

- Progress information, such as rebuilds, synchronizing, and migrations.
- Problems, such as defunct physical drives and PFA errors.

- Changes to the local configuration, such as creating a hot-spare drive or defining a logical drive.

Each event has an **event type** indicating its severity: error, warning, or informational. You can configure a recipient to be notified of errors, errors and warnings, or errors, warnings, and informational events.

The Email Notification Manager logs events to a logging file, RAIDSMTP.LOG. If this file exceeds 200 KB, the ServeRAID Manager copies the file to RAIDSMTP.OLD and creates a new RAIDSMTP.LOG. If there is a RAIDSMTP.OLD already, the ServeRAID Manager overwrites it. The Email Notification Manager appends SMTP transport errors in a separate log file, SMTPERR.LOG.

The Email Notification Manager interface

The Email Notification Manager consists of the following:

Toolbar
Provides quick-path icons for common tasks.

Email notification list
Displays the recipients configured to receive e-mail notifications.

More information

- The Email Notification Manager menubar
- The Email Notification Manager toolbar
- The Email Notification Manager message format
- Using the management station agent

Modifying logical drives using the Configuration wizard

Understanding logical-drive migration: Logical-drive migration is a powerful and flexible feature of ServeRAID controllers. You can use logical-drive migration to:
- Change the RAID level of existing logical drives
- Increase array free space, so you can create additional logical drives
- Increase the size of an existing logical drive

You can migrate a logical drive while the server is running with only minor performance degradation for users during the process. You can start the migration immediately or schedule it for a later date and time.

HostRAID controllers do **not** support logical drive migration.

Note: To perform a logical-drive migration, the following must be true:
- For ServeRAID SCSI controllers, no more than seven logical drives currently exist.

  During logical-drive migration, the ServeRAID Manager creates one logical drive for temporary use. (You must have an extra logical drive onto which the data can migrate.) The ServeRAID Manager sets this logical drive to the system state. When migration is complete, this logical drive is removed.
- The logical drive that you have selected to migrate must be in the okay state.

- For ServeRAID SCSI controllers, the logical drive cannot be configured with RAID level-00, 10, 1E0, 50, or 5EE.

If a physical drive fails during a logical-drive migration and you are migrating between RAID levels other than RAID level-0, the migration will continue and complete. Then, you must replace and rebuild the failed physical drive.

Logical-drive migrations also can recover from a power failure. If power is lost to the server during a logical-drive migration, the migration will restart as soon as power is restored, and complete without data corruption.

Note: (Cluster and failover configurations only) A logical drive will not failover while undergoing logical-drive migration.

More information

- Things to consider when changing the RAID level
- Example: Increasing free space
- Example: Increasing logical drive size
- Understanding scheduled tasks
- Using the Task Manager

Example: Increasing free space: Create an array with three physical drives. Then, create two logical drives (labeled 1 and 2) in the array such that you have no available free space. The logical drive stripes look similar to the following.

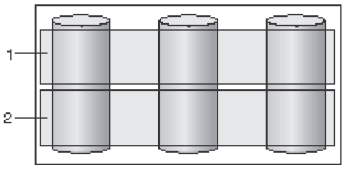

Later, you decide to create another logical drive in the array. First, you must create some free space by adding more physical drives to the array.

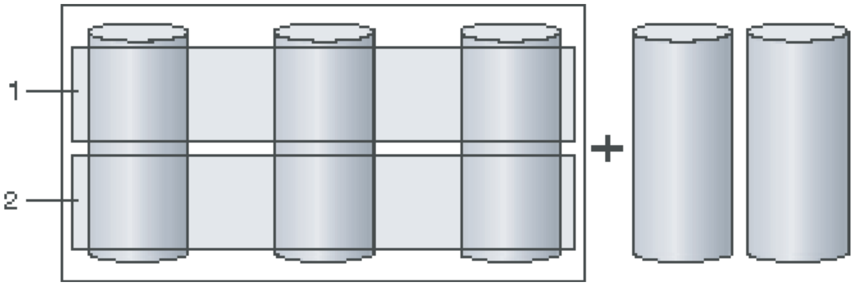

If you add two more physical drives to the array, the logical drives migrate such that they are striped across all five drives. The ServeRAID Manager retains the size of the logical drive by decreasing the size of each block on a physical drive. The space in the logical drives does not increase, so there is free space across the bottom portion of each physical drive (labeled 3).

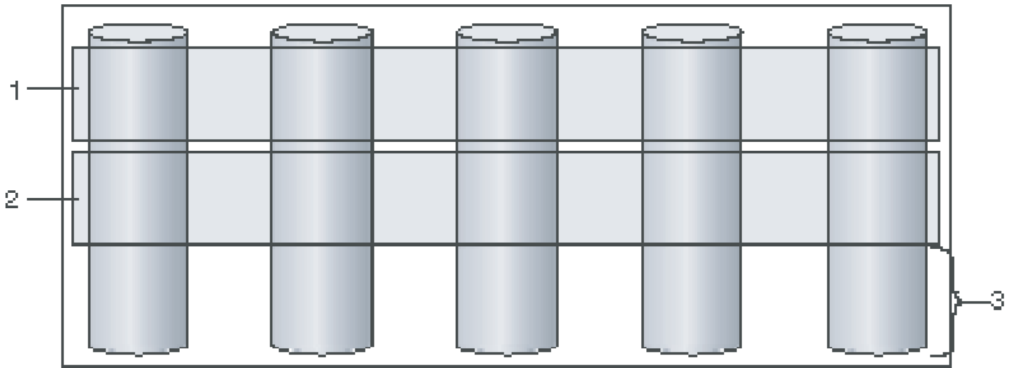

When the operation is complete, use ″Create logical drive″ to create a new logical drive in the free space.

When increasing free space, you can add up to three physical drives, as long as the sum of the physical drives does not exceed the maximum number of physical drives allowed for the array. The maximum number of physical drives you can have in an array is limited by the controller stripe-unit size.

Example: Increasing the logical drive size:

Create an array with three physical drives. Then, create two logical drives (labeled 1 and 2) in the array such that you have no available free space. The logical drive stripes look similar to the following.

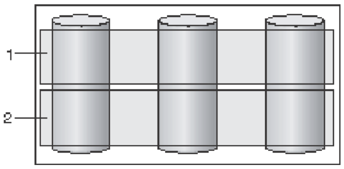

Later, you decide you need additional space in the logical drives in the array. First, you must add more physical drives to the array.

Note: Blocks are labeled as **x** and * in these examples. All blocks labeled **x** are the same size in logical drive 1. All blocks labeled * are the same size in logical drive 2.

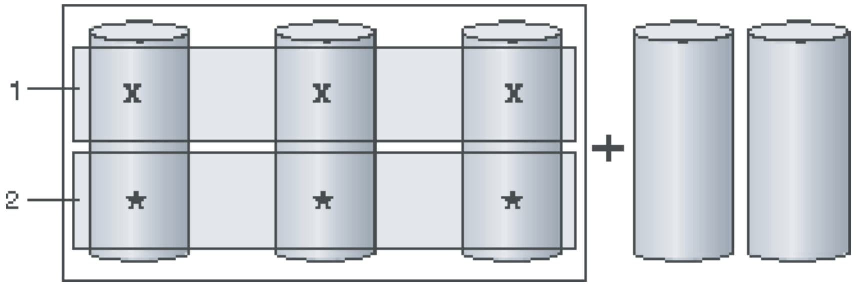

If you add two additional physical drives to the array, the logical drives migrate such that they are striped across all five drives. The ServeRAID Manager expands the size of each logical drive in proportion to the original configuration. If there originally was free space in the array, the free space is increased also.

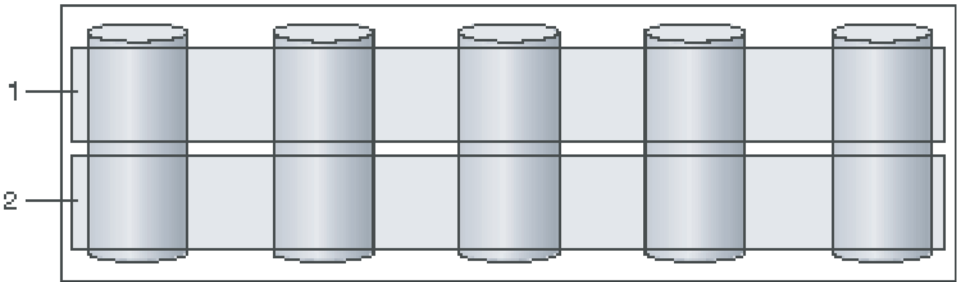

When increasing logical drive space, you can add up to three physical drives, as long as the sum of the physical drives does not exceed the maximum physical drives allowed for the array. The maximum number of physical drives you can have in an array is limited by the controller stripe-unit size.

SR-8i SR-7t For the ServeRAID-8i controller and ServeRAID-7t controller, the ServeRAID Manager does not enforce the three physical drive limit when increasing logical drive size.

More information

- Understanding stripe-unit size
- Understanding creating basic arrays
- Creating logical drives (action)
- Increasing logical drive size (action)
- Increasing logical drive size (SAS, SATA) (action)

Understanding scheduled tasks: The ServeRAID Manager Task Scheduler allows you to run lengthy operations at a convenient time, such as a logical drive migration or synchronizing a logical drive. All schedulable tasks work the same way: To schedule a migration, for example, you perform the standard set of steps; then, in the last step, you can choose to run the task immediately or to schedule it for another time.

You can run the task at the given date and time or make it a recurring task. If the task is recurring, it runs on a regular basis according to the interval you choose: the same time each day, the same time each week, or the same time each month. Not all tasks can be recurring tasks. For example, a logical drive migration cannot be a recurring task.

The Task Manager is a tool that you can use to manage scheduled tasks. It allows you to:

- Modify a task's schedule
- Delete a scheduled task
- View task properties

Start time validation

When you schedule a task, it executes on the agent machine which may be a remote machine in a different time zone. To prevent scheduling problems due to time differences between the agent and client machines, the Task Scheduler performs the following checks:

- To accommodate differences in times zones, you can schedule a task up to 23 hours in the past, pending approval by the agent.

- If the agent determines that the task is scheduled in the past (based on the local time on the agent machine), it will reject the task and you will be prompted to select a different date and time.

Missed start times

If a non-recurring task misses its start time, the Task Manager flags it with an error. You must modify the task schedule if you want to run it at another time.

If a recurring task misses its start time, it is rescheduled to run at the next scheduled interval. Example: a recurring task runs every Monday at 1:00 AM. The agent machine is powered off at that time. When you restart the agent machine, the missed task is rescheduled to run the following Monday at 1:00 AM.

Note: To accommodate temporary or brief interruptions on the agent machine, each task has a 30-minute grace period built into its start time. Example: a task is scheduled to run at 10:00 PM. A power outage lasts 20 minutes, from 9:55 PM to 10:15 PM. The task runs at 10:15 PM.

Error conditions

When a scheduled task does not run successfully, the Task Manager displays an error icon and updates the ″Status″ column with an Error. If a task misses its start time, it is flagged with a warning error condition. You must reschedule the task to clear the error. If the error is non-recoverable it is considered critical. You must delete the task from the Task Manager; to reschedule it, you must create a new task. To see details about the error, use View task properties; examine the Error Description field for a description of what went wrong and what caused the task to fail.

More information

- Synchronize logical drives (action)
- Understanding logical drive migration
- Using the Task Manager
- Starting the Task Manager
- Deleting a scheduled task
- Modifying a scheduled task
- Viewing task properties
- Using the ServeRAID Manager agent

Extending a partition on a logical drive: (Microsoft Windows 2000, Windows XP, Windows Server 2003, and Windows NT 4.0 only)

If you have increased the size of a logical drive, you can extend the partition on that logical drive to use the newly added space.

Notes:

- To complete this procedure, you must be logged on as one of the following:
  - An administrator
  - A member of the Administrators group
- If your server is connected to a network, network policy settings might prevent you from completing this procedure.

Windows 2000, Windows XP, Windows Server 2003

1. After completing ″Increase logical drive space,″ shut down and restart the server.
2. In the lower-left corner of the desktop, click **Start** → **Programs** → **Administrative Tools** → **Computer Management**. The Computer Management window opens.
3. Right-mouse click the volume and click **Extend Volume**.
4. Follow the Extend Volume wizard instructions.

Notes:

- You **cannot** extend the following types of volumes:
  - Striped
  - Mirrored
  - RAID level-5
  - FAT or FAT32 format
  - System
  - Startup (boot)
  - Simple or extended volumes that were upgraded from basic to dynamic.
- You can extend a volume only if the volume is formatted using NTFS **or** does not contain a file system.
- You can extend a simple or extended volume only if the volume was created as a dynamic volume.
- You can extend a simple volume if the extension will occur within its original physical drive **or** onto additional physical drives.
- After a volume is extended onto multiple drives (spanned), it cannot be mirrored or striped.
- You can extend simple and spanned volumes on dynamic disks onto a maximum of 32 dynamic disks.
- After a spanned volume is extended, you cannot delete any portion of the volume without deleting the entire spanned volume.

Windows NT 4.0

1. After completing ″Increase logical drive space,″ shut down and restart the server.
2. In the lower-left corner of the desktop, click **Start** → **Programs** → **Administrative Tools (Common)** → **Disk Administrator**. The Disk Administrator window opens.
3. Hold the Ctrl key down and right-click both the old volume and the new free space on the logical drive; then, click **Extend Volume Set**.
4. Type the total size you want for the partition.
5. Click **OK**.
6. Right-mouse click the volume and click **Commit Changes Now**. A confirmation window opens.
7. Click **Yes**. Another confirmation window opens.
8. Click **Yes**.

Notes:

- You can extend a volume only if it is formatted using NTFS. You must convert FAT volume sets to NTFS before you can extend them.

- You can extend volume sets onto a maximum of 32 physical drives.
- You cannot extend the volume where the Windows NT 4.0 system files reside.
- After a volume is extended, you cannot delete any portion of the volume without deleting the entire volume.

More information

- Increasing logical drive size (action)
- Increasing logical drive size (SAS, SATA) (action)

Configuration wizard notes and attentions: While you are creating the configuration, the ServeRAID Manager reports notes and attentions regarding the configuration in the event viewer. These events can help you create a better or more optimized configuration.

Array storage space still available.

Explanation: You can create another logical drive in this array or make a logical drive in this array larger.

Physical drives contain unusable capacity.

Explanation: You could redefine your arrays with the same size drive capacities to avoid wasting space.

There is(are) *n* ready drive(s) still available. (Where *n* is a number)

Explanation: You can create another array, add the remaining ready drives to one of the new arrays, or create a hot-spare drive.

This hot-spare drive will not work for any defined array.

Explanation: This hot-spare drive is not large enough to replace a failed drive in any array, or there are no redundant logical drives.

Modifying logical drives on ServeRAID SCSI controllers (ServeRAID series 3, 4, 5, 6, and 7K):

Changing the RAID level:

Note: This action is not supported on the integrated RAID controller.

You can change the RAID levels of currently defined logical drives. To use ″Change RAID level,″ all of the logical drives within the array **must** be the same RAID level.

Note:

1. If the logical drive is RAID level-00, 10, 1E0, 50, or 5EE, you cannot change the logical drive to another RAID level. If the logical drive is RAID level-0, 1, 1E, 5, or 5E, you cannot change the logical drive to RAID level-00, 10, 1E0, or 50.
2. You can schedule the logical-drive migration for a later date and time when you review the configuration changes.

   3. You cannot change the RAID level while the logical drive is being synchronized.
   4. (Cluster and failover configurations only) A logical drive will not failover while undergoing logical-drive migration.

1. In the Physical devices view, click ▤ (array).
2. Right-click **Logical-drive migration** and then the appropriate **Change RAID level**.

   Note: If the array already contains the maximum number of physical drives (based on the stripe-unit size), the following actions are not available:
   - **Change RAID level from RAID 0 to RAID 5**
   - **Change RAID level from RAID 1 to RAID 5**
3. If you click one of the following:
   - **Change RAID level from RAID 0 to RAID 5**
   - **Change RAID level from RAID 1 to RAID 5**

   The Configuration wizard opens with the ″ Create arrays″ window. Add the physical drives to the array; then, review and apply your new configuration.

   If you click **Change RAID level from RAID 5 to RAID 0**, the Configuration wizard automatically removes the physical drives and opens the ″ Configuration summary″ window. Review and apply the new configuration.

   If you click **Change RAID level from RAID 5E to RAID 5**, the Configuration wizard does **not** add or remove a physical drive. Instead, the wizard changes the spare space to become free space. Use this free space to create a new logical drive.

More information

- Things to consider when changing RAID levels
- Understanding scheduled tasks

Increasing free space:

Note: This action is not supported on the integrated RAID controller.

Using this action you can add one, two, or three physical drives to an existing array so that you can create another logical drive in the array.

When you create an array, logical drives are striped across all physical drives in that array. To create more free space in an existing array, define additional physical drives to add to the array. Then, the ServeRAID Manager migrates the logical drives such that the data is spread across the existing *and* new physical drives.

Note:

   1. To increase free space in an array, you need at least one ready drive that is *at least* as big as the largest drive in the array.
   2. You can schedule the logical-drive migration for a later date and time when you review the configuration changes.
   3. (Cluster and failover configurations only) A logical drive will not failover while undergoing logical-drive migration.

1. In the Logical devices view, click ▤ (array).
2. Right-click **Logical-drive migration** → **Increase free space**. The Configuration wizard opens with the ″Create arrays″ window.

3. Add the physical drives to the array. When complete, click **Next**. The ″Configuration summary″ window opens.
4. Review and apply your new configuration.
5. After the logical-drive migration is complete, create a new logical drive in the new free space.

More information

- Example: Increasing free space
- Understanding logical-drive migration
- Understanding scheduled tasks

Increasing logical-drive size:

Note: This action is not supported on the integrated RAID controller.

Using this action you can add one, two, or three physical drives to an existing array so you can expand the size of the logical drives in the array. When increasing logical drive space, the ServeRAID Manager migrates the logical drives such that the logical drives gain additional space, much like adding paper to a notebook.

When you create an array, logical drives are striped across all of the physical drives in that array. To increase the size of the logical drives in an existing array, define additional physical drives to add to the array. Then, the ServeRAID Manager migrates the logical drives such that the data is spread across the existing **and** new physical drives.

Note:

1. To increase the logical drive size, you need at least one ready drive that is **at least** as big as the largest drive in the array.
2. You can schedule the logical-drive migration for a later date and time when you review the configuration changes.
3. (Cluster and failover configurations only) A logical drive will not failover while undergoing logical-drive migration.

1. In the Logical devices view, click (array).
2. Right-click **Logical-drive migration** → **Increase logical-drive size**. The Configuration wizard opens with the ″Create arrays″ window.

   Note: This option is not available if the logical drive exceeds 2 terabytes using any available physical drive.
3. Add the physical drives to the array. When complete, click **Next**. The ″Configuration summary″ window opens.
4. Review and apply your new configuration.
5. After the logical-drive migration is complete, create additional partitions to use the newly added logical-drive space. If the server is a Microsoft Windows system, you can extend the size of a pre-existing partition with the newly acquired logical-drive space.

More information

- Example: Increasing logical drive size
- Extending a partition on a logical drive
- Understanding scheduled tasks

Adding physical drives in an existing configuration: You can add physical drives (also called logical-drive migration or LDM) to dynamically modify an existing array and its logical drives.

Note: You cannot add more than three physical drives. The maximum number of physical drives you can have in an array is still limited by the controller stripe-unit size.

1. Click the appropriate **Array** tab in the right pane:

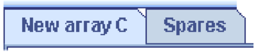

Then, from the list of ready drives, select the drives you want to move to the array:	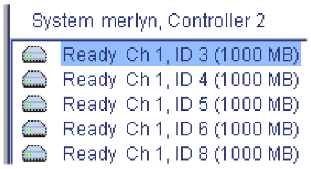 System merlyn, Controller 2 Ready Ch 1, ID 3 (1000 MB) Ready Ch 1, ID 4 (1000 MB) Ready Ch 1, ID 5 (1000 MB) Ready Ch 1, ID 6 (1000 MB) Ready Ch 1, ID 8 (1000 MB)

2. Click

   >> (Add selected drives) to add the drives to the array. You can click

   >> (Add all drives) to move **all** ready drives to an array.

   Notes:

   a. The list includes only physical drives that are **at least** the size of the smallest physical drive in the selected array. When adding to an existing array, you cannot add physical drives that are smaller than the smallest drive currently in the array.
   b. When adding to an existing array, if you add physical drives that are larger than the physical drives currently in the array, you will not be able to use all the space on the new drives.
   c. You cannot add a physical drive to the array if the resulting logical drive size will exceed 2 terabytes.
   d. Existing physical drives in the selected array are labeled **Online**. Newly added physical drives are labeled **New online**. You cannot remove online drives from an existing array.
3. After you add the physical drives to your arrays, click **Next**. The ″ Configuration summary″ window opens.

To leave the Configuration wizard, click **Cancel**.

More information

- Understanding logical-drive migration
- Things to consider when changing RAID levels
- Understanding scheduled tasks
- Changing the RAID level of each logical drive in an array (action)
- Increasing the array free space so you can create another logical drive in the array (action)
- Increasing the logical drive space of each logical drive in the array (action)

- Example: Increasing free space
- Example: Increasing logical drive size

Confirming your logical-drive migration configuration: Use the configuration summary to review all the changes that the ServeRAID Manager will apply to your modified logical drive configuration. You can start the migration immediately or schedule it for another time.

1. Review the information that is displayed in the ″Configuration summary″ window. It describes how the new configuration will affect the logical drives and free space. To change the configuration, click **Back**.
2. To start migration now, click **Apply**. Click **Yes** when asked if you want to apply the new configuration.

   Note: Logical-drive migration is a lengthy process. You cannot perform any other actions on the affected controller until the migration is complete. The ServeRAID Manager displays a progress indicator in the status bar while the operation is in progress. When the migration is complete, the configuration is saved in the ServeRAID controller and in the physical drives.
3. To schedule the migration for a later date and time, click **Schedule**. The ServeRAID Manager displays the Scheduler pane:
4.
   - From the calendar display, select the day, month, year, and time you want the task performed.
   - Click **Apply**.
   - Click **Yes** when asked if you want to schedule the migration.

More information

- Configuration wizard notes and attentions
- Understanding logical-drive migration
- Understanding scheduled tasks
- Using the Task Manager

Modifying logical drives on ServeRAID SAS and SATA controllers (ServeRAID-7t, ServeRAID-8i):

Changing the RAID level: You can change the RAID levels of currently defined logical drives.

1. In the Logical devices view, click 🛢 (logical drive).
2. Right-click **Expand or change logical drive**. The Configuration wizard opens with the ″Choose RAID level″ window.
3. Choose a new RAID level for the logical drive; then, click **Next**. The ″Create logical drives″ window opens.

   Note: Only valid RAID level migrations appear in the list. Migration requirements for each RAID level are described here.
4. Add the physical drives needed to support the new RAID level (if any). Click **Advanced settings** to set the logical drive size; optionally, select a new stripe-unit size. Then, click **Next**. The ″Configuration summary″ window opens.

   Note: When you add drives, the migration wizard shows the current layout of the data, including the current logical drive size. For most migrations, you will need to increase the logical drive size setting; the default size is not generally adequate once you add drives or segments.

5. Review and apply the new configuration.

More information

- Things to consider when changing RAID levels
- Increasing logical-drive size (SAS, SATA) (action)
- Understanding scheduled tasks

Increasing logical-drive size: Use this action to add space to an existing logical drive so you can increase its size.

When you create a logical drive, it is striped across all of the physical drives used by the logical drive. To increase the size of an existing logical drive, you add segments to the logical drive, which can be on the same or different disks. Then, the ServeRAID Manager migrates the logical drive such that the data is spread across the existing **and** new physical drives.

Note:

1. The capacity of the modified logical drive must match or exceed its current capacity.
2. The maximum size of a logical drive is 2 terabytes.

1. In the Logical devices view, click (logical drive).
2. Right-click **Expand or change logical drive**. The Configuration wizard opens with the ″Choose RAID level″ window. Click **Next**. The ″Create logical drives″ window opens.
3. Add segments to the logical drive. Optionally, click **Advanced** to select a new stripe-unit size for the logical drive.

   Note: To remove or replace a physical drive, click the drive you want to remove (indicated by ✖); then, click the drive you want to replace it with. To cancel your changes and start over, click Revert logical drive.

   When you are ready to continue, click **Next**. The ″Configuration summary″ window opens.
4. Review and apply your new configuration or schedule the migration for another time.
5. After the logical-drive migration is complete, create additional partitions to use the expanded logical-drive space. If the server is a Microsoft Windows system, you can extend the size of a pre-existing partition with the newly acquired logical-drive space.

More information

- Example: Increasing logical drive size
- Extending a partition on a logical drive
- Changing the RAID level (SAS, SATA) (action)
- Understanding scheduled tasks

Confirming your logical-drive migration configuration: Use the configuration summary to review all the changes that the ServeRAID Manager will apply to your modified logical drive configuration. You can start the migration immediately or schedule it for another time.

1.

Review the information that is displayed in the ″Configuration summary″ window. It describes how the new configuration will affect the logical drives and free space. To change the configuration, click **Back**.

2. To start the migration now, click **Apply**. Click **Yes** when asked if you want to apply the new configuration.

   Note: Logical-drive migration is a lengthy process. The ServeRAID Manager displays

   (in animation) in the Logical devices view while the operation is in progress. When the migration is complete, the configuration is saved in the ServeRAID controller and in the physical drives.

3. To schedule the migration for a later date and time, click **Schedule**. The ServeRAID Manager displays the Scheduler pane:
   - From the calendar display, select the day, month, year, and time you want the task performed.
   - Click **Apply**.
   - Click **Yes** when asked if you want to schedule the migration.

More information

- Configuration wizard notes and attentions
- Understanding logical-drive migration
- Understanding scheduled tasks
- Using the Task Manager

Modifying logical drives on ServeRAID enclosures:

Changing the RAID level: You can change the RAID level of currently defined logical drives in an array.

1. In the Logical devices view, click (array).
2. Right-click **Expand or migrate array**. The Configuration wizard opens with the ″Choose RAID level″ window.
3. Choose a new RAID level for the array; then, click **Next**. The ″Modify array″ window opens.

   Note: Only valid RAID level migrations appear in the list. Migration requirements for each RAID level are described here.
4. Add the physical drives needed to support the new RAID level (if any). Optionally, click **Advanced settings** to set the drive capacity display units (MB, GB, TB). Then, click **Next**. The ″Configuration summary″ window opens.
5. Review and apply the new configuration.

More information

- Things to consider when changing RAID levels
- Increasing logical-drive capacity (enclosure)

Increasing logical drive capacity: Use this action to increase the size of logical drives in an enclosure. To increase the size of one or more drives in an enclosure, use the Configuration wizard. To increase the size of a specific logical drive, use the Increase logical drive capacity action on the logical drive object.

Increasing logical drive size with the Configuration wizard

1. In the Enterprise view, click (enclosure).
2. Right-click **Configure storage**. The Configuration wizard opens with the ″ Configuration path″ window.
3. Click **Advanced options**; then, select **Increase logical drive capacity**.

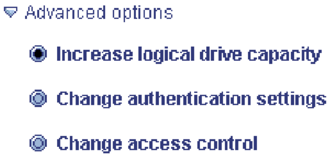

▽ Advanced options

● Increase logical drive capacity

○ Change authentication settings

○ Change access control

4. Click **Next**. The ″Increase logical drive capacity″ window opens.
5. Click the appropriate array tab.

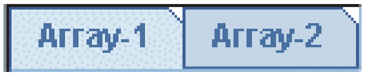

6. For each logical drive in the array, type amount space you want to add to the logical drive in the **Additional (MB)** field.

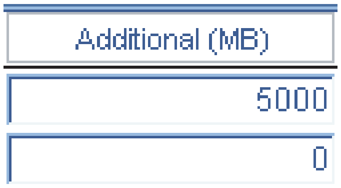

Additional (MB)

5000

0

7. Repeat steps 5 and 6 to increase logical drive capacity for other arrays.
8. When you are ready to continue, click **Next**.
9. Review the configuration summary and apply the changes.

Increasing the size of a single logical drive

1. In the Logical devices view, click (logical drive).
2. Right-click **Increase logical drive capacity**. The Configuration wizard opens with the ″Increase logical drive capacity″ window.
3. Type amount space you want to add to the logical drive in the **Additional (MB)** field.
4. When you are ready to continue, click **Next**.
5. Review the configuration summary and apply the changes.

More information

- Changing authentication settings
- Changing access control settings

Adding drives to an existing array: Use this action to add physical drives to an existing array without changing the RAID level.

When you create an array, it is striped across all the physical drives used by the array. To increase the size of an existing array, you add segments to the array, which can be on the same or different disks. Then, the ServeRAID Manager migrates the array such that the data is spread across the existing **and** new physical drives.

Note:

1. The capacity of the modified array must match or exceed its current capacity.
2. The maximum size of an array is 2 terabytes.

   3. This action is not supported for all RAID levels.

1. In the Logical devices view, click (array).
2. Right-click **Expand or migrate array**. The Configuration wizard opens with the ″Choose RAID level″ window.
3. Click **Next**. (Do **not** change the RAID level!) The ″Modify array″ window opens.
4. Add segments to the array. Optionally, click **Advanced settings** to set the drive capacity display units (MB, GB, TB).
5. When you are ready to continue, click **Next**. The ″Configuration summary″ window opens.
6. Review and apply your new configuration.

More information

- Extending a partition on a logical drive

Changing the authentication settings:

Note: The ServeRAID Manager supports authentication for iSCSI initiators only.

Use this action to modify the authentication method, user list, and advanced settings for each logical drive in the enclosure.

1. In the Logical devices view, click (logical drive).
2. Right-click **Configure target information**. The Configuration wizard opens with the ″Define authentication method″ window.
3. Update the authentication settings for each logical drive.
4. If you enabled authentication (by choosing CHAP or SRP), click Add User; then, add or remove users in the user list.
5. When you are ready to continue, click **Next**.
6. Review the configuration summary and apply the changes.

More information

- Adding users to the user list
- Changing access control settings

Defining the authentication method:

Note: The ServeRAID Manager supports authentication for iSCSI initiators only.

Use the Authentication information window to define the authentication method (if any) for each user permitted to use the logical drives in the array. Optionally, you can enable Radius authentication (an external authentication service) and define the SLP (Service Location Protocol) Scope name.

1. Select a logical drive from the list on the left.
2. From the **Authentication type** list, select the method used to authenticate users of this logical drive. You can choose:
   - **None** - Do not authenticate users
   - **CHAP** - Challenge Handshake Authentication Protocol
   - **SRP** - Secure Remote Password

3. If you enabled authentication (by choosing CHAP or SRP), click Add User ; the ″Global user name and password management″ window opens. Then, add users to the user list.
4. To enable Radius authentication for the CHAP authentication method, click **Radius authentication**; then, enter the following:
   - In the Primary server field, enter the host name or TCP/IP address of the authentication service; then, in the Port field, enter the server's startup port.
   - In the Secondary server field, enter the host name or TCP/IP address of the authentication service; then, in the Port field, enter the server's startup port.

   Note: Radius authentication is supported by the CHAP authentication method only.
5. To modify the default SLP Scope name, click **Advanced settings**; then, in the SLP Scope Name field, type a unique scope name or choose an existing name from the drop-down list.
6. Repeat steps 1 through 5 to enable authentication for additional logical drives.

   To replicate the settings for all logical drives, click Replicate settings .
7. When you are ready to continue, click **Next**. The ″ Define initiators″ window opens.

To return to the ″ Create logical drives″ window, click **Back**. To leave the Configuration wizard, click **Cancel**.

More information

- Adding users to the user list

Changing the rebuild rate:

Note:

1. This action is supported in bootable-CD mode only.
2. This action is not supported on the integrated RAID controller.

Use this action to change the rebuild rate. The rebuild rate determines the rate that the data from a failed physical drive is rebuilt to a new or hot-spare drive.

1. In the Enterprise view, click (controller).
2. Right-click **Change rebuild rate** and then click a rebuild rate. The following choices are available:
   - High
   - Medium
   - Low

When the rebuild rate is High, the rebuild I/O request gets high priority in the controller execution order.

If you change the rebuild rate from High to Medium in a heavily loaded system, the rebuild time can increase but provide better system performance.

If you change the rebuild rate from High or Medium to Low in a moderate to heavily loaded system, the rebuild time can increase, but provide better system performance.

Changing access control settings: Use this action to add, modify, and delete initiators that are permitted to access an enclosure. You can also change the logical drive assignments for an initiator.

1. In the Logical devices view, click (logical drive).
2. Right-click **Configure access control list**. The Configuration wizard opens with the ″Define initiators″ window.
3. Add, delete, or modify an initiator. When you are ready to continue, click **Next**.
4. Update logical drive assignments; then, click **Next**.
5. Review the Configuration summary and apply the changes.

Adding an initiator

To add an initiator:

1. Click Add. Depending on the initiator type, the ″Add iSCSI initiator″ window opens or the ″Add Port Name″ window opens.
2. Define the initiator alias, IQN or port name, and optional advanced settings.

Deleting an initiator

To delete an initiator:

1. Select an initiator from the list: simply point to the initiator and click left.
2. Click Delete.

Modifying an initiator

To modify an initiator:

1. Select an initiator from the list.
2. Click Modify. The ″Modify initiator″ window opens.
3. Type a new initiator alias.

   Note: You cannot change the initiator IQN or port name. The ServeRAID Manager updates the IQN and port name from the initiator alias.
4. (iSCSI initiators only) Optionally, change the advanced settings.
5. Click **OK**.

More information

- Modifying the authentication method

Defining the user list: You can define a user list for each logical drive. You must add at least one user to the list if you enabled authentication for the logical drive.

1. In the **User name** field, enter a user name.
2. In the **Password** field, enter the user’s password.
3. In the **Confirm password** field, enter the password again.
4. Click **Add**.
5. Repeat steps 1-4 to add additional users.

   Note: To remove a user from the list, select the user; then, click **Delete**.
6. When you are finished adding and deleting users, click **Cancel**.

Tuning your system for optimal performance

Fine-tuning your system: When fine-tuning your controller settings for optimal performance, consider the applications that you intend to run on your server. Controller settings are usually sensitive to the **types** of applications running on the server, not the server workload or the number of users using the server. You must investigate your application's design, especially its input/output behavior, to make your decision.

When configuring your controller, your server application environment can influence the following:

- The RAID level you select for your server
- The controller stripe-unit size setting
- The write-back cache mode setting

More information

- Selecting a RAID level
- Application environment groups reference
- Understanding write-cache mode for logical drives
- Changing the write-cache mode on a logical drive (action)
- Understanding write-cache mode for physical drives
- Changing the write-cache mode on a physical drive (action)
- Understanding stripe-unit size
- Changing the stripe-unit size (action)
- Understanding adaptive read-ahead cache mode
- Understanding enabled and disabled read-ahead cache mode
- Enabling and disabling read-ahead cache mode (action)

Fine-tuning the stripe-unit size: A new controller stripe-unit size is set at the factory to 8 KB. If you need to change this setting, you **must** change the stripe-unit size before you store data in the logical drives. After you store data in the logical drives, you cannot change the stripe-unit size without destroying data in the logical drives.

SR-8i SR-7t HostRAID The ServeRAID-8i, ServeRAID-7t, and HostRAID controllers do not support 8 KB stripe-unit size.

Note: You must use the ServeRAID Manager in bootable-CD mode to change the stripe-unit size setting.

Environment	stripe-unit size
Groupware (such as Lotus Notes or Exchange)	16 KB
Transaction processing database	16 KB
Decision support database	16 KB
Thin client environments	8 KB
File server (Microsoft Windows NT, Windows 2000, Windows Server 2003, Novell NetWare)	16 KB
File server (Other)	8 KB
Web server	8 KB
Other	8 KB

More information

- Changing the stripe-unit size (action)
- Fine-tuning your system

Fine-tuning the write-cache mode: (Novell NetWare only)

If you are preparing to install Novell NetWare 5.x from the startable *Novell NetWare 5.x* CD, you must set the write-cache mode for all logical drives to write through. Complete the following steps to accomplish this:

1. In the Logical devices view, right-click (logical drive).
2. Click **Change write-cache mode to write-through**.
3. Repeat steps 1 and 2 for each logical drive.

Selecting a RAID level and tuning performance: Disk arrays are used to improve performance and reliability. The amount of improvement depends on the application programs that you run on the server and the RAID levels that you assign to the logical drive.

Each RAID level provides different levels of fault-tolerance (data redundancy), utilization of physical drive capacity, and read and write performance. In addition, the RAID levels differ in regard to the minimum and maximum number of physical drives that are supported.

When selecting a RAID level for your system, consider the following factors.

Note: Not all RAID levels are supported by all ServeRAID controllers.

RAID level	Data redundancy	Physical drive capacity utili- zation	Read performance	Write performance	Built-in spare drive	Min. number of drives	Max. number of drives
RAID level-0	No	100%	Superior	Superior	No	1	16
RAID level-1	Yes	50%	Very high	Very high	No	2	2
RAID level-1E	Yes	50%	Very high	Very high	No	3	16
RAID level-5	Yes	67% to 94%	Superior	High	No	3	16
RAID level-5E	Yes	50% to 88%	Superior	High	Yes	4	16
RAID level-5EE	Yes	50% to 88%	Superior	High	Yes	4	16
RAID level-6	Yes	50% to 88%	Very High	High	No	4	16
RAID level-00	No	100%	Superior	Superior	No	2	60
RAID level-10	Yes	50%	Very high	Very high	No	4	16
RAID level-1E0	Yes	50%	Very high	Very high	No	6	60
RAID level-50	Yes	67% to 94%	Superior	High	No	6	60 (SCSI) 128 (SAS, SATA)
RAID level-60	Yes	50% to 88%	Very High	High	No	8	128

RAID level	Data redundancy	Physical drive capacity utili- zation	Read performance	Write performance	Built-in spare drive	Min. number of drives	Max. number of drives
Spanned Volume	No	100%	Superior	Superior	No	2	48
RAID Volume	No	50% to 100%	Superior	Superior	No	4	48

Physical drive utilization, read performance, and write performance depend on the number of drives in the array. Generally, the more drives in the array, the better the performance.

More information

- Understanding RAID technology
- Selecting the logical drive size
- Selecting the RAID level by array capacity
- Creating logical drives (action)
- Configuring RAID and creating arrays (action)
- Creating logical drives in the wizard
- Things to consider when changing the RAID level

Things to consider when changing the RAID level:

SR-SCSI

ServeRAID SCSI controllers

If you have a...	And you want...	Then...
RAID level-0 array with two or more drives	RAID level-5	Add one and **only** one drive
RAID level-1 array with two and **only** two drives	RAID level-5	Add one and **only** one drive
RAID level-5 array	RAID level-0	The ServeRAID Manager removes the **last** drive in the array
RAID level-5 Enhanced array	RAID level-5	The ServeRAID Manager does not add or remove a drive

SR-8i SR-7t

ServeRAID SATA and SAS controllers

If you have a...	And you want...	Then...
RAID level-0 logical drive with two or more drives	RAID level-5	Add at least one drive
RAID level-5 logical drive	RAID level-0	Optionally, remove one drive
RAID level-0 logical drive	RAID level-10	**Double** (at least) the drive count
RAID level-6 logical drive	RAID level-5	Optionally, remove one or two drives
RAID level-5 logical drive	RAID level-5EE	Add at least one drive

If you have a...	And you want...	Then...
RAID level-5EE logical drive	RAID level-5	Optionally, remove one drive
RAID level-5 logical drive	RAID level-10	Make sure you have (source-1)*2 drives total
RAID level-1 logical drive	RAID level-5	Add at least one drive
RAID level-1 logical drive	RAID level-10	Add at least two drives

ServeRAID Enclosures (networked storage)

If you have a...	And you want...	Then...
RAID level-0 logical drive with two or more drives	RAID level-5	Add at least one drive
RAID level-1 logical drive	RAID level-5	Add at least one drive
RAID level-1 logical drive	RAID level-10	Add at least two drives in multiples of two

More information

- Changing the RAID level (action)
- Changing the RAID level (SAS, SATA, HostRAID) (action)
- Changing the RAID level (enclosures) (action)
- Understanding logical-drive migration

Application environment groups reference:

Environment	Applications
Groupware	Lotus Notes Microsoft Exchange Other
Transaction processing	DB2 Informix Oracle SQLServer Sybase Other
Decision support or data warehousing	DB2 Informix Oracle SQLServer Sybase Other
Thin client environments	Citrix WinFrame or MetaFrame Microsoft Terminal Server Other
File server	Novell NetWare Microsoft Windows 2000 Microsoft Server 2003 Other network operating systems (for example, OpenServer, Linux)
Web server	Apache Microsoft IIS Netscape Commerce Server Other

Understanding write-cache mode for physical drives: When using the write-cache-mode option, you can choose from two available settings.

Write back

For the write back setting, the controller sends data to the physical drive for storage. Subsequently, the physical drive sends a confirmation to the controller **before** actually storing the data. Doing so increases performance, but also contains an element of risk.

Attention:

1. It is possible to lose data if a power outage occurs while using the write back setting. Consider carefully whether to enable write back on a physical drive. Depending on how you use the system, write back might be undesirable.
2. If you set this feature to write back, wait at least 10 seconds after your last operation before you turn off your system. Failure to follow this practice might result in lost data.

Write through

For the write-through setting, the controller sends data to the physical drive for storage. Subsequently, the physical drive stores the data; then, sends a confirmation to the controller. Using write through can decrease performance, but has no risk of losing data.

More information

- Changing the write-cache mode on a physical drive (action)
- Fine-tuning your system

Understanding write-cache mode for logical drives: When using the write-cache-mode option, you can choose from two available settings.

Write back

For the write-back setting, the operating system sends data to the controller to write to a storage device. Subsequently, the controller sends a confirmation to the operating system *before* actually writing the data to the storage device. Doing so increases performance, but also contains an element of risk. For example, if there is a power failure, the data currently in the controller cache is lost. This is no risk when using a controller with a battery-backup cache. The battery preserves the data in the controller cache in the event of a power failure.

Attention:

1. It is possible to lose data if a power outage occurs while using the write back setting without a battery-backup cache device. If your controller does not have a battery-backup cache installed and enabled, consider carefully whether to enable write back on a logical drive. Depending on how you use the system, write back might be undesirable.
2. If you do not have a battery-backup cache installed and enabled and you set this feature to write back, wait at least 10 seconds after your last operation before you turn off your system. Failure to follow this practice might result in lost data.

Write through

For the write-through setting, the operating system sends data to the controller to write to a storage device. Subsequently, the controller writes the data to the storage device; then, sends a confirmation to the operating system. This setting can decrease performance, but contains no risk of losing data.

More information

- Changing the write-cache mode on a logical drive (action)
- Fine-tuning your system

Understanding stripe-unit size: With RAID technology, data is **striped** across an array of physical drives. This data-distribution scheme complements the way the operating system requests data.

The granularity at which data is stored on one drive of the array before subsequent data is stored on the next drive of the array is called the **stripe-unit size** .

You can set the stripe-unit size to 8 KB, 16 KB, 32 KB, or 64 KB. You can maximize the performance of your ServeRAID controller by setting the stripe-unit size to a value that is close to the size of the system I/O requests. For example, performance in transaction-based environments, which typically involve large blocks of data, might be optimal when the stripe-unit size is set to 32 KB or 64 KB. However, performance in file and print environments, which typically involve multiple small blocks of data, might be optimal when the stripe-unit size is set to 8 KB or 16 KB.

SR-8i SR-7t HostRAID The ServeRAID-7t, ServeRAID-8i, and HostRAID controllers do not support an 8 KB stripe-unit size. The ServeRAID-7t and ServeRAID-8i controllers support these additional stripe-unit sizes: 128 KB, 256 KB, 512 KB, and 1024 KB.

The collection of stripe units, from the first drive of the array to the last drive of the array, is called a *stripe*.

SR-SCSI HostRAID After you configure an array and store data on the logical drives, you cannot change the stripe-unit size without destroying data in the logical drives.

You can set the stripe-unit size to 8 KB, 16 KB, 32 KB, or 64 KB. The default setting is 8 KB data bytes.

- When the stripe-unit size is 8 KB or 16 KB, the maximum number of physical drives in an array is 16.
- If you have a ServeRAID-3H or ServeRAID-3HB controller using ServeRAID firmware (version 4.0, or later) and the stripe-unit size is 32 KB or 64 KB, the maximum number of physical drives in an array is 16. Otherwise, when the stripe-unit size is 32 KB or 64 KB, the maximum number of physical drives in an array is 8.
- If you have a ServeRAID-4 controller and the stripe-unit size is set to 32 KB or 64 KB, the maximum number of physical drives in an array is 16.

More information

- Changing the stripe-unit size (action)
- Application environment groups reference

Understanding enabled and disabled read-ahead cache mode for logical drives:

Note: You can enable and disable the read-ahead setting without destroying data in a logical drive by using the ServeRAID Manager in bootable-CD mode.

Enabled read-ahead cache mode

The controller transfers data from the logical drive to its local cache in increments equal to the stripe-unit size. This provides excellent overall performance when workloads are steady and sequential. However, if the workload is random or the system I/O requests are smaller than the stripe-unit size, reading ahead to the end of the stripe might degrade performance.

Disabled read-ahead cache mode

The controller transfers data from the logical drive to its local cache in increments equal to the system I/O request size, without reading ahead to the end of the stripe. This provides excellent overall performance when the workload is random or the system I/O requests are smaller than the stripe-unit size.

More information

- Enabling and disabling read-ahead cache mode (action)
- Changing the stripe-unit size (action)
- Understanding stripe-unit size

Understanding adaptive read-ahead cache mode: The ServeRAID controller continually reevaluates whether to transfer data from disk to its local cache in increments equal to the stripe-unit size *or* in increments equal to the system I/O request size.

The ServeRAID-4 controllers and the ServeRAID-3HB controller come with the adaptive read-ahead mode as a standard feature. The ServeRAID-3H and ServeRAID-3L controllers must use new ServeRAID firmware (version 3.50, or later) to implement adaptive read-ahead cache mode. You can download and install the new firmware for the ServeRAID-3H and ServeRAID-3L controllers.

If you have the controller and firmware for adaptive read-ahead mode installed, the ServeRAID Manager reports Adaptive on the controller properties pane. Otherwise, the properties pane states either Enabled or Disabled.

Note: With this new firmware, the non-Adaptive read-ahead modes are no longer available in the ServeRAID Manager. You *can* override Adaptive read-ahead, but not through the ServeRAID Manager. Instead, use the IPSSEND command-line program. This program comes with the device option and is on the *IBM ServeRAID Support* CD. For more information, refer to the *IBM ServeRAID User's Reference*.

Read-ahead cache mode settings: Depending on your controller's level of firmware, there are three settings for read-ahead cache mode:

- Enabled or Disabled
- Adaptive

You can change the read-ahead setting without destroying data in a logical drive using the ServeRAID Manager in bootable-CD mode.

Changing the stripe-unit size:

Note:

1. This action is supported in bootable-CD mode only.
2. This action is not supported on the ServeRAID-7t controller, the ServeRAID-8i conroller, and the integrated RAID controller.
3. On ServeRAID-7t controller and ServeRAID-8i controller, use the configuration wizard to change the stripe-unit size.

Use this action to set the size of the chunk of data that the controller reads from each physical drive at a time. To maximize overall performance, choose a stripe-unit size that is close to the size of the system I/O request.

Attention: After you configure an array and store data on the logical drives, you cannot change the stripe-unit size without destroying data in the logical drives.

1. In the Enterprise view, click (controller).
2. Right-click **Change stripe-unit size** and then click a stripe-unit size. The following choices are available:

   8 KB With this stripe-unit size, an array can have a maximum of 16 physical drives.

   16 KB With this stripe-unit size, an array can have a maximum of 16 physical drives.

   32 KB This setting is available if the controller supports an array with 16 physical drives using a 32 KB stripe-unit size, or the controller contains no arrays with more than 8 physical drives. Otherwise, this setting is unavailable.

   64 KB This setting is available if the controller supports an array with 16 physical drives using a 64 KB stripe-unit size, or the controller contains no arrays with more than 8 physical drives. Otherwise, this setting is unavailable.

More information

- Fine-tuning your system
- Understanding stripe-unit size
- Understanding enabled and disabled read-ahead cache mode

Changing the write-cache mode on a logical drive:

Note:

1. This action is supported in bootable-CD mode for ServeRAID SCSI controllers.
2. This action is supported in interactive mode only for the ServeRAID-8i controller and ServeRAID-7t controller.
3. This action is not supported on the integrated RAID controller.

Use this action to change the write-cache mode setting of a logical drive to write through or write back. The write-cache mode determines if the controller writes data to the drive before or after sending a confirmation to the operating system.

Note:

1. (RAID level-x0 only) All logical drives in a spanned array must be either write back or write through. That is, the write-cache mode settings for the logical drives in the spanned array *cannot* be a mixture of write back and write through.
2. (Failover environment only) A fault-tolerant controller pairing must be configured as write through. This is because the logical drives are configured as ″Shared.″ For more information, refer to the *IBM ServeRAID User's Reference* .
3. (Cluster environment only) If the logical drives are configured as ″Shared,″ the controllers must be configured as write through. For more information, refer to the *IBM ServeRAID User's Reference* .

1. In the Logical devices view, click (logical drive).
2. Right-click **Change write-cache mode to write back or write through**.
3. If you click write back and you do **not** have a battery-backup cache, click **Yes** to confirm your choice because there is a risk of losing data.

   If you click **write back** and you do have a battery-backup cache, the ServeRAID Manager changes the mode.

   If you click **write through**, the ServeRAID Manager changes the mode.

More information

- Understanding write-cache mode for logical drives

Changing the write-cache mode:

Note: This action is not supported on enclosures with a dual controller configuration.

Use this action to change the write-cache mode setting of all the physical drives in an array. The write-cache mode determines if a physical drive stores data before or after sending a confirmation to the controller.

Note: It is possible to lose data if a power outage occurs while using the write-back setting. Consider carefully whether to use the write-back setting on a physical drive. Depending on how you use the system, the write-back setting might be undesirable.

1. In the Logical devices view, click (array).
2. Right-click **Configure write cache → write back or write through**.
3. Click **Yes** to confirm.

More information

- Understanding write-cache mode for physical drives

Enabling and disabling read-ahead cache mode: Use this action to set the read-ahead cache mode to Enabled or Disabled. If the read-ahead cache mode is set to Adaptive, you cannot use this action.

Note: You **can** override the Adaptive setting, but not through the ServeRAID Manager. Instead, use the IPSSEND command-line program. This program is on the *IBM ServeRAID Support* CD.

1. In the Logical devices view, click (array).
2. Right-click **Configure read cache** → **enabled**or **Configure read cache** → **disabled**.
3. Click **Yes** to confirm.

More information

- Understanding read-ahead cache mode
- Understanding adaptive read-ahead cache mode

Managing storage devices

Managing direct attached storage devices:

Understanding FlashCopy backup: The FlashCopy function creates a quick backup copy of data. It sets up a link between the source and target logical drives; then, it creates a backup of the source data on the target drive. Any changes made to the source drive after you create a FlashCopy backup are not reflected on the target drive. You can use the backup copy of data for tape backup, drive cloning, and multi-server rollout.

There are two primary FlashCopy functions: backup and nobackup. The FlashCopy backup function copies the entire contents of the source drive to the target drive so that entire logical drives can be moved from one server to another. The FlashCopy nobackup function creates a temporary copy of a drive for tape drive backup and reference purposes. The FlashCopy nobackup function is less I/O-intensive than the FlashCopy backup function.

Before using FlashCopy, consider the following requirements:

- You can perform a FlashCopy backup operation on only one controller at a time.
- The source and target logical drives can have the same or different RAID level (they do not have to match).
- The source and target logical drives must be on the same controller.
- You can create a maximum of four independent FlashCopy backups per controller.
- There is no limit on the size of the source and target logical drive. However, for both FlashCopy backup and nobackup operations, the size of the target drive must be greater than or equal to the source drive.
- You cannot perform any action on a FlashCopy source or target logical drive (such as synchronizing the drive). You can only delete a FlashCopy backup.

More information

- Creating a FlashCopy backup
- Removing a FlashCopy backup

Understanding copy back mode:

Note: Copy back is supported on the ServeRAID-8i, ServeRAID-7k, ServeRAID-6M, and ServeRAID-6i/6i+ controllers only. It is not supported in cluster or failover pair configurations.

Copy back is a method of restoring a logical drive's original configuration after you replace a failed drive in an array. It allows you to restore the data to its prior location, before the logical drive was rebuilt from its spare.

An example

A RAID Level-5 logical drive consists of three physical drives and a spare. When a drive in the array fails, the spare drive is used to rebuild the logical drive. When you replace the failed drive, copy back moves the data from the former spare to the newly replaced drive. Then, the former spare resumes its original role as the spare drive.

Copy back is enabled by default; it starts automatically when the ServeRAID controller detects that a failed drive is replaced. The ServeRAID Manager displays a progress indicator in the status bar while the operation is in progress. To disable copy back, choose Disable copy back mode from the controller object Action menu.

You cannot perform any other actions on the controller until copy back is completed.

Note: Copy back is disabled by default when you upgrade the ServeRAID software from a previous release.

More information

- Enabling and disabling copy back mode (action)

Enabling and disabling copy back mode:

Note: Copy back is supported on the ServeRAID-8i, ServeRAID-7k, ServeRAID-6M, and ServeRAID-6i/6i+ controllers only. It is not supported in cluster or failover pair configurations.

Use this action to change the copy back-mode setting. This setting determines if the ServeRAID Manager restores a logical drive's original configuration after you replace a failed drive in an array.

Copy back is enabled by default except when you upgrade the ServeRAID software from a previous release; then, copy back mode is disabled by default.

1. In the Enterprise view, click (controller).
2. Right-click **Enable or Disable copy back mode.**

More information

- Understanding copy back mode

Creating a FlashCopy backup: Use this action to create a FlashCopy backup of data on a logical drive. You can create a FlashCopy with backup (full copy) or without backup (temporary copy).

1. In the Logical devices view, click (Logical drive).

2. Right click **Create FlashCopy** → **with backup to** → **the logical drive name** or **Create FlashCopy** → **without backup to** → **the logical drive name**.

More information

- Understanding FlashCopy backup
- Removing a FlashCopy backup

Removing a FlashCopy backup: Use this action to remove a FlashCopy backup of a logical drive.

1. In the Logical devices view, click (FlashCopy).
2. Right click **Remove FlashCopy**.

More information

- Understanding FlashCopy backup
- Creating a FlashCopy backup

Managing networked storage devices (enclosures):

Shutting down an enclosure: Use this action to shut down the controllers in an enclosure.

Attention: Use care when you shut down an enclosure. Users cannot access the data on the arrays and logical drives until the enclosure is restarted.

Follow these steps to shut down an enclosure:

1. In the Enterprise view, click (enclosure) that you want to shut down.
2. Right-click **Shut down enclosure.**.
3. Click **Yes** when asked to confirm that you want to shut down the enclosure.

More information

- Updating enclosure software
- Restarting an enclosure

Restarting an enclosure: Use this action to restart the controllers in an enclosure.

1. In the Enterprise view, click (enclosure) that you want to restart.
2. Right-click **Restart enclosure**.
3. Click **Yes** when asked to confirm that you want to restart the enclosure.

   Attention: Restarting the enclosure may take several minutes. Data on the controller(s) are unavailable during that time.

More information

- Updating enclosure software
- Shutting down an enclosure

Changing controller date and time: Follow these steps to change the controller date and time:

1. In the Enterprise view, click (enclosure).
2. Right-click **Change controller date and time**. The ″ Change controller date and time″ window opens.
3. Select the new date, time, and time zone from the calendar and time controls.
4. Click **OK**. The ServeRAID Manager will apply the update to the controller.

More information

- Updating enclosure software
- Shutting down a controller

Configuring host information: Use this action to configure the network settings for an external storage enclosure. You can set the host name, domain, primary and secondary DNS servers, and the default gateway for network access.

1. In the Enterprise view, click ▦ (enclosure).
2. Right-click **Configure network details**. The ″Configure Host Information″ window opens
3. In the **Host name** field, type the enclosure's host name.
4. In the **Domain name** field, type the enclosure's network domain; for example, `mydomain.com`.
5. In the **Primary Domain Name Server** field, type the TCP/IP address of the enclosure's primary DNS.
6. (optional) In the **Secondary Domain Name Server** field, type the TCP/IP address of enclosure's secondary DNS.
7. In the **Default Gateway** field, type the TCP/IP address of the default gateway for network access.
8. Click **OK**.

More information

- Configuring Ethernet ports

Configuring Ethernet ports: Use this action to configure the Ethernet settings for the controllers in an enclosure. For iSCSI initiators, you can configure three ports per controller: the management port and two iSCSI ports. For fibre channel initiators, you can configure the management port for each controller.

For the management port (iSCSI or fibre channel), you can configure the port link speed. For iSCSI ports, you can configure network settings, such as Maximum Transfer Unit (MTU), TCP/IP address, and subnet mask. If your network uses dynamic IP addresses, you can enable DHCP.

Configuring the management port

To configure the management port:

1. In the Physical devices view, click a controller in the enclosure.

   Attention: Be sure to switch to enclosure view first.
2. Right click **Configure Ethernet port** → **ETH 0 (management)**. The Configure Ethernet port window opens.
3. From the **Link speed** drop-down list, select the port link speed, in megabytes. To automatically detect the link speed (1 GB or less), select **AUTO**.
4. Click **OK**.

Configuring iSCSI ports (iSCSI initiators only)

To configure the iSCSI ports:

1. In the Physical devices view, click a controller in the enclosure.

   Attention: Be sure to switch to enclosure view first.
2. Right-click **Configure Ethernet port** → **ETH 2 (iscsi)** or **Configure Ethernet port** → **ETH 3 (iscsi)**. The Configure Ethernet port window opens.

3. From the **Link speed** drop-down list, select the port link speed, in megabytes, as described above.
4. From the **Maximum transmission unit** drop-down list, select the MTU packet size, in megabytes; either 1500 MB or 9000 MB.
5. If your network uses dynamic IP addresses, click **DHCP**; then, skip to step 9.
6. In the **IP address** field, enter the iSCSI port's TCP/IP address.
7. In the **Subnet mask** field, enter the TCP/IP number of the device's TCP/IP subnet.
8. (optional) In the **Broadcast address** field, enter the TCP/IP address for sending messages to machines on the network.
9. Click **OK**.

More information

- Configuring host information

Updating controller and enclosure software: The ROM Update wizard guides you through the process of updating the software for your direct-attached and network-attached storage devices:

- For direct-attached storage devices, the wizard updates the BIOS software for all ServeRAID controllers of the same type on the local and remote systems. You can update only one type of controller at a time.
- For external (network-attached) storage devices, the wizard updates the enclosure and controller software for all controllers in the enclosure (single or dual controller configurations).

Before you can use the ROM update wizard, you must download the latest software images from your vendor's software support site on the World Wide Web. For direct-attachd storage devices, the image upgrade files typically come in sets of two or more and have a `.ufi` file extension. For external storage enclosures, the (single) image upgrade file has a `.upgrade` file extension.

To update the controller software for direct-attached or network-attached storage devices:

1. In the Enterprise view, click (Direct attached storage object) or (Networked storage object)
2. Right-click **Update controller images**. The ROM Update wizard opens.
3. Click **Next**.
4. Choose the ROM image update files.
5. Click **Next**; then, select the controllers or enclosures you want to update.
6. Click **Next**; then, review the update summary.
7. Click **Apply**. The ServeRAID Manager applies the software update to the selected controllers or enclosures.
8. Restart the server(s) or enclosures to activate the new ROM image.

Recovering storage with fail back:

Note: This action is not supported on enclosures with a single controller configuration.

For enclosures with a dual controller configuration, use this action to restore an array to its preferred owner. When a controller fails in an enclosure (or is intentionally removed), its arrays and logical drives automatically fail over to the other controller. This action moves the arrays and logical drives back to the controller to which they were originally assigned.

Note: You cannot use fail back until the failed controller is replaced or repaired.

To fail back storage in an enclosure:

1. In the Enterprise view, click (enclosure).
2. Right-click **Fail back storage**.
3. Click **Yes** when prompted to confirm the action.

More information

- Moving an array to a different controller

Managing Foreign arrays and Alien arrays:

Foreign arrays

You can export RAID data and transfer the drives and RAID configuration to another enclosure. When you export an array, it is called a

foreign array

. The array icon is grayed out in the Logical devices view. You cannot perform any action on a foreign array, except view array components.

When you physically remove the component drives from the enclosure, the ServeRAID Manager removes the foreign array icon from the console. When you insert the drives in another enclosure, the ServeRAID Manager displays a foreign array icon in the Logical devices view. You must then import the RAID data to complete the transfer. You may then use the array to store and manage data.

Alien arrays

If the ServeRAID Manager only partially completes the array configuration process, the resulting array is called an

alien array

. For example, an alien array might be created if power is interrupted during array creation. You cannot use an alien array to store or manage data. The array icon is grayed out in the Logical devices view. The only action you can perform on an alien array is to delete it.

More information

- Exporting an array
- Importing an array
- Deleting an array

Managing FlashCopy backups:

Note: You cannot use the ServeRAID Manager console create a FlashCopy for an enclosure-based logical device. To create a FlashCopy backup, you must use the command line interface provided with your IBM DS300 or DS400 storage enclosure. Or you can use the `sstool` command provided with the FlashCopy agent.

A FlashCopy backup creates a point-in-time backup of logical drive data. It sets up a link between the source and target logical drives. When data on the source drive changes, the differences are captured on the target drive, maintaining the data as it looked at the time the backup was created. You can use the backup copy of data for tape backup, drive cloning, and multi-server rollout.

Initially, a FlashCopy backup starts at 10% the size of the FlashCopy source. As new data are written to the source drive, the FlashCopy backup grows to the size of FlashCopy source. This is known as the FlashCopy virtual size.

The ServeRAID Manager displays both the virtual size and actual size of a FlashCopy backup, depending on the situation:

- In the Logical devices view, it displays the virtual size:

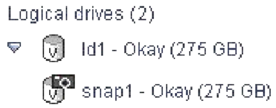

Logical drives (2)
ld1 - Okay (275 GB)
snap1 - Okay (275 GB)

- In the Configuration wizard, it displays the actual size of the *first* FlashCopy backup. It displays 0 for all other FlashCopy backups.
- In the logical drive Properties panel, it displays both the virtual size and actual size.

You cannot perform any action on a FlashCopy target. You can only delete a FlashCopy target, using the Delete logical drive action.

More information

- Deleting a logical drive

Updating the software license key: Use this action to enable new features on the specified enclosure. To enable new features, you must obtain a feature key from your ServeRAID Manager sales or support representative.

To update the software license key:

1. In the Enterprise view, click (enclosure).
2. Right-click **Enter software key**. The ″Enter software key″ window opens.
3. In the Feature key field, enter the feature key.
4. Click **OK**. A confirmation window opens.
5. Click **OK** to confirm.

More information

- Updating enclosure software

Updating controller and enclosure software: The ROM Update wizard guides you through the process of updating the software for your direct-attached and network-attached storage devices:

- For direct-attached storage devices, the wizard updates the BIOS software for all ServeRAID controllers of the same type on the local and remote systems. You can update only one type of controller at a time.
- For external (network-attached) storage devices, the wizard updates the enclosure and controller software for all controllers in the enclosure (single or dual controller configurations).

Before you can use the ROM update wizard, you must download the latest software images from your vendor’s software support site on the World Wide Web.

For direct-attachd storage devices, the image upgrade files typically come in sets of two or more and have a `.ufi` file extension. For external storage enclosures, the (single) image upgrade file has a `.upgrade` file extension.

To update the controller software for direct-attached or network-attached storage devices:

1. In the Enterprise view, click (Direct attached storage object) or (Networked storage object)
2. Right-click **Update controller images**. The ROM Update wizard opens.
3. Click **Next**.
4. Choose the ROM image update files.
5. Click **Next**; then, select the controllers or enclosures you want to update.
6. Click **Next**; then, review the update summary.
7. Click **Apply**. The ServeRAID Manager applies the software update to the selected controllers or enclosures.
8. Restart the server(s) or enclosures to activate the new ROM image.

Adding a remote system

Note: This action is not supported when using the following:

- ServeRAID Manager in bootable-CD mode
- ServeRAID Manager as a plug-in application

Use this action to connect to a remote system or management station and add it to the Enterprise view.

Tip: Before adding a remote system, verify that the system is running the required software. To add a remote system with direct-attached storage, the ServeRAID Manager must be running on that system. To add a management station with a network-attached storage enclosure, the ServeRAID Manager management station software must be running on that system.

1. From the Remote menu, select **Remote** → **Add** or click (Add) on the toolbar. The Add managed system window opens.
2. From the Type drop-down list, select:
   - **Managed System**, to add a remote system with direct-attached storage
   - **Management Station**, to add a remote management station with a network-attached storage enclosure
3. Type the remote host name or TCP/IP address.
4. (Managed system only) Type the remote system startup port number. The default port number is 34571.
5. Type your user name and password.

   Note: The password is case sensitive.
6. If you want to save the user name and password, select the **Save user name/Password** box. Once you successfully connect to the remote system, the ServeRAID Manager stores this information in a file along with the host name so you do not have to type your user name and password every time you run the ServeRAID Manager.
7. Click **Connect**.

More information

- Failing to add a remote system
- Successfully adding a remote system
- Specifying remote access settings
- Removing a remote system

Viewing the ServeRAID Manager event log

Note: This action is not supported in bootable-CD mode.

Use this action to view events in the ServeRAID Manager agent event log.

1. From the Actions menu, click **Agent actions** → **View agent event log**. The Agent event log window opens.
2. If you want to save the event log to a file, click **File** → **Save As**. The default is Events.txt.
3. Click **File** → **Close** to close event log window.

More information

- Configuring the ServeRAID Manager agent
- Configuring the ServeRAID Manager agent general settings
- Using the ServeRAID Manager agent

Using the ServeRAID Manager agent

Starting the ServeRAID Manager agent

The ServeRAID Manager agent is a monitoring agent for ServeRAID controllers that uses less memory than the ServeRAID Manager graphical user interface (that is, the console). You can manage and configure a server running the agent through a ServeRAID Manager console running on another system.

Note: If the system administrator chose during installation to start the agent as a background service (daemon), the agent is already running. Verify that the agent is not running before performing this procedure.

Use the following procedure to start the ServeRAID Manager agent on your system.

Windows

On Windows systems, the ServeRAID Manager agent is installed as a background service. It starts automatically when the system is started. To start the agent manually, use the Windows Administrative tools. See Verifying that the agent is running.

OS/2

To start the ServeRAID Manager agent on a Windows or OS/2 system:

1. Change to the directory where you installed the ServeRAID Manager program by typing:

   ```
   cd \RaidMan
   ```
2. Press Enter.
3. Type the following:

RaidAgnt

4. Press Enter.

NetWare

To start the ServeRAID Manager agent on NetWare:

1. From the NetWare console, type the following:

   LOAD RaidAgnt
2. Press Enter.

OpenServer, UnixWare, and Open UNIX

To start the ServeRAID Manager agent on OpenServer, UnixWare, and Open UNIX:

1. Change to the directory where you installed the ServeRAID Manager program by typing one of the following:

For OpenServer	cd /opt/RaidMan
For UnixWare or Open UNIX	cd /opt/RaidMan

2. Press Enter.
3. Type the following:

   sh RaidAgnt.sh
4. Press Enter.

Linux

On Linux systems, the ServeRAID Manager agent runs as a background process (daemon). It starts automatically when the system is started. It is loaded from the following location:

```
/etc/init.d/raid_agent
```

More information

- Using the ServeRAID Manager agent
- Configuring the ServeRAID Manager agent
- Verifying that the ServeRAID Manager agent is running

Viewing the ServeRAID Manager event log

Note: This action is not supported in bootable-CD mode.

Use this action to view events in the ServeRAID Manager agent event log.

1. From the Actions menu, click **Agent actions** → **View agent event log**. The Agent event log window opens.
2. If you want to save the event log to a file, click **File** → **Save As**. The default is Events.txt.
3. Click **File** → **Close** to close event log window.

More information

- Configuring the ServeRAID Manager agent
- Configuring the ServeRAID Manager agent general settings
- Using the ServeRAID Manager agent

Receiving events from a removed system: When a remote system is in the Enterprise view, you always receive events that occur on that remote system. With the Remove remote system action, you can choose between ″Continue to receive events from remote system″ or ″Do not continue to receive events from remote system.″ If you choose the first option, you can monitor for problems on remote systems without having the system in your Enterprise view.

The default is to continue receiving events.

If you choose **not** to continue receiving events, the ServeRAID Manager on your local system must connect to the remote system. If the ServeRAID Manager successfully connects, it removes your local system from the remote system notification list. Because the notification list determines what systems receive events from that remote system, your local system stops receiving events.

If the ServeRAID Manager cannot connect to the remote system and is therefore unable to remove your system from the notification list, another window opens asking if you want to remove the system from the Enterprise view even though you will continue receiving events.

- If you click **Yes**, the ServeRAID Manager removes the remote system from the Enterprise view tree and you continue receiving events.
- If you click **No**, the remote system remains in the Enterprise view tree and you continue receiving events. You can reattempt removing your system from the notification list at a later time.

Event viewer description of events: The event viewer description of events can inform you of the following:

- A rebuild is started
- A configuration is applied
- A failed drive is detected
- Other potential problems that might occur to your managed systems.

Events for remote systems display in the event viewer when the following occurs:

1. Your local system is defined in a remote system’s notification list.
2. You have added the remote system to your Enterprise view using ″ Add remote system.″

More information

- Using the ServeRAID Manager interface > Event viewer
- Using the Notification Manager

Configuring the ServeRAID Manager agent

You can configure the ServeRAID Manager agent port number and alarm. You can also configure the ServeRAID Manager agent to log events to the operating system event log. To configure the agent, edit the file RaidAgnt.pps or adjust the General settings in the ServeRAID Manager agent console. RaidAgnt.pps is located in the same directory where you installed the ServeRAID Manager. If the RaidAgnt.pps file does not exist when you start the ServeRAID Manager agent, a new file is created with the default settings.

Note:

1. The RaidAgnt.pps file is preserved during a ServeRAID Manager upgrade.

2. If you change settings while the ServeRAID Manager agent is running, you must stop and restart the agent to make the changes take effect.
3. The ServeRAID Manager agent alarm is not supported by the ServeRAID Manager in bootable-CD mode.

Configuring the port number

The default port number for the ServeRAID Manager agent and console (client) is 34571. To change the port number, edit the following line in the file RaidAgnt.pps:

agent.startupPortNum=34571

For accessing remote systems, the ServeRAID Manager uses four consecutive ports starting from the startup port: 34571, 34572, 34573, and 34574. If your system has a conflict with these ports, change the startup port to a different port number.

Configuring the alarm

When the ServeRAID Manager agent is started, it reads the alarm settings. When a critical or fatal event occurs in the RAID subsystem, the ServeRAID Manager agent triggers its alarm, if enabled. The alarm continues at the specified interval until you either:

- Correct all the critical and fatal problems (if agent.auto.off.alarm is set to true).
- Delete the file alarm.on. This file is located in the same directory where you installed the ServeRAID Manager. When you delete this file, the alarm stops until the next critical or fatal event occurs. The alarm.on file is automatically created each time the alarm starts.

You can configure the following alarm settings:

- **agent.enable.alarm** Specifies whether the agent alarm is enabled or disabled. Set this value to true to enable the alarm or false to disable the alarm. The default value is false.
- **agent.interval.alarm** Specifies the interval (in seconds) between audible alarms. The default is 300 seconds (5 minutes).
- **agent.auto.off.alarm** Specifies whether the alarm should turn off automatically when no more problems are detected. Set this value to true to cause the alarm to turn off automatically, or false to manually turn the alarm off. If you set this value to false, you must delete the file alarm.on to turn off the alarm. The default is true.

Configuring event logging

The ServeRAID Manager agent logs warning and fatal events to the operating system event log. You can view the event log from the ServeRAID Manager console. For Windows systems, you can also view the event log with the Windows Event Viewer. For Unix systems, events are logged according to the /etc/syslog.conf settings. (For more information, see the syslogd(8) man page.)

You can enable and disable event logging by editing the following line in the file RaidAgnt.pps. Set the value to true to enable event logging (the default) or false to disable event logging.

agent.enable.logEventsWithOS = true

More information

- Using the ServeRAID Manager agent
- Configuring the ServeRAID Manager agent general settings
- Configuring NetWare user authentication
- Verifying that the ServeRAID Manager agent is running
- Starting the ServeRAID Manager agent
- Specifying remote access settings
- Viewing the ServeRAID Manager agent event log

Configuring the ServeRAID Manager agent > general settings

Note: This action is not supported in bootable-CD mode.

Use this action to configure The ServeRAID Manager agent general settings, including the base port for the agent and console. The agent can log events to the operating system event log, sound an alarm when an event occurs, and broadcast events to users. If you enable event broadcasting, agents on Windows machines will display pop-up alert dialogs when events occur. On Linux machines, a message will be broadcast to all connected console processes using the ’wall’ command.

You can adjust the agent general settings with the following steps:

1. In the ″ServeRAID Manager Agent″ window, click the **General settings** tab.
2. In the Agent base port number field, enter the port number for the ServeRAID Manager agent and console. The default port is 34571.

   Note: For accessing remote systems, the ServeRAID Manager uses four consecutive ports starting from the base port: 34571, 34572, 34573, and 34574. If your system has a conflict with these ports, change the base port to a different port number.
3. To log events to the operating system event log, click **Save events in OS event log**.

   Note: You can use the View agent event log option to view the event log. You can also use operating system administration tools to view the event log.
4. To broadcast events to users, click **Broadcast events to logged-in users**.
5. Select the check box to enable or disable the alarm. When a critical or fatal event occurs in the RAID subsystem, the ServeRAID Manager agent triggers an alarm, if enabled.
6. If you enabled the alarm, adjust the length of time (in seconds) between alarms. The default is 300 seconds (5 minutes).
7. Click **Save changes**.
8. Restart the ServeRAID Manager for the settings to take effect.

Note: To load settings from the ServeRAID Manager agent configuration file, click Refresh from agent . Use this option if you change settings in the configuration file while the agent is running. If you change the port number, you must stop and restart the agent for the change to take effect.

More information

- Configuring the ServeRAID Manager agent
- Starting the ServeRAID Manager agent
- Verifying that the ServeRAID Manager agent is running

- Viewing the ServeRAID Manager agent event log

Using the ServeRAID Manager agent

The ServeRAID Manager agent includes the following components:

- Notification Manager
- Email Notification Manager
- Task Manager
- General settings

Because the agent runs as a background process, it uses less memory resources than the ServeRAID Manager console. However, the ServeRAID Manager agent is useful in the following situations:

- The agent can relay events to any ServeRAID Manager console installed on the local system or on a remote system. As events occur, the agent sends the events to each remote system. Use the Notification Manager from the agent console to add systems to the agent destination list.
- The agent can notify users by email when events occur on the local system. Use the Email Notification Manager to add recipients to the email notification list.
- The agent can run maintenance tasks on a regular basis, such as a logical drive migration or synchronization. Use the Task Manager to manage scheduled tasks.
- The agent can log events to the operating system event log, sound an alarm when an event occurs, and broadcast critical or fatal events to users. Use the General settings from the agent console to configure these settings.

  Note: ServeRAID Manager also logs events to a log file. Each component of the agent (except for the Task Manager) has its own log file.

The ServeRAID Manager agent monitors and generates events for critical or fatal problems in the ServeRAID Manager configuration every 5 seconds. These changes include, but are not limited to:

- defunct drives
- PFA drives
- Failed battery
- Offline or critical logical drives
- Failed controllers
- Enclosure problems
- Non-warranted drives. An event is sent at startup, console connection, and every 30 days

More information

- Configuring the ServeRAID Manager agent
- Configuring the ServeRAID Manager agent general settings
- Starting the ServeRAID Manager agent
- Verifying that the ServeRAID Manager agent is running
- Using the Notification Manager
- Using the Email Notification Manager
- Using the Task Manager
- Viewing the ServeRAID Manager agent event log (action)

Using Security Manager

Security Manager menu bar

The Security Manager menu bar has the following functions:

File →

Close Closes the Security Manager.

View →

Toolbar
Turns the toolbar on and off. The default is on (that is, selected).

Actions →

Add user

Delete user

Modify user

Help →

Information about this window
View context-sensitive information for the current window.

Search
Searches for one or more specified words in ServeRAID Manager Assist and displays a list of topics that include the words.

Contents
Presents the ServeRAID Manager Assist contents. You can use the contents to acquaint yourself with ServeRAID Manager Assist topics.

About ServeRAID Manager
Reports the ServeRAID Manager version number, copyright, and legal information.

Security Manager toolbar

The Security Manager toolbar has the following functions:

Add user

Delete user

Modify user

Information about this window View context-sensitive information for the current window.

Adding a user in the Security Manager

Use this action to permit a user access to the management station.

1. Click (Add user). The ″Add user″ window opens.
2.

   Fill in all the entry fields for the user you want to add. **User name**, **Password**, and **Confirm password** are required fields.

   Note: The password is case sensitive.

   .
3. Click **Add**.

More information

- Using the Security Manager
- Deleting a user in the Security Manager
- Modifying a user in the Security Manager

Deleting a user in the Security Manager

Use this action to revoke access to the management stations for a selected user.

1. Click the user that you want to delete from the security list.
2. Click (Delete user).
3. Click **Yes** to confirm. The ServeRAID Manager deletes the user from the security list.

More information

- Using the Security Manager
- Adding a user in the Security Manager
- Modifying a user in the Security Manager

Modifying a user in the Security Manager

Use this action to modify a user’s password.

1. Click the user you want to modify in the security list.
2. Click (Modify user). The ″User properties″ window opens.
3. Enter a new password; then, enter it again to confirm.
4. Click **Modify**.

More information

- Using the Security Manager
- Adding a user in the Security Manager
- Deleting a user in the Security Manager

Using SNMP Trap Manager

SNMP Trap Manager menu bar

The SNMP Trap Manager menu bar has the following functions:

File →

Close Closes the SNMP Trap Manager.

View →

Toolbar
Turns the toolbar on and off. The default is on (selected).

Actions →

Add system

Delete system

Modify system

Send test trap

Help →

Information about this window
View context-sensitive information for the current window.

Search
Searches for one or more specified words in ServeRAID Manager Assist and displays a list of topics that include the words.

Contents
Presents the ServeRAID Manager Assist contents. You can use the contents to acquaint yourself with ServeRAID Manager Assist topics.

About ServeRAID Manager
Reports the ServeRAID Manager version number, copyright, and legal information.

SNMP Trap Manager toolbar

The SNMP Trap Manager toolbar has the following functions:

Add system

Delete system

Modify system

Information about this window View context-sensitive information for the current window.

Adding a system in the SNMP Trap Manager

Use this action to add a system to the SNMP traps list. You can add *one* system only, such as an administrator's workstation. The system is notified of the traps that occur on the enclosures attached to the management station.

1. Click (Add system). The "Add system" window opens.
2. Type the host name or TCP/IP address of the system you want to add.
3. Type the remote system SNMP trap port number. The default port is 162.
4. Type the Community for the system.
5. From the Version drop-down list, select the SNMP version number. The default is SNMPv1.
6. Select the type of events you want to be notified of: error, warning, or informational.

7. Click **Add**.

More information

- Using the SNMP Trap Manager
- Deleting a system from the SNMP traps list
- Modifying system properties in the SNMP traps list
- Sending a test trap to a system in the SNMP traps list

Deleting a system in the SNMP Trap Manager

Use this action to delete a system from the SNMP traps list. After deleting the remote system, the SNMP Trap Manager no longer notifies the remote system of traps occurring on the enclosures connected to the management station.

1. In the SNMP traps list, click the system you want to delete.
2. Click (Delete system).
3. Click **Yes** to confirm. The SMNP Trap Manager deletes the system from the SNMP traps list.

More information

- Using the SNMP Trap Manager
- Adding a system to the SNMP traps list
- Modifying system properties in the SNMP traps list
- Sending a test trap to a system in the SNMP traps list

Modifying a system in the SNMP Trap Manager

Use this action to modify system properties in the SNMP trap list, such as the host name, TCP/IP address, or port number.

1. In the SNMP traps list, click the system that you want to modify.
2. Click (Modify system). The ″System properties″ window opens.
3. Change the system properties in the appropriate fields.
4. Click **Modify**.

More information

- Using the SNMP Trap Manager
- Adding a system to the SNMP traps list
- Deleting a system from the SNMP traps list
- Sending a test trap to a system in the SNMP traps list

Managing software and firmware

Confirming your software update

Use the software update summary to review all the changes that the ServeRAID Manager will apply to your controllers or enclosures:

1.

   Review the information displayed in the ″Update summary″ window.
2. When you are ready to continue, click **Apply**; then, click **Yes** when asked if you want to apply the software update.
3. When the update is complete, click **OK**.
4. Restart your server(s) or enclosures to activate the new software.

Changing BIOS-compatibility mapping

Note:

1. This action is supported in bootable-CD mode only.

2. This action is not supported on the integrated RAID controller.

Use this action to change BIOS-compatibility mapping. **Extended** indicates 8 GB Extended; **Limited** indicates 2 GB Limited.

The ServeRAID controller allows the migration of drives from the PCI RAID Adapter or Streaming RAID Adapter/A configurations. Using ″Change BIOS-compatibility mapping,″ you can set your ServeRAID controller to be compatible with these older drive configurations by selecting either 2 GB or 8 GB mapping. The default is 8 GB.

1. In the Enterprise view, click (controller).
2. Right-click either **Change BIOS-compatibility mapping** → **Extended** either or **Change BIOS-compatibility mapping** → **Limited**.
3. Click **Yes** to confirm the setting.

Updating BIOS and firmware code

Before configuring the controller, you must have the latest BIOS and firmware code installed on your server. Complete the following steps to update the levels of BIOS and firmware code:

1. Insert the *IBM ServeRAID Support* CD into the server CD-ROM drive, and turn on the server.

   The ROM Update Wizard automatically starts. The ROM (read-only memory) Update Wizard is a program that updates the BIOS and firmware code on your ServeRAID controllers. The wizard automatically identifies and scans each controller.

   If the BIOS and firmware code do not require updating, the wizard automatically stops and the ServeRAID Manager program starts. Use the ServeRAID Manager program to continue with configuring the controller.

   If the BIOS and firmware code require updating, a report screen opens with the following information:
   - Controller types found.
   - Controller slot number, if known.
   - Firmware version.
   - BIOS version.
   - Update status. If a controller has outdated BIOS or firmware code, the ROM Update Wizard marks the controller as a candidate for update.

   The ROM Update Wizard asks if you want to update. You decide whether to update, but you must update all or none of the controllers in your server; you cannot selectively update.
2. If you want to update your controllers, click **Update**. If the wizard detects an error, an error message appears and you are prompted to insert a diskette into your diskette drive. The wizard saves details about the error to a file on the diskette.

   If you do not want to update your controllers, click **Cancel**.
3. When all updates are completed, scroll through the Features window. This window reports the changes that the ROM Update Wizard applied to your controllers.
4. Leave the *ServeRAID Manager Support* CD in the CD-ROM drive; shut down and restart the server.

Selecting ROM update images

To select ROM update image files:

1. Click Add ; then, select the software image files from the file system.

   Note: To remove an image from the file list, select the file, then click Remove .
2. When you are ready to continue, click **Next**; the ″ Select controller″ window opens.

Managing SNMP devices

This topic describes the SNMP Browser and SNMP devices tasks in IBM Director.

Configuring SNMP trap forwarding

This topic describes how to configure SNMP trap forwarding in IBM Director.

You can forward SNMP traps in one of two ways: either through the Event Action Plan Builder or by configuring the SNMPServer.properties file. If you use the Event Action Plan Builder, IBM Director events are converted to SNMP traps.

In the Event Action Plan Builder, SNMP trap forwarding is configured by selecting one of the following events, then right-clicking and selecting Customize:

- Send an SNMP Trap to a NetView Host
- Send an SNMP Trap to an IP Host

Complete the following steps to forward SNMP traps without modification:

1. Using a text editor, edit a file named SNMPServer.properties in the IBM\Director\data\snmp directory.
2. To forward SNMPv1 traps:
   a. Remove the # character at the beginning of this line:
      ```
      #snmp.trap.v1.forward.address.1=
      ```
   b. Type the IP address of the SNMPv1 trap destination after the equal sign (=).
   c. Remove the # character at the beginning of this line:
      ```
      #snmp.trap.v1.forward.port.1=
      ```
   d. Type the port number of the SNMPv1 trap destination after the equal sign (=).
3. To forward SNMPv2 traps:
   a. Remove the # character at the beginning of this line:
      ```
      #snmp.trap.v2.forward.address.1=
      ```
   b. Type the IP address of the SNMPv2 trap destination after the equal sign (=).
   c. Remove the # character at the beginning of this line:
      ```
      #snmp.trap.v2.forward.port.1=
      ```
   d. Type the port number of the SNMPv2 trap destination after the equal sign (=).
4. **Optional:** To set a second or a third destination, edit the applicable lines in the SNMPServer.properties file.
5. Save the file.
6. Stop and restart IBM Director Server.

Note: Do not configure a trap destination that is sending traps to the management server, as you might create a loop.

Creating an SNMPv3 profile

This topic describes how to create an SNMPv3 profile in IBM Director.

Complete the following steps to create an SNMPv3 profile:

1. In IBM Director Console, click **Options** → **Server Preferences**.
2. In the Server Preferences window, click the **SNMP** tab.
3. Click **Add**.
4. In the Add Profile window, type the profile name.
5. Type the user name.
6. Select the authentication protocol.
7. If you select an authentication protocol other than the default **None**, type the password in both the **Password** and **Confirm Password** fields.
8. If you select a privacy protocol other than the default **None**, type the password in both the **Password** and **Confirm Password** fields.
9. Type a context name.
10. Type a context engine ID.
11. Click **Add**.

Managing MIB files

The MIB Management window allows you to select MIBs to compile when IBM Director Server starts. You can also select the MIBs you want to load into memory.

Compiling a MIB file

This topic describes how to compile a MIB file in IBM Director.

The SNMP Browser initially displays a tree view of the MIB structure for the selected SNMP devices. If no compiled MIBs are available on IBM Director Server to format the information, or if the device returns information that is not found in a compiled MIB, the information is displayed in a dotted-decimal numeric format. IBM Director includes various MIB files that are typically needed for SNMP browsing for commonly defined devices. These files are located in the Director\proddata\snmp directory. When IBM Director Server starts, it loads a default set of compiled MIBs that are commonly used. If you have updated the list of MIBs to load, your most recent selections are loaded.

MIB data is stored in its own persistent-storage file, *MIB definition name*.mibdata, in the Director\data\snmp directory. By deleting these files and snmpcompiledmibs.dat, you can remove all MIB data in IBM Director but retain other persistent-storage data.

Complete the following steps to compile a MIB file:

1. On the IBM Director Console window, click **Tasks** → **SNMP Browser** → **Manage MIBs**.
2. On the MIB Management window, click **File** → **Select MIB to Compile** .
3. In the Select MIB to Compile window, specify the directory and file name of the MIB file that you want to compile, and click **OK**. A status window indicates the progress of the compilation.

Selecting MIB files to load into memory

If you want to select specific MIBs to load into memory, use the SNMP browser task.

Complete the following tasks to select specific MIB files to load into memory:

1. On the IBM Director Console window, select **Tasks** → **SNMP Browser** → **Manage MIBs**.
2. On the MIB Management window, select **File** → **Select MIB to load**
3. From the **Available MIBs** list, select the MIBs that you want to load into memory.
4. Click **Add** to transfer the selected MIBs to the **Loaded MIBs** list.
5. Click **OK**.

Setting an attribute value

This topic describes how to set an attribute value in IBM Director.

You can set a user-defined value for an attribute that displays a 📝 icon. The community name on your SNMP device must also allow the change. Those attributes that are displaying a 📄 icon are read-only.

To set a value for an SNMP attribute, expand the tree and select a settable attribute. The current value is displayed in the Value pane. Type the new value and click **Set**.

Note: To discover SNMP devices, you must define an SNMP device to use as a discovery seed, or an SNMP service must be installed and running on the management server.

Setting discovery parameters

This topic describes how to set discovery parameters for SNMP devices in IBM Director.

Complete the following steps to set discovery parameters for SNMP devices:

1. In IBM Director Console, click **Options** → **Discovery Preferences**.
2. In the Discovery Preferences window, click the **SNMP Devices** tab.
3. Click **SNMP version** to select **SNMPv1**, **SNMPv2c**, or **SNMPv3**.
4. If you selected SNMPv1 or SNMPv2c, use the **Add**, **Import**, **Replace**, and **Remove** buttons to create your lists of IP addresses, corresponding subnet masks, and community names. If you selected SNMPv3, use the **Add**, **Import**, **Replace**, and **Remove** buttons to create your lists of IP addresses, corresponding subnet masks, and profile names.

Using the SNMP Browser

This topic describes how to use the SNMP Browser task in IBM Director.

You can use the SNMP Browser task to view and configure the attributes of SNMP devices, for example, hubs, routers, or other SNMP-compliant management devices. You can use the SNMP Browser for SNMP-based management, troubleshooting, or monitoring the performance of SNMP devices.

To start the SNMP Browser, in the IBM Director Console Tasks pane, drag the **SNMP Browser** task onto an SNMP device.

In the SNMP Browser window Device Information pane, expand the tree to view the SNMP information.

When you select an attribute in the Device Information pane, the right pane splits and displays the Value and Details panes. The Value pane displays the value of the selected attribute. The Details pane displays the characteristics of the selected attribute, including, for example, the type and access status of the device attribute and a description of the device attribute. If a snap-in is available for the selected attribute, it is displayed in the Selected Object pane in place of the default value and characteristics information.

Managing SMI-S storage devices

This topic describes tasks you can perform to manage SMI-S storage devices that have been discovered by IBM Director.

Adding a new SMI-S storage device

This topic describes how to add a new SMI-S storage device. You can choose to add a supported SMI-S storage device in a locked or unlocked state.

Complete the following steps to add a new SMI-S storage device.

1. In IBM Director Console, right-click in an empty area within the Group Contents pane to display the context menu. Select **New** → **SMI-S Storage Devices**. The Add SMI-S Storage Devices dialog appears.
2. Fill in the following required fields on the Add SMI-S Storage Devices dialog.
   a. In the **Name prefix** field, type the leading part of the name to assign to the SMI-S storage device. The trailing part of the name is determined by the Template Name specified in the Discovery Preferences window on the SMI-S Storage Devices page.
   b. In the **Network Address** field, type the IP Address or host name of the storage device’s SMI-S provider.
3. To unlock the storage device when you add it, select the **Advanced** check box and fill in the following fields. Note that if you do not select the **Advanced** check box, the storage device appears in the Group Contents pane in an unlocked state, and you will need to request access to unlock it.
   a. From the **Type** pull down list, select **SNIA:Array**, which indicates an array of storage (disk drives).
   b. In the **Namespace** field, type the name of the CIM root namespace in the SMI-S provider CIMOM. If not specified, all root namespaces known to IBM Director are tried. These are the values of the cim_namespace.*properties in the Director/data/Smis*.properties files.
   c. In the **User ID** field, type a user identification that satisfies the requirements of the SMI-S provider CIMOM.
   d. In the **Password** field, type a corresponding password that satisfies the requirements of the SMI-S provider CIMOM.
4. Click **OK**. The resulting SMI-S Storage Devices object appears in the Group Contents pane.

Launching Storage Manager for IBM DS storage array devices

IBM DS Storage Manager, which is an application external to IBM Director, is used to manage DS storage array devices (previously called FAStT storage array devices). This topic describes how to launch Storage Manager from the IBM Director Console for IBM DS storage array devices.

To launch Storage Manager, you use the Storage Manager command task provided in the External Application Launch task. However, before you can launch Storage Manager on a Windows server, the following environment variables must be set:

JAVA_FAStT
Defines the JRE directory associated with the Storage Manager product. The default installation on Windows is: ″C:\Program Files\Common Files\IBM_FAStT\jre\1.4″.

STORAGE_MANAGER
Defines the DS working directory. The default Windows installation is: ″C:\Program Files\IBM_FAStT\client″

There is no need to set environment variables for Storage Manager on a Linux system.

To start the Storage Manager Client, follow these steps:

1. From IBM Director Console, in the **Group Contents** pane, right-click on a storage managed object.
2. Click **External Application Launch** → **Storage Manager**.

For information about IBM DS Storage Manager, see the documentation that comes with that product.

Setting discovery preferences for SMI-S storage devices

This topic describes how to change discovery preferences for SMI-S storage devices.

Complete the following steps to set discovery preferences for SMI-S storage devices:

1. In IBM Director Console, click **Options** → **Discovery Preferences**.
2. In the Discovery Preferences window on the SMI-S Storage Devices page, select the preferences that you want IBM Director to use.
   a. In the **Auto-discover period (hours)** field, select the frequency with which IBM Director Server attempts to discover SMI-S storage devices automatically. The possible range is one to 168 hours (seven days). By default, autodiscovery is disabled.
   b. In the **Presence Check period (minutes)** field, select the frequency that IBM Director Server checks the status of each SMI-S storage device. A presence check detects whether a storage device is online or offline. The possible range is one to 240 minutes (four hours); this feature also can be disabled. By default, a presence check is set to occur every 15 minutes.
   c. Under **Service Location Protocol (SLP) Profiles**, select the SLP profiles that IBM Director Server will discover. To restore the original default parameters, select the **Reset to default value** check box.
   d. Under **Naming Conventions Template for SMI-S Devices**, add the parameters to include in the **Template Name** by selecting a parameter in the **Available Parameters** column, and then clicking **Add** to place it in the

Selected Parameters column. To remove a parameter from the template name, select a parameter in the **Selected Parameters** column, and then click **Remove**. To restore the original default parameters, select the **Reset to default value** check box.

3. To save your selections, click **OK**.

Unlocking SMI-S storage devices

This topic provides information about how to unlock SMI-S storage devices. An SMI-S storage device must be unlocked before IBM Director can communicate with the device and obtain Common Information Model (CIM) data from it.

To unlock an SMI-S storage device, complete the following steps:

1. Open IBM Director Console.
2. In the **Group Contents** pane, right-click the storage managed object you want to unlock.
3. Click **Request Access**.
4. Type a valid login ID and password that will enable access to the SMI-S provider for the storage managed object.

Viewing attributes for SMI-S storage devices

Attributes for SMI-S storage devices can include four sets of data values: Default, CIMOM, Base and Extended. This topic describes how to view these attributes.

To display the attributes for an SMI-S storage device, in the **Group Contents** pane, double-click on a storage managed object.

Managing digital-signature certificates

IBM Director allows you to monitor the expiration of digital-signature certificate for Level-1 managed systems and regenerate certificates that expire.

Configuring digital-signature certificate polling setting

This topic describes how to configure the how often IBM Director Server polls Level-1 managed systems for a valid digital-signature certificate.

Complete these steps to configuring the polling setting:

1. From a command line on the management server, open this file using a text editor:

   ```
   data/CertificateExpirationManager.properties
   ```
2. Change the **polling_interval** setting to the desired frequency, in seconds. The default is 86 400 (24 hours).
3. Save and close the file.

Configuring digital-signature certificate notification setting

This topic describes how to configure the notification setting that determines when IBM Director Server generates an event to notify you when a digital-signature certificate for a Level-1 managed system is about to expire.

Complete these steps to configuring the notification setting:

1. From a command line on the management server, open this file using a text editor:

data/CertificateExpirationManager.properties

2. Change the **advance_notify_in_hours** setting to the desired number of hours in advance that IBM Director Server is to generate a notification event. The default is 240.
3. Save and close the file.

Regenerating an expired digital-signature certificate

This topic describes how to regenerate an expired digital-signature certificate.

To regenerate an expired digital-signature certificate, use the **dircli certmgr –e** ***cert_sign_req*** command. This command generates a certificate-signing request file using the absolute path and filename specified by *cert_sign_req*. The generated certificate-signing request file is a standard base64-encoded file that can be sent to a third-party signer of security certificates.

For more information about this command, see “certmgr” on page 537.

Managing operating system accounts

The System Accounts task allows you to manage accounts for your operating system with functions such as deleting users, adding groups, and editing group membership.

Adding a group

The System Accounts task allows you to add an operating system group using IBM Director.

Complete the following steps to add an operating system group:

1. Drag the **System Accounts** task onto a managed system or group that supports System Accounts.
2. In the System Accounts window, click the **Groups** tab.
3. Click **Add**.
4. In the Group Configuration page, type the group name in the **Name** field.
5. **Optional:** In the **Description** field, type a description.
6. Click **Accept**.
7. Click **Apply**.

Deleting a user

The System Accounts task allows you to delete a user from the operating system using IBM Director.

Complete the following steps to delete a user:

1. Drag the **System Accounts** task onto a managed system or group that supports System Accounts.
2. In the System Accounts window, click the **Users** tab.
3. In the **Users** field, select a user name.
4. Click **Delete**. A window opens that displays the following message:
   ```
   Note: User not deleted until the apply button is clicked!
   ```
5. Click **Close** to close the window.
6. Click **Apply**.

Editing group membership

The System Accounts task allows you to edit group membership for your operating system accounts.

Complete the following steps to add a user to a group or to remove a user from an operating system group:

1. Drag the **System Accounts** task onto a managed system or group that supports System Accounts.
2. In the System Accounts window, click the **Groups** tab.
3. In the **Users** field, select the group you want to edit.
4. Click **Properties**.
5. If you are adding a user to the group, complete the following steps:
   a. In the **Non-members** field of the Group Configuration page, select a user name.
   b. Click < to move the selected non-member to the **Members** field.

   If you are removing a user from a group, complete the following steps:
   a. In the **Members** field, select a user name.
   b. Click > to move the selected member to the **Non-members** field.
6. Click **Accept**.
7. Click **Apply**.

Distributing software

This topic describes the Software Distribution task in IBM Director.

You can use the Software Distribution task to import applications and data, build a software package, and distribute the package to IBM Director managed systems: Level-0 managed systems, Level-1 managed systems, and Level-2 managed systems. There are two editions of software distribution: Standard and Premium. To use the Premium Edition, you must have purchased and installed IBM Director Software Distribution Premium Edition on the management server.

With IBM Director Software Distribution Standard Edition, you can import only software that is distributed by IBM and build a software package that uses only the IBM Update Assistant wizard. With the Premium Edition, you can:

- Import non-IBM or IBM software and build software packages that use the following wizards:
  - InstallShield Package wizard (Windows)
  - Microsoft Windows Installer wizard (Windows)
  - RPM Package wizard (AIX and Linux)
  - AIX InstallP wizard (AIX)
- Import non-IBM or IBM software and build a software package by using the Custom Package Editor
- Import a software package created in IBM Director by using the Previously Exported Package wizard
- Export a software package for use on another management server
- Restore i5/OS libraries, objects and installed programs

Note:

1. By default, Software Distribution uses TCP. If you disable TCP-session support on a managed system, Software Distribution uses UDP.
2. If you use the IBM Update Assistant wizard to create a package for the System Availability Agent for Linux, then distribute the package to a managed system, the dependency failure will not be detected by IBM Director.

Understanding software distribution

This topic describes how to distribute software in IBM Director.

You must follow three steps to distribute software packages to IBM Director managed systems:

1. Obtain the software.
2. Import the software into IBM Director Server and build a software package.
3. Distribute the software package to managed systems using one of the following methods:
   - Streamed distribution
   - Redirected distribution

   Note: Redirected distribution is only available for Level-2 managed systems.

Streamed distribution

A *streamed distribution* copies the software package from the management server to the managed system and then installs the software package onto the managed system.

When you use streaming to distribute a software package to a Level-2 managed system, if a network connection is lost during the transmission, IBM Director attempts to resume the distribution from the point at which the transmission was interrupted. If the streaming operation can be resumed, retransmission time is saved. Otherwise, the entire package must be sent again.

When you use streaming to distribute a software package to a Level-0 managed system, or a Level-1 managed system, if a network connection is broken during the transmission, the entire package must be sent again.

Redirected distribution

With *redirected distribution*, a file-distribution server called a *redirector share* functions as a storage location for a software package. The redirector share caches a software package. After a package has been cached on a redirector share, the cached package is used for future distributions, which can reduce the amount of time that is required to distribute a software package. A software package is cached on a redirector share only when the package is distributed.

The benefit of redirected distribution is the reduction of network congestion. With redirected distribution, the managed system receives only the minimum installation code that is needed to access the share and install the software from the management server.

Note: If the installation is interrupted, for example, if the connection is lost, the installation must be started again.

During a redirected distribution, IBM Director Server first determines which of its defined redirector shares that the managed systems to which the software package is being distributed can access. Then, IBM Director Server determines whether the software package is already cached on any of the mutually accessible redirector shares. If the package is not cached, IBM Director Server searches its list of shares to determine which share has enough free space to save the package.

To use redirected distribution, IBM Director must be set up to use a file-distribution server. You can use either an FTP-based share or a universal naming convention (UNC)-based share. See the *IBM Director Installation and Configuration Guide* for more information about setting up a share.

Note:

1. The redirector shares keep an archive of all redirected software packages. To avoid exceeding available space on the shares, you should periodically examine the shares and delete cached software packages that are no longer needed, by using the File Distribution Servers Manager.
2. Because a system account cannot write to a Microsoft network share, you cannot distribute software packages to a managed system that uses a network share. If a package is distributed to a folder on a Microsoft network share, the distribution fails, and the system log reports a lack of hard disk space. Modify the distribution to distribute to a local drive.

For software that uses Microsoft Windows Installer or InstallShield Professional as the installation utility, when you use the redirected distribution method, the software package is installed directly from the file-distribution server automatically. However, you can specify that the package stream from the file-distribution server by selecting the applicable check box in the “Distribution Preferences” window for a managed system or a group.

You must install the software packages by using the applicable wizard.

Importing software and building software packages

This topic describes how to import files and build a software package in IBM Director.

You can use the following wizards or the Custom Package Editor to import files and build a software package:

- IBM Update Assistant wizard
- InstallShield Package wizard
- Microsoft Windows Installer Package wizard
- RPM Package wizard
- AIX InstallP Package wizard
- i5/OS Restore Library Package wizard
- i5/OS Restore Licensed Program Package wizard
- i5/OS Restore Object Package wizard

You can import files and packages from the following hardware only:

- A universal naming convention (UNC)-based share
- A local hard disk drive of the management console
- A local hard disk drive of the management server

Creating software packages to distribute

This topic describes the IBM Update Assistant.

The IBM Update Assistant is a wizard that imports software that is distributed by IBM into IBM Director and creates the software package or packages for distribution. The IBM Update Assistant can import these types of software packages:

- Packages in the Solution Install format, an architecture that provides a universal way to package and distribute software
- Packages in the UpdateXpress for xSeries format

Solution Install packages enable software distribution to Level-0 managed systems, Level-1 managed systems, and Level-2 managed systems. Solution Install packages contain the following elements:

- Installable unit (IU) of software that is self-extracting, self-installing and runs in silent, unattended mode
- Installable unit deployment descriptor (IUDD) XML file, the PackagedIU.xml file that describes the dependencies of the package
- Artifact XML file that describes the installation instructions for the IU
- Optional: An XML file in xSeries format that can enable Software Health Check for the package

A Solution Install package may be in the form of a JAR file, a Zip file, or a Solution Install formatted directory structure that can be accessed through the *ComponentName*PackagedIU.xml file.

Note: The *ComponentName* prefix is optional.

UpdateXpress packages can only be distributed to Level-2 managed systems. UpdateXpress packages contain the following elements:

- The software-update file that is self-extracting, self-installing and runs in silent, unattended mode
- An XML file that describes the software-update file and how to install it

Complete the following steps to import the software and create one or more software packages:

1. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.
2. (Standard) in the Software Distribution Manager window, double-click IBM Update Assistant. (Premium) Expand the **Wizards** tree. Double-click **IBM Update Assistant**.
3. In the IBM Update Assistant wizard, specify whether the files are on the local management console or on the management server by clicking the applicable button.
4. Type the location of the XML, JAR, or Zip file that describes the software package or packages that you want to import, or click **Browse** to locate the file.
5. Select the **Make Category Private** check box to make the new category visible to only the IBM Director account that created it.
6. Click **Next**. If one software package is specified in the XML file, the package is displayed in the Packages pane and it is selected for import into IBM Director by default.

   If more than one software package is specified, a tree structure is displayed in the Packages pane. For example, for UpdateXpress, a folder is displayed for each managed-system machine type that is specified in the XML file. Expanding each folder displays a list of the software packages that apply to the specific managed system machine type. If you click a package in the Packages pane, a description of the software package is displayed in the Details pane. By default, no software packages are selected for import into IBM Director, which is indicated by the red X beside each package in the Packages pane.

7. Double-click the package or packages in the Packages pane to select the package you want to import. If you want to select all of the packages, or just those that are deemed critical by IBM, you can right-click the folder and click either **Select All Items** or **Select Critical Items**. The red X beside a package in the Packages pane changes to a green check mark to indicate that the package will be imported.
   a. (Managed systems running Windows only) In the Options pane, you can specify an alternative installation script to run by typing the path name in the **Alternate install script** field. Options are not displayed when you are working with Server Plus Pack software packages for managed systems running Windows.
   b. (Upgrading IBM Director Agent on managed systems running Linux only) In the Options pane, you can specify an alternative installation script to run by typing the path name in the **Alternate response file/script** field.
   c. (Managed systems running AIX only) In the Options pane, you can specify an alternative installation script to run by typing the path name in the **Alternate response file/script** field.
8. Click **Finish**.

If you import only one software package, the package name is displayed in the IBM Director Console Tasks pane under the **All Software Distribution Packages** category. If you import more than one software package, a software-distribution category is created for each selected software package ; the category name is the name of the folder the packages resided in when displayed in the IBM Update Assistant wizard. Individual software packages are displayed under each category. The packages also are displayed in the IBM Director Console Tasks pane under **All Software Distribution Packages**.

Note: If you want to change the contents of a software-distribution category, use the category editor.

You can distribute the software package or software-package category that contains the packages that you want to distribute now, or you can set a time to distribute the software package or software-package category using Scheduler. Packages in the software-package category are distributed in the order listed in the category editor.

Creating InstallShield software packages

This topic describes how to use the InstallShield Package wizard in IBM Director.

Use this wizard to import the software and build a software package for an application that uses InstallShield Professional as its installation program. You can create packages for applications that use InstallShield Professional 5, 6, or 7 for Windows. InstallShield Professional requires a response file during installation to allow and perform an unattended installation. You can create a response file either by recording an installation or by using an editor. Note that you can distribute a software package that is created with this wizard to managed systems running Windows only.

Most applications do not indicate anywhere in the documentation that they use InstallShield Professional as their installation program. To determine whether an application uses InstallShield Professional, start the installation EXE file (usually setup.exe). When the first window opens (which is the standard InstallShield Setup window), minimize that window; then, right-click the taskbar, and click **About**.

If you see the word *InstallShield* in this window, use the InstallShield Package wizard in the Software Distribution task to build a software package.

Next, determine whether a response file is included with the software that you want to distribute. To determine whether a response file is included with the software for which you want to build a package, search for an ISS file (typically setup.iss). The response file is plain-text format so that you can edit the response file for use in your specific environment. If a response file is included, you must test the response file to make sure that it can be used to install the software on each managed-system type you intend to use it with, and that any customizations that you make are correct. If no response file is included, you must create a response file and test it.

Note: Many software products are not designed for unattended installation, although InstallShield provides the capability. Contact the product vendor if the software does not support unattended installation.

If no response file is included, record one by using the installation command for the software, typically setup.exe or install.exe. For example:

```
setup -r -f1x:\response_filename.iss -f2logfile
```

where:

- *setup* is the installation command for the product.
- *x:\response_filename* is the path where you want to save the response file. If you do not specify the –f1 parameter, InstallShield saves the response file in c:\windows\setup.iss.
- *logfile* is the path where you want to save the installation log file. If you do not specify the –f2 parameter, InstallShield does not create an installation log file.

When the installation command runs, you are prompted for required information. The responses that you provide must reflect how you want the application to be installed on the managed system. For more information about response files, go to http://www.InstallShield.com.

When you build the response file, you also install the software locally. Before you can test the response file, you must uninstall the software. After you uninstall the software, test the recorded response file or the response file that is included with the software. Type the following command:

```
setup -s -f1x:\response_filename.iss -f2logfile
```

where:

- *setup* is the installation command for the product.
- *x:\response_filename* is the path of the response file that you recorded or the response file that is included with the software. If you do not specify the –f1 parameter, InstallShield assumes that the response file is in c:\windows\setup.iss.
- *logfile* is the path where you want to save the log file. If you do not specify the –f2 parameter, InstallShield does not create a log file.

When the command is completed, check the system log file. If the software was installed successfully, the result code is 0. If the software was not installed successfully, you cannot distribute it by using IBM Director.

Complete the following steps to import the software and create a software package:

1. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.
2. In the Software Distribution Manager window, expand the **Wizards** tree. Double-click **InstallShield Package**.
3. In the InstallShield Package wizard, specify whether the files are on the local management console or on the management server by clicking the applicable button. Then, type the location of the setup program and the response file in the applicable fields, or click **Browse** to locate the setup program or response file. Click **Next**.
4. The package name is entered in the **Package Name** field automatically. If you want to use a different name, type the package name.
5. **Optional:** Specify additional command-line parameters that are specific to the application that you are importing by typing the applicable command-line parameters.
6. **Optional:** To install the software under a different user name and password, click **Advanced**. Type the applicable information and click **OK**.
7. Click **Finish**. Individual software packages are displayed under the **All Software Distribution Packages** category.

You can distribute the software package now or schedule a later time for distribution.

Creating Microsoft Windows Installer software packages

This topic describes how, in IBM Director, to import the software and build a software package for an application that uses Microsoft Windows Installer.

Use this wizard to import the software and build a software package for an application that uses Microsoft Windows Installer as its installation program.

Note: To determine whether an application uses Windows Installer technology, search for an MSI file in the top level of the application directory.

Using this wizard, you can change some installation parameters and use a Microsoft software transformation (MST) file. You can use this wizard to build software packages for distribution only to managed systems running Windows.

Complete the following steps to import the software and create a software package or packages:

1. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.
2. In the Software Distribution Manager window, expand the **Wizards** tree. Double-click **Microsoft Windows Installer Package**.
3. In the Microsoft Windows Installer Package wizard, in the **Package Name** field, type the package name.
4. Specify whether the files are on the local management console or on the management server by clicking the applicable button. Then, type the location of

the program file, or click **Browse** to locate it. Click the applicable button to install or uninstall the software package. Click **Next**.

5. **Optional:** Specify a Microsoft software transformation (MST) file by typing the location of the transform file in the applicable field or clicking **Browse** to locate it. You can also specify additional Windows Installer parameters by typing the parameters in the applicable field.
6. **Optional:** To install the software under a different user name and password, click **Advanced**.
7. In the Advanced Options window, type the user ID and password in the applicable fields and click **OK**.
8. Click **Next**. A summary is displayed.
9. Click **Finish**. The package name is displayed in the IBM Director Console Tasks pane under the **All Software Distribution Packages** category.

You can distribute the software package now or schedule a later time for distribution.

Creating an RPM software package

This topic describes how to use the RPM Package wizard in IBM Director.

Use the RPM Package wizard to import the software and build a software package for an application that uses Red Hat Package Manager (RPM) for its installation program. The RPM program is the common installer for all IBM Director-supported Linux operating systems. An RPM is an archive of files that are specific to an application. Using this wizard, you can create and distribute a single software package that contains one or more RPMs. You can use this wizard to build RPM software packages for distribution only to managed systems running Linux or AIX.

Complete the following steps to create an RPM software package:

1. Complete the following steps to import the software and create a software package:
   a. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.
   b. In the Software Distribution Manager window, expand the **Wizards** tree. Double-click **RPM Package**.
   c. In the **Package Name** field, type the package name.
   d. In the **Target OS** field, select **Linux** or **AIX**.
   e. Select **Install** to install the software package.
   f. Click **Next**.
   g. Specify whether the files are on the local management console or on the management server by clicking the applicable button. Then, select the RPMs that you want to import by clicking **Add**. A separate window opens in which you can select the files that you want to import. You can select more than one RPM to import at a time.
   h. Click **Finish**. The software package name is displayed in the IBM Director Console Tasks pane under the **All Software Distribution Packages** category.

   You can distribute the software package now or schedule a later time for distribution.
2. Complete the following steps to uninstall a software package:
   a. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.

b. In the Software Distribution Manager window, expand the **Wizards** tree. Double-click **RPM Package**.

c. In the **Package Name** field, type the package name.

d. In the **Target OS** field, select **Linux** or **AIX**.

e. Click **Uninstall** to uninstall the software package.

f. Click **Next**.

g. Select the RPM that you want to uninstall by clicking **Add** and entering the RPM name.

h. Click **Finish**.

Creating an AIX InstallP software package

This topic describes how to use the AIX InstallP Package wizard in IBM Director.

Use the AIX InstallP Package wizard to install an AIX InstallP format package.

Complete the following steps to import the software and create a software package:

1. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.
2. In the Software Distribution Manager window, expand the **Wizards** tree. Double-click **AIX InstallP Package** to start the AIX InstallP Package wizard.
3. In the **Package Name** field, type the package name.
4. Select **Install** to install the software package.
5. **Optional:** Select the **Verify install or uninstall** check box.
6. Click **Next**.
7. Specify whether the files are on the local management console or on the management server by clicking the applicable button. Then, select the InstallP filesets that you want to import by clicking **Add**. A separate window opens in which you can select the files that you want to import. You can select more than one fileset to import at a time.
8. Click **Finish**.

Creating an i5/OS Restore Library Package

This topic describes how to use the i5/OS Restore Library Package wizard in IBM Director.

An i5/OS library is an object that contains other i5/OS objects in the file system. Use the i5/OS Restore Library Package wizard to build a package to restore a library to a managed system running i5/OS.

Complete the following steps to create an i5/OS Restore Library Package:

1. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.
2. In the Software Distribution Manager window, expand the **Wizards** tree. Double-click **i5/OS Restore Library Package**.
3. In the **Package Name** field, type the package name.
4. Specify whether the files are on the local management console or on the management server by clicking the applicable button.
5. Type the name of the stream file that contains the library or click **Browse** to locate it.
6. Type the name of the library that you want to restore from the stream file.

7. **Optional:** Type any additional library restore parameters.
8. **Optional:** Click **Advanced**.
9. In the Advanced Options window, type the user ID and password; then click **OK**.
10. Click **Next**.
11. Click **Finish**.

Creating an i5/OS Restore Licensed Program Package

This topic describes how to use the i5/OS Restore Licensed Program Package wizard in IBM Director.

Use the i5/OS Restore Licensed Program Package wizard to build a package to restore a program to a managed system running i5/OS.

1. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.
2. In the Software Distribution Manager window, expand the **Wizards** tree. Double-click **i5/OS Restore Licensed Program Package**.
3. In the **Package Name** field, type the package name.
4. Specify whether the files are on the local management console or on the management server by clicking the applicable button.
5. Type the name of the stream file that contains the licensed program or click **Browse** to locate it.
6. Type the name of the licensed program that you want to restore from the stream file.
7. Click **Next**.
8. Click **Finish**.

Creating an i5/OS Restore Object Package

This topic describes how to use the i5/OS Restore Object Package wizard in IBM Director.

Use the i5/OS Restore Object Package wizard to build a package to restore an object to a managed system running i5/OS.

1. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.
2. In the Software Distribution Manager window, expand the **Wizards** tree. Double-click **i5/OS Restore Object Package**.
3. In the i5/OS Restore Object Package wizard, in the **Package Name** field, type the package name.
4. Specify whether the files are on the local management console or on the management server by clicking the applicable button.
5. Type the name of the stream file that contains the object or click **Browse** to locate it.
6. Type the name of the object to restore from the stream file.
7. Type the name of the objects.
8. Type the name of the library that contains the objects.
9. Type the object types.
10. Click **Next**.
11. **Optional:** Type any additional object parameters.
12. Click **Finish**.

Creating a custom software package

This topic describes how to use the Custom Package Editor in IBM Director.

Use the Custom Package Editor to import the software and build a software package without using a wizard. You can specify the files, target directory names and paths, and installation programs or batch files that perform the software installation.

Complete the following steps to import and build a software package:

1. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.
2. In the Software Distribution Manager window, double-click **Custom Package Editor**.
3. On the General page, specify the package name and any distribution options and prerequisites.
4. On the Files page, specify the files to use by navigating to each file in the Source File System pane and clicking **Add**. You can change whether the files are displayed from the local management console or from the management server by selecting from the list at the top of the pane.

   Note: If you want to include all subdirectories in a parent directory, select the **Include subfolders** check box, or, if maintaining the file structure is important, select the **Save full path information** check box. Then, select the directory and click **Add**.
5. Specify distribution of a software package to a managed system running an AIX, Linux, i5/OS, or Windows operating system by selecting the applicable check box on the applicable page.
6. **Optional:** In the Execute Pre-Distribution pane, click **Advanced**.
7. **Optional:** In the Pre-Distribution window, select an application to run on the managed system before the software distribution occurs. You can select multiple applications and specify the order in which they run. When complete, click **OK**.
8. **Optional:** In the Execute Pre-Distribution pane, select the **File exists on target system** check box if the specified application exists on the managed system.
9. **Optional:** (i5/OS only) In the Execute Pre-Distribution pane, click **Native** or **QShell** to select how the application is to be run.
10. **Optional:** In the Execute Post-Distribution pane, click **Advanced**
11. **Optional:** In the Post-Distribution window, select an application to run on the managed system before the software distribution occurs. You can select multiple applications and specify the order in which they run. When complete, click **OK**.
12. **Optional:** (i5/OS only) In the Execute Post-Distribution pane, click **Native** or **QShell** to select how the application is to be run.
13. **Optional:** (AIX, Linux, and i5/OS only) Click **File Permissions**.
14. **Optional:** (AIX, Linux, and i5/OS only) In the File Permissions window, set the file and directory permissions for the software distribution and click **OK**.

    Note: (Managed systems running Linux only) Files that are copied to a managed system running Linux are set with the default permission or with the account that IBM Director Agent runs as, which is the root account. Use the File Permissions window to set alternative permissions for the software distribution.

15. **Optional:** (Windows only) Click **Do Nothing** or **Restart Computer** to select whether to restart the managed system after the software distribution is completed.
16. **Optional:** (Windows only) Click **Windows NT/2000/XP/2003 Configuration**.
17. **Optional:** (Windows only) In the Windows NT/2000/XP/2003 Configuration window, specify that changes in Windows system files, INI files, and Registry keys are to be distributed to the managed system and click **OK**.
18. Click **OK**. The package name is displayed in the IBM Director Console Tasks pane under the **All Software Distribution Packages** category.

You can distribute the software package now or schedule a later time for distribution.

Creating a package for delayed delivery

If you want to create a package during off-peak hours to be distributed at a later time, you can preload the software package on a file-distribution server.

Normally, software packages are copied to the file-distribution server only when you initiate a software distribution. If you want to preload a file-distribution server, but wait to distribute the package, you can create a package for delayed delivery.

1. In the Tasks pane of IBM Director Console expand **Software Distribution** → **All Software Distribution Packages**.
2. Right-click the software package that you want to preload to the file-distribution server and select **Create Package for Delayed Delivery**.
3. If you want to copy the package to a local disk on the IBM Director Server, select **Copy to Disk** on the Create Package for Delayed Delivery window. The **Copy to Disk** option can save network bandwidth by allowing you to transfer the software package to a file-distribution server during peak hours using any method you choose, independent of IBM Director; you then can distribute it to managed systems at a time that is appropriate for your environment. If you want to copy the package directly to one of the file-distribution servers defined for the IBM Director Server, select **Copy to File Distribution Server** on the Create Package for Delayed Delivery window. Then, click **Next**.
4. If you selected the **Copy to Disk** option, select the location where you want to save the package on the IBM Director Server. If you selected the **Copy to File Distribution Server** option, select the file-distribution server that you want to preload with the software package. Then, click **Finish**.
5. If you selected the **Copy to Disk** option, you must transport and install the package information that was saved to disk to the file-distribution server from which you want to distribute the package. If you have not done so already, you must update IBM Director Server preferences to add the file-distribution server and direct the managed system to use this file-distribution server.

Creating and editing software-package categories

This topic describes how to create a new software-package category in IBM Director.

You can use the software-package category function in Software Distribution to create new categories of software packages or to edit existing categories of software packages.

Complete the following steps to create a new software-package category:

1. In the IBM Director Console Tasks pane, right-click the **Software Distribution** task and click **New Package Category**.
2. In the **Category Name** field, type a category name.
3. In the Available Packages pane, click a package; then, click **Add**. The order in which the software packages are displayed in the Selected Packages pane specifies the order of delivery when that category is distributed. To modify the order in which software packages are delivered, in the Package Names column, select a package; then, drag the package to its new location.
4. **Optional:** Set the managed system to restart after delivery of a specific software package by selecting the **Reboot** check box for that package in the Selected Packages pane.
5. **Optional:** (Windows only) To restart the managed system after all software packages in that category are delivered, select the **Reboot at end of Category Distribution** check box.
6. **Optional:** Select the **Make Category Private** check box to make the new category visible to only the IBM Director account that created it.
7. Click **OK** to save the new software-package category.

Distributing a software package

This topic describes how to distribute a software package or software-package category in IBM Director.

You can distribute a software package or software-package category immediately or schedule a later time for distribution.

Notes:

1. Group-distribution preferences and individual managed-system distribution preferences are independent of each other. That is, when you distribute a software package to a group, the group-distribution preferences apply to all the managed systems within the group. If you distribute a software package to an individual managed system, the managed-system distribution preferences apply.
2. If you distribute a software-package category to a group of managed systems, each software package within that category is delivered individually to each managed system in the group. The package that is listed first in the category is the first to be distributed ; use the category editor to see how the software packages are ordered for distribution. After the first package has been distributed, each succeeding package is delivered to each managed system until all software packages have been distributed.

Complete the following steps to distribute a software package or software-package category:

1. In the IBM Director Console Tasks pane, drag the software package or software-package category onto the managed system or group to which you want to distribute the package.
2. Click **Execute Now**, or click **Schedule** to schedule the distribution for a later time.

Editing a software package

This topic describes how to edit an existing software package in IBM Director.

Complete the following steps to edit a software package:

1. In IBM Director Console, click **Tasks** → **Software Distribution** → **All Software Distribution Packages** → *Software Package* → **Edit**.
2. Make your changes in the package editor.

When you attempt to open a package, you might receive a message indicating that the package is locked by another process. This means that another user is editing the package, it is being copied to a file-distribution server, or it is in the process of being distributed to one or more managed systems. The package remains locked until the other process is completed. However, it is possible for a package to remain locked when no process or user is using it. For example, if a system was turned off while a package was being edited, the package will remain locked for 5 to 10 minutes.

Editing software-package categories

This topic describes how to edit a software-package category in IBM Director.

Complete the following steps to edit an existing software-package category:

1. In the IBM Director Console Tasks pane, expand the **Software Distribution** task.
2. Right-click the package category that you want to edit and click **Open**.
3. To add a package to the category, click a package in the Available Packages pane of the Edit Package Category window, and then click **Add**. To delete a software package from the category, right-click a software package in the Selected Packages pane, and then click **Remove**. The order in which the software packages are displayed in the Selected Packages pane specifies the order of delivery when that category is distributed. To modify the order in which software packages are delivered, select a package in the Package Names column, and then drag the package to its new location.
4. **Optional:** To specify that the managed system is to restart after delivery of a specific software package, select the **Reboot** check box for that package in the Selected Packages pane. Or, to restart the managed system after all software packages in that category are delivered, select the **Reboot at end of Category Distribution** check box.
5. **Optional:** Select the **Make Category Private** check box to make the new category visible to only the IBM Director account that created it.
6. Click **OK** to save any changes that you made to an existing category.

Exporting a software package (Premium Edition)

This topic describes how to export a software package in IBM Director.

If you have IBM Director Software Distribution Premium Edition, you can export a software package for use on another management server or to back up a software package. Software packages are exported in SPB format and can be imported into Tivoli Configuration Manager.

Note:

- Exporting a software package is not supported when IBM Director Server is installed on a server running i5/OS.

- Software Distribution does not support exporting packages to a network share. If a package is exported to a network share, the export fails, and the following error message is displayed: `Unable to export package`. Modify the export to export to a local drive.

Complete the following steps to export a software package:

1. Right-click a software package and click **Export**.
2. In the **File Name** field of the Export Software Distribution window, type a file name and click **Save**.

Importing a previously created software package (Premium Edition only)

The Previously Exported Package wizard imports software package block (SPB) format files into IBM Director. Create these files by exporting an IBM Director software package. If you want to import a software package that was created in IBM Director, you must use this wizard.

Note: You cannot use the Previously Exported Package wizard to import SPB files that were created by Tivoli software or signed package (BFP) format software packages that were created with IBM Director version 3.1 or earlier.

You can distribute the software package immediately or schedule a later time for distribution.

Complete the following steps to import a software package:

1. In the IBM Director Console Tasks pane, double-click the **Software Distribution** task.
2. In the Software Distribution Manager window, expand the **Wizards** tree. Double-click **Previously Exported Package**.
3. In the Previously Exported Package wizard, specify whether the files are on the local management console or on the management server by clicking the applicable button. Then, type the location of the SPB file, or click **Browse** to locate it.
4. Click **Next**.
5. Click **Finish**. The package name is displayed in the IBM Director Console Tasks pane under the **All Software Distribution Packages** category.

Maintaining software health

The Software Health Check function of the Software Distribution task allows you to check whether your managed systems need certain updates.

The following topics provide more information about using Software Health Check:

Checking for managed systems that need updates

If you want to locate managed systems that need updates for IBM Director Agent or for xSeries drivers or firmware, you can use the Software Health Check to find the managed systems.

After you have imported a package using IBM Update Assistant, you can use Software Health Check to locate the managed systems that require the update.

Note: The XML file associated with the package must be enabled for Software Health Check.

Complete the following steps to perform a software health check for a selected package:

1. In the **Tasks** pane of the IBM Director Console , expand **Software Distribution** → **All Software Distribution Packages**.
2. Right-click the package or category for which you want to locate managed systems that need the update and select **Perform Health Check**.

   Note: If **Perform Health Check** is not available to select, Software Health Check is not supported for the software package.

   When the Software Health Check is complete, the Software Health Check Log window shows the managed systems that require the update, as well as the current version of the installed software. A dynamic group that contains all of the managed systems that need the update is also created. The group name is the title of the software package, preceded by *Health:*. For example, the group name for an update titled, *IBM Hard Disk Drive Update Program (Windows)* is *Health: IBM Hard Disk Drive Update Program (Windows)*. If a managed system that needs the update comes online after Software Health Check is performed, the dynamic group is automatically updated.

   Note: If Software Health Check is performed on a software-package category, each package in the category is checked. A single dynamic group is created that contains each managed system that requires an update for any of the packages in the category.

Updating out-of-date managed systems

If you want to update out-of-date managed systems identified by running Software Health Check, you can distribute the update package to the corresponding dynamic group that was generated by Software Health Check.

Complete the following steps to distribute an update to managed systems that were identified by Software Health Check:

1. In the **Groups** pane of IBM Director Console, expand **Software Health Groups** and locate the group that was generated for the update package by Software Health Check. The group is titled: **Health:** *Update Package Name*.
2. In the **Tasks** pane of IBM Director Console, expand **Software Distribution** → **All Software Distribution Packages** and locate the update software package or software-package category that you want to distribute.
3. From the **Tasks** pane, drag the software package or software-package category onto the Health group with the corresponding title.
4. Click **Execute Now**, or click **Schedule** to schedule the distribution for a later time.

After the software is distributed and inventory is collected, the dynamic group will change to reflect the current status of the software update. If the package update on each managed system is successful, then after inventory is collected, no systems will appear in the dynamic group. You can perform a custom collection on the following inventory items to update the dynamic group:

- Device Drivers
- Director Systems
- Firmware
- Operating System

- Physical Enclosure

You can delete the dynamic group when you no longer need it.

Viewing health check reports

If you want to see the managed systems that require a particular update, you can view the Health Check Report for that software package or software-package category.

After Software Health Check is run for a particular software update, you can view the Health Check Report to see which managed systems require the update and the version of the software that is currently installed. The Health Check Report also displays the severity level of the required update. Complete the following steps to view the Health Check Report

1. In the **Tasks** pane of IBM Director Console, expand **Software Distribution** → **All Software Distribution Packages**.
2. Right-click the update software package or software-package category and select **Health Check Report**.

   Note: The report reflects the last time Software Health Check was performed. If you want an updated report, rerun Software Health Check. The dynamic group will not be recreated if it still exists, but a new report will be generated.

Restricting software-package access

This topic describes how to restrict access to a software package in IBM Director.

You can restrict access to a software package by specifying a user name and password combination that a user must type to gain access to the package.

To enable this option, follow these steps:

1. Right-click the package, and then click **Secure Package**.
2. Type a user ID and password for the user that you want to allow to modify the package, and then click **OK**.

Viewing software-package contents

This topic describes how to view the contents of a software package in IBM Director.

You can view the contents of a software package, including the package files, the managed-system type for which the package was created, and whether a restart on the target system is set to occur after package installation.

To view the contents of a package, follow these steps:

1. In the IBM Director Console Tasks pane, expand the **Software Distribution** task.
2. Right-click the package for which you want to see the contents, and then click **View Package Summary**.

Viewing the software-distribution history for a software package

This topic describes how to view the distribution history for a selected software package in IBM Director.

Complete the following steps to view the distribution history for a selected software package:

1. In the IBM Director Console Tasks pane, expand the **Software Distribution** task to view the list of software packages.
2. Right-click the software package for which you want to view the history, and click **View Distribution History**.

Viewing software-package creation and distribution status

This topic describes how to view software-package creation and distribution status in IBM Director.

Using the Package Audit Log, you can determine the status of software-package creation and distribution. Three levels of detail are provided to assist you in tracking and troubleshooting.

To access the log, in the IBM Director Console Tasks pane, right-click the **Software Distribution** task and click **Package Audit Log**.

Viewing details about file-distribution servers and software packages

This topic describes how to view details about file-distribution servers and software packages in IBM Director.

Using the File Distribution Servers Manager, you can view details about file-distribution servers and the software packages that are stored on a file-distribution server.

To access the File Distribution Servers Manager, follow these steps:

1. In the IBM Director Console Tasks pane, right-click the **Software Distribution** task.
2. Click **Manage File Distribution Servers**.

The Software Packages group box displays the software packages that are stored on the selected file-distribution server. In the **File Distribution Server Details** group box, Maximum Managed Systems indicates the maximum number of managed systems that can access the file-distribution server at one time.

You can perform several tasks in the File Distribution Servers Manager window as follows:

- To view the file-distribution maintenance log, click **File** → **Maintenance Log**.
- To test access to the file-distribution servers, click **Actions** → **Test Access to All File Distribution Servers**. To test access to an individual file-distribution server, click the file-distribution server in the **File Distribution Servers** group box; then, click **Actions** → **Test Access to Selected File Distribution Server(s)**.
- To refresh a software package from the file-distribution server, click the package in the **Software Packages** group box; then, click **Actions** → **Refresh Package on File Distribution Server**.
- To delete a software package from the file-distribution server, click the package in the **Software Packages** group box; then, click **Actions** → **Remove Package from File Distribution Server**.

Changing software-distribution server preferences

This topic describes how to change your software-distribution server preferences in IBM Director.

You can change your software-distribution server preferences, such as the maximum number of managed systems on which streaming can occur concurrently, streaming bandwidth, and redirected distribution options.

Complete the following steps to change your software-distribution server preferences:

1. In IBM Director Console, click **Options** → **Server Preferences**.
2. In the Server Preferences window, click the **Software Distribution** tab.
3. Change the applicable selections.
4. Click **OK**.

Changing software-distribution server preferences for a managed system or group

This topic describes how to change your software-distribution server preferences for a managed system or group in IBM Director.

You can change your software-distribution server preferences, such as the maximum number of managed systems on which streaming can occur concurrently, streaming bandwidth, and redirected distribution options.

Complete the following steps to change your software-distribution preferences for a managed system or group:

1. In the IBM Director Console Group Contents pane, right-click a managed system or group of managed systems, and then click **Distribution Preferences**.
2. In the Distribution Preferences window, to stream the software-distribution package from the management server to the managed system, select **Always stream to system(s)**. Or, to stream the software-distribution package from a shared directory (share) to the managed system, select **Use File Distribution Server Shares**. If you select **Use File Distribution Server Shares**, you can select specific shares to use from the shares that are defined on the File Distribution Servers page of the Server Preferences window. Click **Add** to select a share.
3. Select the **Restrict share selection to list** check box to use only the shares that are listed in the **Shares** field.

   Note: If this check box is selected and the managed system is unable to connect to any of the defined shares, the software distribution will fail.
4. Select the **Stream from File Distribution Server** check box to copy the contents of a software-distribution package to the managed system before installation.
5. Click **OK** to update the distribution preferences for the managed system or group.

Deploying z/VM virtual servers

This topic provides information about using the Virtual Server Deployment task.

Starting the Virtual Server Deployment task

This topic describes how to start working with the Virtual Server Deployment task.

Before you start: z/VM Center and its subtask Virtual Server Deployment are not necessarily included in each IBM Director setup. Depending on the platform where your IBM Director Server is installed, you must install z/VM Center as an optional feature or as an extension of IBM Director.

To be able to use z/VM Center, you must also install a license key. For more details see the z/VM Center setup information in the IBM Director information center athttp://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/vsd0_t_install_key.html

You always open the Virtual Server Deployment task for the scope of a specific z/VM system. Starting from the IBM Director Console, complete these steps:

1. Locate the z/VM System with which you want to work in the Group Contents pane.

   Tip: Select z/VM Systems in the Groups pane to only display z/VM Systems in the Group Contents pane.

   If the z/VM System you want to work with is not shown in the Group Contents pane, you might need to run a managed object discovery. Complete these steps to discover z/VM systems:

   a. In the Group Contents pane, right-click into an empty space.

   b. Click **Discover** → **z/VM Systems**.
2. If you are not logged on to the z/VM system with which you want to work, you need to log on now. z/VM systems that you are not logged on to, are marked with a padlock icon (🔒). Complete these steps to log on:

   a. Right-click the z/VM system; then click **Request access...**.

   b. Specify the Linux administrator ID and password for the z/VM manageability access point. By default, this is root.
3. Right-click the z/VM system; then click **z/VM Center** → **z/VM Virtual Server Deployment**.

Note: Alternatively, you can also start the Virtual Server Deployment task by:

- Expanding the z/VM Center task in the Task pane and dragging the z/VM Virtual Server Deployment task onto the z/VM system in the Group Contents pane.
- Dragging the z/VM system in the Group Contents pane onto the z/VM Virtual Server Deployment task in the Tasks pane.

Working with the z/VM system

This topic lists the tasks that affect the entire z/VM system.

Setting up a profile for the z/VM system

Before z/VM Center can work with a z/VM system, you need to set up a profile that includes access credentials for working with the systems management API.

Before you start: z/VM Center and its subtask Virtual Server Deployment are not necessarily included in each IBM Director setup. Depending on the platform where your IBM Director Server is installed, you need to install z/VM Center as an optional feature or as an extension of IBM Director.

Complete the following steps to set up a profile:

1. Start the IBM Director Console.
2. Locate the z/VM System in the Group Contents pane.

   Tip: Select z/VM Systems in the Groups pane to only display z/VM systems in the Group Contents pane.

   If the z/VM System you want to work with is not shown in the Group Contents pane, you might need to run a managed object discovery. Complete these steps to discover z/VM systems:

   a. In the Group Contents pane, right-click into an empty space.
   b. Click **Discover → z/VM Systems**.

3. If you are not logged on to the z/VM system you want to work with, you need to log on now. z/VM systems that you are not logged on to, are marked with a padlock icon (🔒). Complete these steps to log on:
   a. Right-click the z/VM system; then click **Request access...**
   b. Specify the Linux administrator ID and password for the z/VM manageability access point. By default, this is root.
4. Right-click the z/VM system; then click **z/VM Center → z/VM Virtual Server Deployment**.
5. In the z/VM Virtual Server Deployment widow, expand the top node of the z/VM System resource tree. The top node represents the z/VM system.
6. Click the z/VM Profile node below the top node. The z/VM Profile pane is displayed with a prompt to specify the access credentials for the z/VM systems management API.
7. Click **Configure**.
8. In the Set Access Credentials wizard, specify the access credentials.
   a. Type the z/VM manageability access point administrator ID in the **User ID** field. This is the z/VM user ID that you have authorized when setting up your z/VM for z/VM Center. For more information refer to the following topic in the IBM Director information center publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/vsd0_t_prepare_map_auth.html
   b. Type the password in the **Password** field.
   c. Type the IP address with which you access the z/VM systems management API in the **VSMSERVE Server** field. This is usually the IP address of the z/VM TCPIP service machine. Be sure to type the correct address as defined on z/VM. The IP address must be in IPv4 dotted-decimal format.
   d. Click **Finish**.
9. From the **Disk Pool** list, select the disk pool you want to be the default disk pool.
10. Click **Save**.

Providing access credentials for the z/VM systems management API

You configure the access credentials for the z/VM systems management API on the z/VM system. If your access credentials for the z/VM system are incomplete or no longer valid, only a limited subset of the z/VM Center functions are available to you.

Complete the following steps to provide the access credentials to z/VM Center:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the top node. The top node represents the z/VM system.

3. Click the z/VM Profile node below the top node. If your current access credentials are complete and valid, the z/VM Profile pane is displayed with the profile data. If your credentials are incomplete or no longer valid, you are informed that you must update the credentials.
4. In the z/VM Profile pane, click **Configure**.
5. In the Set Access Credentials wizard, make the changes according to your z/VM configuration.
   a. If changed or missing, type an administrator user ID that is authorized for the z/VM systems management API in the **User ID** field. The user ID has to be authorized for the z/VM systems management API and requires at least a subset of the privileges of the default IBM user classes B and E.
   b. If changed or missing, type the current password in the **Password** field.
   c. If changed or missing, type the IP address with which you access the z/VM systems management API in the **VSMSERVE Server** field. This is usually the IP address of the z/VM TCP/IP service machine. Be sure to type the correct address as defined on z/VM. The IP address must be in IPv4 format.

      The IP Address field is separated into four sections according to the dotted-decimal format. You can edit each section separately. Double-click a section to start editing it.

      Tip: Double-click the first section to edit it. When you are done with editing a section, press the period (.) key to progress to the next section of the IP address. Press Enter when you have completed the address.
6. Click **Finish**.

Viewing the z/VM system configuration

This topic describes how to view the configuration of the z/VM system that hosts your z/VM virtual servers.

Complete the following steps to view the configuration of your z/VM system:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, select the top node if it is not already selected.

Editing the description of the z/VM system

You can edit a description of the z/VM system. All other z/VM system information that you can view in the z/VM system is obtained from z/VM and cannot be changed from z/VM Center.

Complete the following steps to edit the description of your z/VM system:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, select the top node.
3. In the **Description** field, type your description of the z/VM system. You can change or extend an existing description. You can click **Refresh** to discard any changes you have made and restore the description that was last saved.
4. Click **Save**. Your changes do not take effect until you click **Save**.

Editing the profile for working with the z/VM systems management API

The profile includes the IP address and credentials for the z/VM systems management API of the z/VM system. It also includes the default for the disk pool from which the storage requirements for operating system instances and templates are satisfied.

The access data that you provide in this task are used by z/VM Center to access the z/VM systems management API. You cannot change this data from z/VM Center. Specify the data as defined on the z/VM system.

Complete the following steps to edit a profile:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the top node. The top node represents the z/VM system.
3. Click the z/VM Profile node below the top node. If your current access credentials are complete and valid, the z/VM Profile pane is displayed with the profile data. If your credentials are incomplete or no longer valid, you are informed that you must update the credentials.
4. **Optional:** In the z/VM Profile pane, click **Configure**.
5. **Optional:** In the Set Access Credentials wizard, specify the access credentials.
   a. If changed or missing, type an administrator user ID that is authorized for the z/VM systems management API in the **User ID** field. The user ID has to be authorized for the z/VM systems management API and requires at least a subset of the privileges of the default IBM user classes B and E.
   b. If changed or missing, type the current password in the **Password** field.
   c. If changed or missing, type the IP address with which you access the z/VM systems management API in the **VSMSERVE Server** field. This is usually the IP address of the z/VM TCP/IP service machine. Be sure to type the correct address as defined on z/VM. The IP address must be in IPv4 format.
   d. Click **Finish**.
6. **Optional:** In the **Disk Pool** field select the pool you want to be the default disk pool.
7. **Optional:** In the **Description** field, you can provide a description of the profile. You can edit or extend and existing description.
8. Click **Save**.

Working with z/VM virtual servers

This topic lists the tasks that you can perform with z/VM virtual servers.

Creating a new z/VM virtual server

This topic tells you how to create a new z/VM virtual server from a virtual server template.

Before you start: To create a new z/VM virtual server you need a virtual server template. To view a list of existing virtual server templates in the right pane, select the Virtual Server Templates node in the Provisioning Resources tree.

Complete the following steps to create a new z/VM virtual server:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.

2. In the z/VM System navigation tree, select the z/VM Virtual Servers node. The existing z/VM virtual servers are listed in the right pane.
3. In the right pane, click **New** to start the Create z/VM Virtual Server wizard.
4. Follow the instructions in the wizard.

Viewing a z/VM virtual server configuration

This topic describes how to find and display information on a particular z/VM virtual server.

Complete the following steps to view the configuration of a z/VM virtual server:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the z/VM Virtual Servers node to display a list of z/VM virtual servers.
3. Select the z/VM virtual server with which you want to work. A notebook with the configuration of the selected z/VM virtual server is displayed in the right pane.
4. Click the tabs to navigate among the notebook pages.

   The **Disks**, **Processors**, and **Network Ports** pages might have more than one disk, processor, or port for which you can view configuration data. These pages have a list of the respective resource on the left. Click the item for which you want to view configuration data. All other fields on the page show the data for the disk, processor, or port you have selected in the list on the left.

Activating and deactivating a z/VM virtual server

When you activate a z/VM virtual server, z/VM provides the hardware resources according to its definition. If z/VM Center is aware of an operating system on the z/VM virtual server, the guest operating system is booted. You can think of activating or deactivating a z/VM virtual server as powering on or off the virtual hardware.

You can activate inactive z/VM virtual servers and deactivate active z/VM virtual servers. Deactivating gracefully shuts down the operating system. If the shutdown is not completed within a given time interval, the operating system is stopped by force.

Be sure that the time interval is set such as to allow for a regular shutdown for a functioning operating system. You can query this time interval on z/VM by issuing the QUERY SIGNAL SHUTDOWN control program command. You can use the SET SIGNAL SHUTDOWN command to change the time interval if necessary. See *z/VM CP Command and Utility Reference*, SC24-6081 for details.

Complete these steps to activate or deactivate a z/VM virtual server:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the z/VM Virtual Servers node to display a list of z/VM virtual servers.
3. Select the z/VM virtual server you want to activate or deactivate. The z/VM virtual server configuration is displayed in the right pane. The icon in the **Status** field indicates whether the z/VM virtual server is active or inactive.
4. Proceed according to the action you want to perform:

- If you want to start an inactive z/VM virtual server, select and right-click the z/VM virtual server; then click **Activate**. If z/VM Center is aware of an operating system on the z/VM virtual server **Activate** also boots this operating system.
- If you want to deactivate an active z/VM virtual server, select and right-click the z/VM virtual server; then click **Deactivate**. You will be prompted to confirm that you want to deactivate the z/VM virtual server.

Editing parameters of a z/VM virtual server

You can change the names that z/VM Center uses for a z/VM virtual server and the names it uses for its disks. You can also edit descriptions for the z/VM virtual server and its disks, processors, memory, and ports.

You determine most of the parameters of a z/VM virtual server in the virtual server template it is based on. After the virtual server has been created, all but the names and descriptions of the z/VM virtual server and its resources are fixed. None of the changes described in this task result in changes to parameters on z/VM.

Complete the following steps to edit names and descriptions of a z/VM virtual server and its resources:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the z/VM Virtual Servers node to display a list of z/VM virtual servers.
3. Select the z/VM virtual server you want to work with. A notebook with the configuration of the selected z/VM virtual server is displayed in the right pane.
4. **Optional:** If you want to change the name or description of the z/VM virtual server, complete the following steps on the **Overview** page:
   a. Optional: In the **z/VM Virtual Server** field, type the new name you want to use for the z/VM virtual server.
   b. Optional: In the **Description** field, type your description of the z/VM virtual server. You can change or extend an existing description.
5. **Optional:** If you want to change the name or description of one of the disks that are defined for the z/VM virtual server, click the **Disks** tab; then select the disk for which you want to change the name or description in the **Disks** field. You can change the name or description separately for each disk.
   a. Optional: In the **Name** field, type the new name you want to use for the disk.
   b. Optional: In the **Description** field , type your description of the disk. You can change or extend an existing description.
6. **Optional:** If you want to change the description of one of the processors that are defined for the z/VM virtual server, click the **Processors** tab; then select the processor for which you want to change the description in the **Processors** field. You can change the description separately for each processor. You can change or extend an existing description.
7. **Optional:** If you want to change the description of the memory that is defined for the z/VM virtual server, click the **Memory** tab; then type your description of the memory in the **Description** field. You can change or extend an existing description.
8. **Optional:** If you want to change description of one of the ports defined for the z/VM virtual server, click the **Network ports** tab; then select the port for which

you want to change the description in the **Ports** field. You can change the description separately for each port. You can change or extend an existing description.

9. Click **Save**. Your changes do not take effect until you click **Save**.

Note: You can click **Save** at any time to save changes you have made. You can also click **Refresh** to discard any changes you have not already saved and refresh the notebook with the current data.

Deleting a z/VM virtual server

When you delete a z/VM virtual server you remove all knowledge of it and any installed operating system from z/VM Center. Deleting frees the disks and other resources that are used by a z/VM virtual server for redeployment.

Complete these steps to delete a z/VM virtual server:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the z/VM Virtual Servers node to display a list of z/VM virtual servers.
3. Select and right-click the z/VM virtual server you want to delete; then click **Delete**. You will be prompted to confirm that you want to delete the z/VM virtual server.

Working with operating system instances

This topic lists the tasks that you can perform with operating system instances.

Creating a new operating system instance in a z/VM virtual server

This task describes how to create an operating system instance in an existing z/VM virtual server, using operating system templates.

Before you start:

- The z/VM virtual server must already exist.
- You need an operating system template.
- z/VM Center can only manage a single operating system for each z/VM virtual server. You cannot perform this task for a z/VM virtual server for which z/VM Center is already aware of an operating system. You need to de-register an existing operating system instance before you can create a new operating system instance with z/VM Center.

Complete the following steps to create an operating system instance:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System navigation tree, expand the z/VM Virtual Servers node to display the existing z/VM virtual servers.
3. Expand the node that represents the z/VM virtual server on which you want to create an operating system instance.
4. Select the Operating Systems node below the z/VM virtual server node.

   If an operating system instance on the z/VM virtual server is already known to z/VM Center, it is shown in the right pane. In this case, you need to de-register the known operating system instance before you can continue with this task.
5. Click **New** to start the Create Operating System wizard.

6. Follow the instructions in the wizard.

Viewing the configuration of an operating system instance

The only configuration data of an operating system instance that you can change within z/VM Center are its description and names and descriptions for its resources. All other data you can view only. You use the Re-Register Operating System wizard to make z/VM Center aware of changes to an operating system instance that have been applied outside z/VM Center.

Complete the following steps to view the configuration of an operating system instance:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the z/VM Virtual Servers node to display a list of z/VM virtual servers.
3. Expand the node for the z/VM virtual server you want to work with; then expand the Operating Systems node below it. If there is no operating system instance below the Operating Systems node, none has been installed on the z/VM virtual server or an operating system has been installed outside the Virtual Server Deployment task and needs to be registered.
4. Select the operating system instance to display a notebook with the configuration of the operating system instance in the right pane.
5. Click the tabs to navigate among the notebook pages.

   The **Disks**, and **Network Ports** pages might have more than one disk or port for which you can view configuration data. These pages have a list of the respective resource on the left. Click the item for which you want to view configuration data. All other fields on the page show the data for the disk or port you have selected in the list on the left.

Registering an operating system instance

If you have installed an operating system instance outside z/VM Center, you need to register it to enable z/VM Center to work with it.

Before you start:

- z/VM Center can only manage a single operating system for each z/VM virtual server. You cannot perform this task for a z/VM virtual server for which z/VM Center is already aware of an operating system. You need to de-register an existing operating system before you can register a new operating system.
- To be able to perform this task you need to know:
  - Which disks hold the operating system instance, the applications installed on it, and other data that is closely related to the operating system
  - From which disk the operating system is booted
  - Where applicable, the IP addresses you want to use for any network interfaces
  - The host name used by the operating system instance
  - The port used by each network interface. If you want to use the operating system instance to be registered as a source for operating system templates, you must provide a complete mapping of network interface and port for each network interface that is used by your operating system instance.

z/VM Center needs some basic information about an operating system instance to be able to work with it. It automatically gathers this information when you install an operating system instance using an operating system template. If you install an operating system instance outside z/VM Center, you need to register it to supply the required information. An operating system instance that you install outside z/VM Center is not represented in the z/VM Center task window until you register it.

Be aware that registering does not make any changes to the operating system configuration data. Be sure to enter the correct configuration data, according to the actual operating system configuration.

Complete the following steps to register an operating system instance:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the z/VM Virtual Servers node to display a list of z/VM virtual servers.
3. Expand the node for the z/VM virtual server on which your operating system instance has been installed.
4. Select and right-click the Operating Systems node below the expanded z/VM virtual server node; then click **Register Operating System** to start the Register Operating System wizard.

   If an operating system instance on the z/VM virtual server is already known to the Virtual Server Deployment, this menu item is not selectable.
5. Follow the instructions in the Register Operating System wizard.

Result: The operating system instance is known to z/VM Center, and you can work with it in the same way as operating system instances that have been installed from an operating system template.

Keeping an operating system configuration up-to-date (re-register)

If you change an operating system instance outside z/VM Center, you need to inform z/VM Center about the changes.

z/VM Center cannot detect configuration changes you make to an operating system instance. If you want to use an operating system instance as a source for an operating system template, z/VM Center needs up-to-date configuration data. The data that needs to be up-to-date corresponds to the data that is created when an operating system instance is registered.

Complete the following steps to re-register an operating system instance:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the z/VM Virtual Servers node to display a list of z/VM virtual servers.
3. Expand the node for the z/VM virtual server you want to work with; then expand the Operating Systems node below it.
4. **Optional:** Expand the node for the z/VM virtual server on which the operating system instance that you want to re-register is installed. If there is no operating system instance below the Operating Systems node, none has been installed on the z/VM virtual server or the operating system has been installed outside the Virtual Server Deployment task and has not been registered.
5. Select and right-click the Operating Systems node; then click **Re-Register Operating System** to start the Re-Register Operating System wizard.

6. Make updates that describe your configuration changes as you proceed through the Re-Register Operating System wizard.

After re-registering an operating system instance that has been installed using an operating system template, the re-registered operating system instance is no longer considered to have been derived from the template.

Editing descriptions and names related to an operating system instance

Within z/VM Center, you cannot make changes to the configuration of an operating system instance. This task describes how to change the descriptions for the operating system instance and its resources. It also describes how to change the names that z/VM Center uses for the resources of the operating system instance. Use the Re-Register Operating System wizard to make z/VM Center aware of configuration changes that have taken place outside z/VM Center.

Complete the following steps to edit names and descriptions of an operating system instance and its resources:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the z/VM Virtual Servers node to display a list of z/VM virtual servers.
3. Expand the node for the z/VM virtual server you want to work with; then expand the Operating Systems node below it. If there is no operating system instance below the Operating Systems node, none has been installed on the z/VM virtual server or an operating system has been installed outside the Virtual Server Deployment task and needs to be registered.
4. Select the operating system instance. A notebook with the configuration of the operating system instance is displayed in the right pane.
5. **Optional:** If you want to change the description of the operating system instance, click the **Overview** tab; then type your description in the **Description** field. You can change or extend an existing description.
6. **Optional:** If you want to change the name or description of one of the disks that hold the operating system, installed applications, or other data closely associated with the operating system, click the **Disks** tab; then select the disk for which you want to change the name or description in the **Disks** field. You can change the name or description separately for each disk.
   a. Optional: In the **Name** field, type the new name you want to use for the disk.
   b. Optional: In the **Description** field , type your description of the disk. You can change or extend an existing description.
7. **Optional:** If you want to change description of one of the ports defined for the operating system instance, click the **Network ports** tab; then select the port for which you want to change the description in the **Ports** field. You can change the description separately for each port. You can change or extend an existing description.
8. Click **Save**. Your changes do not take effect until you click Save.

Note: You can click **Save** at any time to save changes you have made. You can also click **Refresh** to discard any changes you have not already saved and refresh the notebook with the current data.

De-registering an operating system instance

When you de-register an operating system instance, you remove all knowledge of it from z/VM Center.

De-registering does not de-install the operating system nor free the disks where the operating system instance and related data reside for redeployment. To free the resources you must delete the z/VM virtual server. You need to de-register an operating system instance before you can create or register a new operating system instance.

Complete these steps to de-register an operating system instance:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the z/VM System resource tree, expand the z/VM Virtual Servers node to display a list of z/VM virtual servers.
3. Expand the node for the z/VM virtual server with the operating system instance that you want to de-register; then expand the Operating Systems node below it.
4. Select and right-click the operating system instance you want to de-register; then click **De-Register Operating System**. You are prompted to confirm that you want to de-register the operating system instance.

Working with templates

This topic lists the tasks that you can perform with virtual server templates and with operating system templates.

There are two types of templates within z/VM Center:

Virtual server templates
are templates on which z/VM virtual servers are based. Most of the parameters of a z/VM virtual server are determined by the parameters of the virtual server template on which it is based. After the z/VM virtual server has been created these parameters are fixed.

Operating system templates
are templates on which operating system instances are based. An operating system template is created from an especially prepared instance of an operating system, the *master system*. To z/VM, an operating system template is a guest virtual machine. Once an operating system template is in place, it can be used to create operating system instances.

Creating a new virtual server template

Virtual server templates are required for creating z/VM virtual servers.

Before you start: To perform this task you need to know the following about the z/VM virtual servers to be based on the template:

- You need to know the user class that defines the privileges of the z/VM virtual servers on z/VM.
- You need to know which prototype, if any, to use to base your z/VM virtual servers on.
- You need to know how many virtual processors to assign to the z/VM virtual servers.

- You need to know the initial virtual memory to be assigned to the z/VM virtual server and the maximum to which the virtual memory can be extended through a z/VM control program DEFINE STORAGE command.
- To z/VM, the z/VM virtual servers that are based on the virtual server template are guest virtual machines. You need a naming scheme for the z/VM user IDs of these guest virtual machines.

Tip: To create a virtual server template that is similar to an already existing template, you can copy the existing template and then edit the copy.

Complete the following steps to create a new virtual server template:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources tree, select the Virtual Server Templates node. The existing z/VM virtual server templates are listed in the right pane.
3. In the right pane, click **New**. The Create Virtual Server Template is started that guides you through creating a new virtual server template.
4. Follow the instructions in the wizard.

Viewing a virtual server template

Use the Virtual Server Template pane to view parameters for a virtual server template.

Complete the following steps to view a virtual server template:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources tree, expand the Virtual Server Templates node to list the available virtual server templates.
3. Select the template you want to work with. A notebook with the template configuration is displayed in the right pane.
4. Use the tabs to navigate among the notebook pages.

Editing parameters of a virtual server template

You can make changes to the parameters and defaults for the z/VM virtual servers that will be based on the template. You can also change the name and description of the template.

Complete the following steps to edit a virtual server template:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources tree, expand the Virtual Server Templates node to list the available virtual server templates.
3. Select the template you want to use. A notebook with the template configuration is displayed in the right pane.
4. Use the tabs to navigate among the notebook pages and type your changes in the fields. You can click **Refresh** to discard any changes you have made and refresh the notebook with the values that were last saved.
5. Click **Save**. Your changes do not take effect until you click **Save**.

Copying a virtual server template

Copying an existing virtual server template can be an efficient way to obtain a new virtual server template that differs only slightly from the original.

Complete these steps to delete a virtual server template:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources tree, expand the Virtual Server Templates node to display a list of the available virtual server templates.
3. Select and right-click the virtual server template you want to copy; then click **Copy**.
4. Follow the instructions in the “Create copy of virtual server template” wizard.

When you have completed the wizard, you can edit your new copy to make it different from the original.

Deleting a virtual server template

You can delete a virtual server template if you no longer need it. Deleting removes its configuration data from z/VM Center.

Complete these steps to delete a virtual server template:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources tree, expand the Virtual Server Templates node to display a list of the available virtual server templates.
3. Select and right-click the virtual server template you want to delete; then click **Delete**. You will be prompted to confirm that you want to delete the template.

Creating a new operating system template

You use operating system templates to create operating system instances on z/VM virtual servers. This topic describes how to create an operating system template using the Create Operating System Template wizard.

Before you start: Operating system templates are created from existing operating system instances that are known to z/VM Center and that have been made into master instances. To perform this task you need to know:

- Which z/VM virtual server the master instance is installed on and when this z/VM virtual server can be shut down safely
- Which disks you want to be shared by all operating system instances that are based on the new template
- Which disk pool you want to use to satisfy the storage needs of the template and of the operating system instances that are based on it

Complete the following steps to create a new operating system template:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources resource tree, click the Operating System Templates node. The available operating system templates are listed in the right pane.
3. In the right pane, click **New** to start the Create Operating System Template wizard.
4. Follow the instructions in the wizard.

Viewing an operating system template

Use the Operating System Template pane to view parameters for an operating system template.

Complete the following steps to view an operating system template:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources tree, expand the Operating System Templates node. A list of the available operating system templates is displayed in the tree.
3. Select the template you want to work with. A notebook with the template configuration is displayed in the right pane.
4. Use the tabs to navigate among the notebook pages.

Editing parameters of an operating system template

Most of the parameters of an operating system template are inherited from the master operating system on which the operating system template is based and cannot be changed in the template. This task describes how to change names and descriptions that z/VM Center uses for an operating system template and its resources. None of the changes described in this task result in changes to parameters on z/VM.

Complete the following steps to edit names and descriptions of an operating system template and its resources:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources tree, expand the Operating System Templates node to display a list of the available operating system templates.
3. Select the operating system template for which you want to change names and descriptions. A notebook with the operating system template data is displayed in the right pane.
4. **Optional:** If you want to change the name or description of the operating system template complete these steps on the **Overview** page.
   a. Optional: In the **Name** field, type the new name you want to use for the operating system template.
   b. Optional: In the **Description** field , type your description of the operating system template. You can change or extend an existing description.
5. **Optional:** If you want to change the name or description of one of the disks that are defined for the operating system template, click the **Exclusive Disks** or the **Shared Disks** tab; then select the disk for which you want to change the name or description in the **Disks** field. You can change the name or description separately for each disk.
   a. Optional: In the **Name** field, type the new name you want to use for the disk.
   b. Optional: In the **Description** field, type your description of the disk. You can change or extend an existing description.

   The **Exclusive Disks** or the **Shared Disks** tabs are selectable only if there are exclusive or shared disks defined for the template.
6. **Optional:** If you want to change the description of one of the ports defined for the operating system template, click the **Network ports** tab; then select the port for which you want to change the description in the **Ports** field. You can change the description separately for each port. You can change or extend an existing description.
7. Click **Save**. Your changes do not take effect until you click **Save**.

Note: You can click **Save** at any time to save changes you have made. You can also click **Refresh** to discard any changes you have not already saved and refresh the notebook with the current data.

Deleting an operating system template

You can delete an operating system template if you no longer need it. Deleting the template removes its configuration data from z/VM Center and frees the disk space it had used for redeployment.

Complete these steps to delete an operating system template:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources tree, expand the Operating System Templates node to display a list of the available operating system templates.
3. Select and right-click the operating system template you want to delete; then click **Delete**. You are prompted to confirm that you want to delete the template.

Working with disk pools

This topic lists the tasks that you can perform with disk pools.

Viewing disk pool parameters

This topic describes how to find and display information on a particular disk pool.

Complete the following steps to view the configuration of a disk pool:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources tree, expand the Disk Pools node to display a list of available disk pools.
3. Select the disk pool of interest.

Editing parameters of a disk pool

You can change the name and description of a disk pool.

Complete the following steps to edit disk pool parameters:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.
2. In the Provisioning Resources tree, expand the Disk Pools node to display a list of available disk pools.
3. Select the disk pool of interest.
4. **Optional:** In the **Name** field, type the new name you want to use for the disk pool. Changing the name does not change any parameters on z/VM.
5. **Optional:** In the **Description** field, type your description of the disk pool. You can change or extend an existing description.
6. Click **Save**.

Changing the default disk pool

The default disk pool is the default provider of disk resources for creating z/VM virtual servers and for templates.

Complete the following steps to change the default disk pool:

1. If it is not started already, start the Virtual Server Deployment task for the z/VM System you want to work with.

2. In the z/VM System resource tree, expand the top node. The top node represents the z/VM system.
3. Click the z/VM Profile node below the top node. If your current access credentials are complete and valid, the z/VM Profile pane is displayed with the profile data. If your credentials are incomplete or no longer valid, follow the instructions on the screen to update you credentials before you continue with the next step.
4. From the **Disk Pool** list, select the disk pool you want to be the default disk pool.
5. Click **Save**.

Preparing a master Linux system

This topic describes how to prepare a Linux instance to serve as a master operating system. A master operating system is required for creating operating system templates.

Before you start:

- You need access to the z/VM MAINT user ID or an alternative user ID that is authorized to issue commands for your directory manager.
- The starting point for your master Linux system can be an instance of:
  - SUSE LINUX Enterprise Server 9 for IBM System z9, zSeries and S/390
  - Red Hat Enterprise Linux AS, version 4.0, for IBM System z9, zSeries and S/390

  The Linux instance must be installed on a z/VM virtual server and can be:
  - A Linux instance that has been created from an operating system template
  - A Linux instance that has been installed outside z/VM Center
- The boot partition for the Linux to be prepared as a master Linux must be the first partition on that DASD that provides a bootmap file at any of the following locations:
  - /bootmap
  - /boot/bootmap
  - /zipl/bootmap
  - /boot/zipl/bootmap
- You need access to a user with root authority on the Linux system.

Complete these steps outside z/VM Center to make a Linux instance into a master Linux system:

1. Establish a terminal session with the Linux instance that you want to prepare as a master Linux. You must log in as a user with root authority.
2. If you want to manage the derived Linux instances with IBM Director, install the appropriate IBM Director Agent, if it has not been installed already.
3. Check if the personalization RPM has been installed on your Linux instance. By default, the personalization RPM is installed with IBM Director Agent. Issue:

   ```
   rpm -qa | grep zVMPersonalization
   ```

   If the command returns a line with `zVMPersonalization` and some other information, the RPM is already installed and you can skip step 4.
4. If it has not already been installed, install the personalization RPM.

   You can find the RPM on your IBM Director Server at *installation_directory*/proddata/zVMCenter/ zVMPersonalization.s390*ver*.rpm, where *installation_directory* is the directory to which you have installed IBM Director Server and *ver* is a string that identifies the version of the RPM.

a. Transfer the RPM to a directory of your choice on the Linux instance that you want to prepare as a master Linux, for example with FTP.

b. Issue an **rpm** command to install the RPM.

5. Install the CPINT RPM. This RPM is shipped with SUSE LINUX Enterprise Server 9.

   For Red Hat Enterprise Linux AS, you can download the RPM from linuxvm.org/Patches/. You need version 2.5.3 or later.

   Note: Be aware that installing the RPM on Red Hat Linux AS might affect any support contract you may have for the distribution.

6. Change the access mode for all disks that you want to be shared to read-only. Be sure not to make any disks read-only that your Linux systems will need to write to. Change the access mode from both the Linux system and the z/VM virtual server where the Linux system is installed. For example, you can change the access mode like this:

   a. On z/VM, establish a session with user ID MAINT.

   b. For each disk that you want to be read-only, issue a command like this:

   ```
   dirmaint for userid mdisk devno rr
   ```

   where *userid* is the z/VM user ID of the guest virtual machine where the Linux is installed and *devno* is the virtual device number used to access the disk.

   c. On Linux, edit the /etc/fstab file. Ensure that the comma separated mount options in the 4th field of each line that represents a read-only file includes "ro". **Example:** To set a disk with device node /dev/dasdc1 read-only change

   ```
   /dev/dasdc1          /tools          ext2    defaults       0 0
   ```

   to

   ```
   /dev/dasdc1          /tools          ext2    defaults,ro    0 0
   ```

   Read-only disks are considered shared disks in the operating system templates that are based on the master Linux. z/VM Center provides a single physical copy of a shared disks for sharing by all Linux instances that are based on the same operating system template. Disk sharing can reduce your disk space requirements and is most effective if you carefully design your Linux file system, separating read-only data on separated disks from read-write data.

7. Establish all ports and network connections that you want to have for the operating system templates (and, consequently, the Linux instances) that are derived from the master Linux system. You need to configure both, the z/VM guest virtual machine and Linux.

8. **Optional:** Ensure that /etc/inittab includes a statement of this form:

   ```
   ca::ctrlaltdel:/sbin/shutdown -h now [message]
   ```

   where *message* is a message you want to be issued to users when a shutdown is initiated. If there is no such statement, Linux instances cannot be shut down gracefully but are stopped by force when deactivated from the IBM Director Console or directly from z/VM.

For more details, refer to the z/VM Center installation and conceptual information in the IBM Director information center at http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/vsd0_t_prepare.html and http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/vsd0_c_concepts.html

If the Linux instance was installed outside z/VM Center, you need to register it before you can use it to create an operating system template.

Auditing z/VM Center information

To audit z/VM Center information you need to ensure that auditing is enabled and that the required auditing categories are selected.

Complete the following steps to ensure that z/VM Center information is audited:

1. From IBM Director Console select **Options** → **Auditing Administration**.
2. On the IBM Director Server Auditing Administration window, ensure that the **Enable auditing** check box is selected.
3. To audit Virtual Server Deployment information, ensure that "CIM" is included in the **Selected categories** list.
4. To also audit Server Complexes information, ensure that both "CIM" and "Remote command execution" are included in the **Selected categories** list.

Managing z/VM server complexes

Starting the Server Complexes task

You can activate the z/VM Server Complexes task on a single z/VM object.

Before you begin: z/VM Center and its Server Complexes subtask are not necessarily included in each IBM Director setup. Depending on the platform where your IBM Director Server is installed, you may need to install the z/VM Center as an optional feature or as an extension of IBM Director.

You can only open the Server Complexes task for the a specific z/VM system. Starting from the IBM Director Console, perform these steps:

1. Locate the z/VM System you want to work with in the Group Contents pane.

   Tip: Select **z/VM Systems** in the Groups pane so it only displays **z/VM Systems** in the Group Contents pane.
2. If the z/VM System you want to work with is not shown in the Group Contents pane, you might need to run a managed object discovery. Perform these steps to discover z/VM system managed objects:
   a. In the Group Contents pane, right-click an empty space below the listed z/VM systems.
   b. Click **All Systems and Devices** → **Discover Systems** → **z/VM Systems**.
3. If you are not logged on to the z/VM system you want to work with, you need to log on now. The z/VM systems you are not logged onto are marked with a padlock icon. Perform these steps to log on:
   a. Right-click the z/VM system, then click **Request Access**.
   b. In the Request Access window, specify your credentials for the Manageability Access Point on that z/VM.
4. Right-click the z/VM system, then click **z/VM Center** → **Server Complexes**.

Note: Alternatively, you can also start the Server Complexes task by:

- Expanding the z/VM Center task in the Task pane and dragging the z/VM Server Complexes task onto the z/VM system in the Group Contents pane, or
- Dragging the z/VM system in the Group Contents pane onto the z/VM Server Complexes task in the Tasks pane.

Setting up server complexes

In the Server Complexes task, you can manage server complexes. You can create new server complexes, duplicate existing ones, or delete them.

Creating a server complex instance

Within a managed z/VM, you can create as many server complexes as you need.

Complete the following steps to create a new server complex instance:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. From the **Action** menu, choose **New server complex** to open the Create a Server Complex window.

   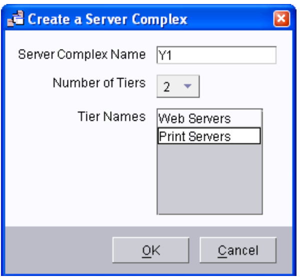

   Server Complex Name: Y1
   Number of Tiers: 2
   Tier Names: Web Servers, Print Servers
3. In the **Server complex name** field, type a unique name with one to ten alphanumeric characters. Avoid non-alphanumeric characters such as %, _, *, etc.
4. From the **Number of tiers** drop down list, choose between one and four tiers. The default is 1.
5. (Optional) Type a name for each tier in the **Tier names** field. These names are used in various places as a convenient reminder of the tier's meaning.
6. Click **OK**.
7. The Progress Indication window shows the progress of the action. If there is an inconsistency, you are given a choice whether to proceed or stop. Stopping saves the new properties but does not apply the configuration. After the operation is completed, click **Close** to close the Progress Indication window.

The new server complex is created and shown (with empty tiers) on the Server Complexes pane to the right of the z/VM Server Complexes window.

Note: If there is already an existing server complex with the desired name, an error message appears and the operation is cancelled.

Duplicating a server complex instance

When you duplicate a server complex instance, you create a new unpopulated server complex instance with the same property values, apart from the name.

Complete the following steps to duplicate a server complex:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select the server complex you want to duplicate.

3. From the **Action** menu, choose **Duplicate** to open the Duplicate Server Complex window. This window is similar to the New Server Complex window, except that the names and number of tiers cannot be changed.
4. In the **Server complex name** field, type a unique name with up to ten alphanumeric characters. Do not use non-alphanumeric characters such as *, &, $, etc.
5. Click **OK**.
6. The Progress Indication window shows the progress of the action. If there is an inconsistency, you are given a choice whether to proceed or stop. Stopping saves the new properties but does not apply the configuration. After the operation is completed, click **Close** to close the Progress Indication window.

The new duplicated server complex is created and displayed (with empty tiers) on the Server Complexes pane to the right of the z/VM Server Complexes window.

Note: If there is already an existing server complex with the desired name, an error message appears and the operation is cancelled.

Note: Alternatively, you can right-click the server complex and choose **Duplicate** from the pop-up menu.

Deleting a server complex instance

When you delete a server complex instance, you also detach all included Linux guest systems.

By deleting a server complex instance, you:

- Detach each included Linux guest system and move them into the Free Linux Guests pane
- Detach each included virtual server
- Delete the IBM Director server complex managed object

Complete the following steps to delete a server complex:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select the server complex you want to delete.
3. From the **Action** menu, choose **Delete**. A warning message is displayed, asking you to confirm the deletion.
4. Click **Yes** to confirm that you want to delete the server complex.
5. The Progress Indication window shows the progress of the action. If there is an inconsistency, you are given a choice whether to proceed or stop. Stopping saves the new properties but does not apply the configuration. After the operation is completed, click **Close** to close the Progress Indication window.

Note: Alternatively, you can right-click the server complex and choose **Delete** from the pop-up menu.

The server complex is deleted and any included Linux guest systems are moved to the Free Linux Guests pane.

Setting up virtual networking properties

Before you start configuring the server complexes, there are a number of tasks you need to perform to ensure that the z/VM networking properties are set up correctly.

Setting connection properties for Linux guest systems

Before you can set the network interfaces for server complexes, you must set up the connection properties for guest LANs, OSAs, or VSWITCHes in the z/VM. These will be used to route the Linux guest systems.

After you have set these connection properties, you can configure the network interfaces for a server complex.

Complete the following steps to set the connection properties:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Click **Action→ z/VM Networking Properties** to open the z/VM Networking Properties window.

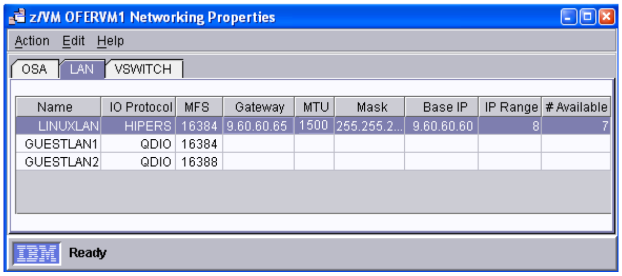

Name	IO Protocol	MFS	Gateway	MTU	Mask	Base IP	IP Range	# Available
LINUXLAN	HIPERS	16384	9.60.60.65	1500	255.255.2...	9.60.60.60	8	7
GUESTLAN1	QDIO	16384						
GUESTLAN2	QDIO	16388						

Ready

3. If this is the first time you are using the Server Complexes task for this z/VM, all the fields will be empty. Click **Action → Refresh from z/VM** to see the existing connections (LANs, OSAs, and VSWITCHes).
4. In the z/VM Networking Properties window, select the **LAN**, **OSA**, or **VSWITCH** tab.
5. Select the connection whose properties you want to edit and click **Edit → Edit Properties** to open the Connection Properties window.

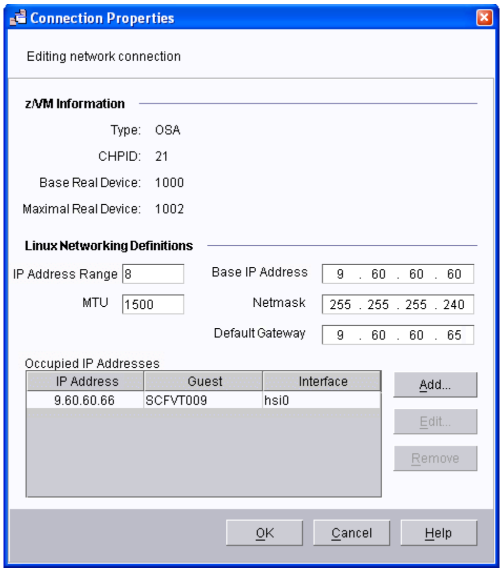

Editing network connection

z/VM Information

Type: OSA
CHPID: 21
Base Real Device: 1000
Maximal Real Device: 1002

Linux Networking Definitions

IP Address Range 8
MTU 1500
Base IP Address 9 . 60 . 60 . 60
Netmask 255 . 255 . 255 . 240
Default Gateway 9 . 60 . 60 . 65

Occupied IP Addresses

IP Address	Guest	Interface
9.60.60.66	SCFVT009	hsi0

a. In the **IP Address Range** field, type the number of IP addresses that can be taken up by the network, starting from the Base IP address, when assigning an IP address to a Linux guest system configured on this connection.

b. In the **MTU** field, type an integer for the maximum transmission unit. In cases of a guest LAN, this must not be larger than the **Maximal frame size** shown in the z/VM information on the upper part of the Connection Properties window.

c. In the **Base IP Address** field, type the lowest IP address that can be used when assigning an IP address to a Linux guest system configured on this guest network.

d. In the **Netmask** field, type the subnet mask that will be used for the Linux routing.

e. (Optional) In the **Default gateway** field, type the default gateway for the connection. This step is optional because you can also type the default gateway in the Server Complex Properties window.

f. Click **OK** to update the connection properties and close the Connection Properties window.

6. In the z/VM Networking Properties window, click **Action**→ **Save** .

Note: You can also double-click a line in the z/VM Networking Properties table to open the Connection Properties window for that connection.

Reserving IP addresses

You can reserve a specific IP address by adding it to the **Occupied IP Addresses** list in the Connection Properties window. This prevents others from using that particular IP address.

Complete the following steps to add an IP address to the list of occupied IP addresses:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Click **Action** → **z/VM Networking Properties**.
3. In the z/VM Networking Properties window, select the **LAN**, **OSA**, or **VSWITCH** tab.
4. Select the connection whose properties you want to edit.
5. Click **Edit** → **Edit Properties** to open the Connection Properties window.
6. Click **Add** to open the Occupied IP Address Properties window.

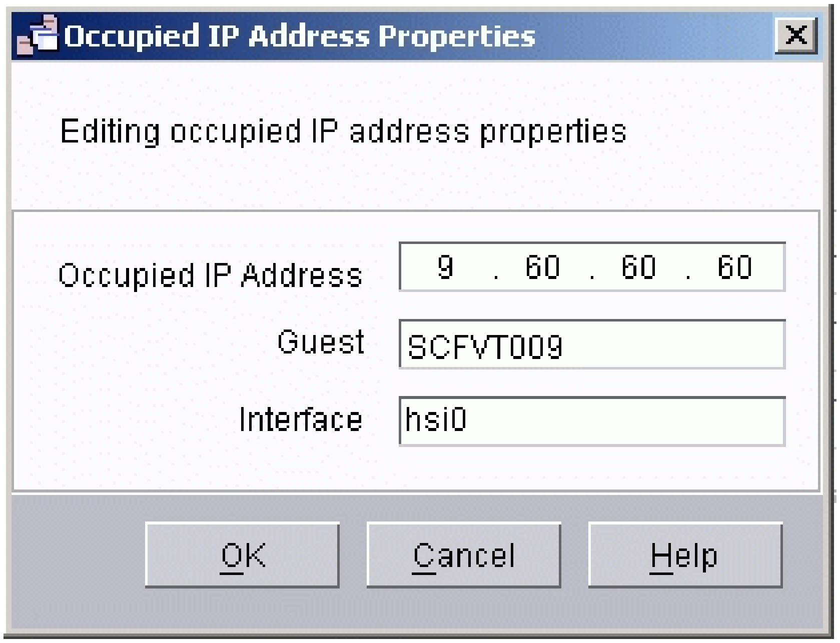

Occupied IP Address Properties

Editing occupied IP address properties

Occupied IP Address: 9 . 60 . 60 . 60

Guest: SCFVT009

Interface: hsi0

7. In the **Occupied IP Address** field, type the occupied IP address. This field is mandatory and its value must be within the IP address range for that connection (as defined in the Connections Properties window).
8. (Optional) In the **Guest** field, type the name of the Linux guest system that owns the occupied IP address.
9. (Optional) In the **Interface** field, type the network interface name of the above Linux guest system.
10. Click **OK** to accept changes and close the Occupied IP Address Properties window.
11. Click **OK** to accept changes and close the Connection Properties window.
12. Click **Action**→**Save** in the z/VM Networking Properties window.

Note: You can also use the **Edit** and **Remove** buttons in the Connection Properties window to update or remove an occupied IP address.

Scanning for occupied IP addresses

You can run a scan for the occupied IP addresses of existing connections in the z/VM. This action updates the list of IP addresses that are already being used by managed systems known to the IBM Director server.

This will prevent you from using an occupied IP address when cloning a new Linux guest system or reconfiguring an existing Linux guest system according to the properties of a server complex network.

Complete the following steps to scan for occupied IP addresses:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Click **Action** → **z/VM Networking Properties** to open the z/VM Networking Properties window.
3. In the z/VM Networking Properties window, click **Action** → **Scan for Occupied IPs**.

This scans the Linux guest systems for used IP addresses and updates the occupied IP list for each connection with the specified properties.

Saving the current networking properties

If you change the networking properties of connections in the z/VM, you must save these changes. This is not done automatically when you click **OK** in the Connection Properties window.

Complete the following steps to save the current networking properties:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Click **Action** → **z/VM Networking Properties** to open the z/VM Networking Properties window.
3. In the z/VM Networking Properties window, click **Action** → **Save**.

Resetting the connection properties

It is impossible to clear all the properties of a selected connection by deleting the fields in the Connection Properties window. To clear these properties, you need to reset the connection.

Complete the following steps to clear all properties of the selected connection:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Click **Action** → **z/VM Networking Properties** to open the z/VM Networking Properties window.
3. In the z/VM Networking Properties window, select a connection.
4. Click **Edit** → **Reset Connection**.

Refreshing networking properties from the z/VM

If you made changes to the networking properties in the Connection Properties window, you will need to refresh the list in the z/VM Networking Properties window so the correct details are displayed. You also need to do this the first time you open the z/VM Networking Properties window.

Complete the following steps to refresh the picture of the networking properties from the z/VM:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Click **Action** → **z/VM Networking Properties** to open the z/VM Networking Properties window.
3. In the z/VM Networking Properties window, click **Action** → **Refresh from z/VM**.

This updates the details of existing connections. It also adds those that are newly discovered and removes those that no longer exist.

Setting and applying server complex properties

In the Server Complexes task, you can set properties for server complexes and apply them to Linux guest systems.

Applying server complex properties to Linux guest systems

There are three ways to apply server complex properties to Linux guest systems.

You can:

- Move a Linux guest system directly into a server complex. This process automatically applies the properties of that server complex to the Linux guest system.
- Change and save the properties of a server complex that already contains Linux guest systems. This applies the new saved properties to the Linux guest systems contained in that server complex.
- Use the **Apply** button on a Server Complex Properties domain page to directly apply the properties of a single domain to the Linux guest systems contained in that server complex (or tier).

Moving Linux guest systems into/out of server complexes:

When you move a Linux guest system into a server complex tier, the properties that you have defined for that server complex tier are applied to the Linux guest system.

The simplest way to apply configuration properties to a Linux guest system is to move a Linux guest system from one server complex tier to another (in the same or a different server complex) or from the Free Linux Guests pane. Once you do this, the Linux guest system is reconfigured according to the properties of the target server complex. First the ‘exiting’ configuration related to the source server complex (if one exists) is applied and then the ‘entering’ configuration related to the target server complex is applied.

Complete the following steps to move a Linux guest system into a server complex tier:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select the Linux guest system that you want to move.
3. Click **Action** → **Move to**.
4. In the **Choose Destination** window, choose from the list of existing server complexes in the **Target server complex** field.
5. If the target is a particular tier, choose from the list of tiers in that server complex in the **Target tier** field.
6. Click **OK**.
7. The Progress Indication window shows the progress of the action. If there is an inconsistency, you are given a choice whether to proceed or stop. Stopping saves the new properties but does not apply the configuration. After the operation is completed, click **Close** to close the Progress Indication window.

Note: You can also perform this action by dragging the Linux guest system icon onto the target tier, or by right-clicking the Linux guest system and clicking **Move to** in the pop-up menu.

Note: To detach a Linux guest system from a server complex, simply do the opposite: either

- Click **Action** → **Detach the guest from the server complex**, or
- Right-click the Linux guest system and click **Move to** in the pop-up menu, or
- Drag the Linux guest system into the Free Linux Guests pane

Changing server complex properties in multiple domains:

You can change the server complex properties in one or more domains. This updates the properties of the Linux guest systems contained in that server complex.

Complete the following steps to change server complex properties in multiple domains:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select the server complex where you want to modify the properties.
3. Click **Action** → **Edit Properties** to open the Server Complex Properties window.
4. Make your changes to the properties of the server complex. (For detailed information about changing properties, see Setting server complex properties.)
5. Click **OK** to apply the configuration changes.
6. If the changes apply to the minidisk or network domains, the Guest Selection window opens. Choose the Linux guest systems on which to apply the reconfiguration and click **OK**.
7. Before the reconfiguration is applied, the new configuration is checked against the old properties. If there is an inconsistency, the Progress Indication window offers a choice to stop or to proceed. Click **Proceed** to override the current configuration.

Directly applying server complex properties in a single domain:

You can directly apply the server complex properties in a single configuration domain to the Linux guest systems contained in that server complex (or tier). This applies the properties immediately, without first checking for consistency.

You can apply server complex properties at any time, with or without changing them. You may want to do this, if, for example, you were unable to apply properties that you modified previously, because a Linux guest system was offline at the time.

Complete the following steps to directly apply server complex properties in a single domain:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select the server complex where you want to apply properties.
3. Click **Action** → **Edit Properties** to open the Server Complex Properties window.
4. In the Server Complex Properties window, choose the domain in which you want to apply properties, and choose whether to apply them to the whole server complex or to a specific tier.
5. (Optional) Make your changes to the server complex properties in the domain. (For detailed information about changing properties, see Setting server complex properties.)
6. Click **Apply**.
7. If you are in the VMRM page, the reconfiguration occurs immediately. In all other domains, the Guest Selection window opens and you can choose the Linux guest systems on which to apply the reconfiguration.
8. Click **OK**.

Setting server complex properties

You can set server complex properties in four domains: VMRM, scripts, network, and minidisks.

Setting VMRM goals for a server complex:

The VMRM (Virtual Machine Resource Manager) performance goal specifies the percentage of CPU and IO resources for a group of virtual machines in a z/VM. By setting the VMRM properties, you set a VMRM goal for Linux guest systems in a server complex and, thus, configure the VMRM to maintain that performance goal.

You can specify these VMRM goals for each tier separately or specify a common goal for all the Linux guest systems in the server complex.

A VMRM goal is considered valid if and only if both:

- The **Importance** field is specified
- Either the **CPU** or **DASD** field is specified.

Complete the following steps to edit the VMRM configuration domain:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select the server complex for which you want to set the VMRM goals.
3. Click **Action** → **Edit Properties** to open the Server Complex Properties window and click the **VMRM** tab to display the VMRM page.

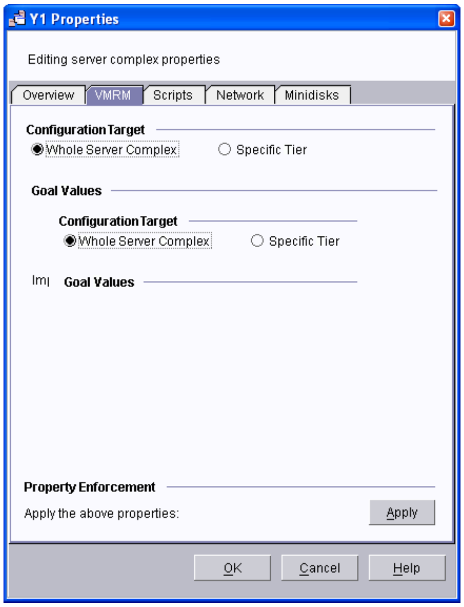

Configuration Target
● Whole Server Complex ○ Specific Tier

Configuration Target
● Whole Server Complex ○ Specific Tier

4. Choose whether the configuration target is the entire server complex or a specific tier.
5. If you chose the **Specific tier** radio button, choose the desired tier. The properties of the selected tier are displayed.
6. In the **CPU** field, type the VMRM CPU velocity goal. This is a value from 1 to 100.
7. In the **DASD** field, type the VMRM DASD (direct access storage device) velocity goal (IO priority). This is a value from 1 to 100.

8. In the **Importance** field, type the VMRM importance of the goal. This is a value from 1 to 10. The Linux guest system with the higher relative value will take precedence in the case of a conflict over resources. Goals with a higher importance value will have higher priority when being maintained by VMRM.
9. Click **OK** to apply the modified settings to the Linux guest systems currently in the server complex. If you modified settings in any of the other domains (except for the Scripts domain), these are also applied to the Linux guest systems in the server complex. Before applying the modified properties, the Server Complexes task runs a consistency check against the old properties.
10. The Progress Indication window shows the progress of the action. If there is an inconsistency, you are given a choice whether to proceed or stop. Stopping saves the new properties but does not apply the configuration. After the operation is completed, click **Close** to close the Progress Indication window.

Note: At any time, with or without changing the property values, you can open this page and click the **Apply** button to directly apply the properties of this domain to Linux guest systems in the server complex. For more information, see Directly applying server complex properties in a single domain.

Note: For further information on VMRM, see **z/VM 5.1 Performance** (SC24-6109-00), chapter 17, ″VMRM SVM Tuning Parameters″.

Setting entering/exiting scripts for a server complex:

By setting the script properties, you can specify an entering script that will be run on a Linux guest system that joins the server complex, and/or an exiting script that will be run on a Linux guest system that leaves the server complex.

Scripts are code added by users to configure a specific environment or application for Linux guest systems placed in a server complex tier. A script can reside on the Linux guest system or on the Director server (from where it is sent to the Linux guest system to be run).

Complete the following steps to edit the scripts configuration domain:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select the server complex where you want to apply the scripts.
3. Click **Action** → **Edit Properties** to open the Server Complex Properties window and click the **Scripts** tab to display the Scripts page.

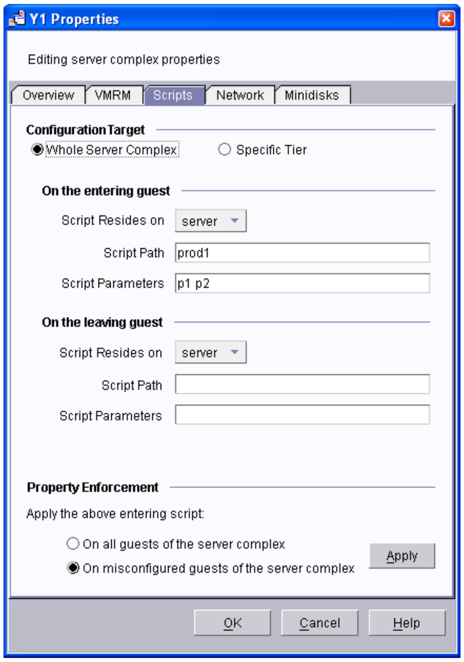

Editing server complex properties

Configuration Target

● Whole Server Complex ○ Specific Tier

On the entering guest

Script Resides on: server

Script Path: prod1

Script Parameters: p1 p2

On the leaving guest

Script Resides on: server

Script Path:

Script Parameters:

Property Enforcement

Apply the above entering script:

○ On all guests of the server complex

● On misconfigured guests of the server complex

4. Choose whether the configuration target is the whole server complex or a specific tier.
5. If you chose the **Specific tier** radio button, choose the target tier. The properties of the selected tier are displayed.
6. For **On the entering guest**, choose either **server** or **guest** in the **Script resides on** field.
7. In the **Script path** field, type the location of the script. If the path starts with / or \, it is considered a full path; otherwise, it is assumed to be the script path relative to the IBM Director data directory.
8. In the **Script parameters** field, type a list of parameters for the script, separated by blank spaces. Words separated by blanks surrounded by double-quotation marks (″) are considered a single parameter.
9. For **On the leaving guest**, choose either **server** or **guest** in the **Script resides on** field.
10. In the **Script path field**, type the location of the script.
11. Click **OK**. The scripts you configured will be applied to Linux guest systems entering or leaving the server complex. If you modified settings in any of the other domains (except for the Scripts domain), these are also applied to the Linux guest systems in the server complex.
12. In the Guest Selection window, choose the Linux guest systems on which to apply the reconfiguration, and click **Yes**.
13. The Progress Indication window shows the progress of the action. After the operation is completed, click **Close** to close the Progress Indication window.

Note: At any time, with or without changing the property values, you can open this page and click the **Apply** button to directly apply the properties of this domain to Linux guest systems in the server complex. For more information, see Directly applying server complex properties in a single domain.

Setting network interfaces for a server complex:

By setting the network properties, you can define up to four network interfaces to be configured on a Linux guest system contained in the server complex.

You can set each interface for configuration on an existing z/VM guest LAN or VSWITCH, or for direct attachment on an OSA card.

Note: Before you can set the networking properties for server complexes, you must set up the connection properties that connections will use for routing the Linux guest systems. (See "Setting connection properties for Linux guest systems" on page 433)

Complete the following steps to edit the network configuration domain:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select the server complex for which you want to define the network interfaces.
3. Click **Action** → **Edit Properties** to open the Server Complex Properties window and click the **Network** tab to display the Network page.

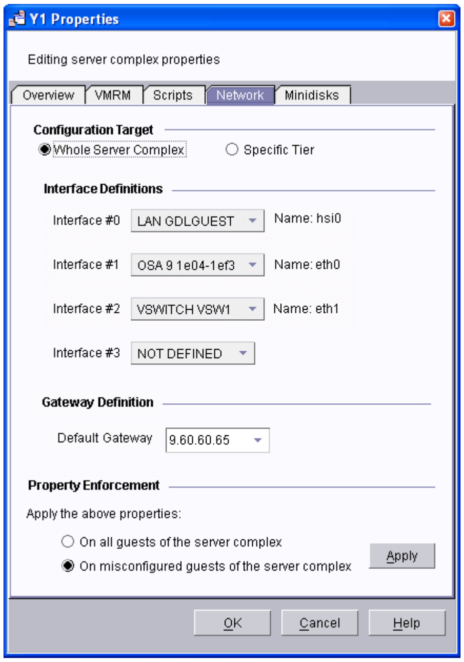

Editing server complex properties

Configuration Target

● Whole Server Complex ○ Specific Tier

Interface Definitions

Interface #0 LAN GDLGUEST Name: hsi0

Interface #1 OSA 9 1e04-1ef3 Name: eth0

Interface #2 VSWITCH VSW1 Name: eth1

Interface #3 NOT DEFINED

Gateway Definition

Default Gateway 9.60.60.65

Property Enforcement

Apply the above properties:

○ On all guests of the server complex

● On misconfigured guests of the server complex

4. Choose whether the configuration target is the whole server complex or a specific tier.
5. If you chose the **Specific tier** radio button, choose the target tier.
6. In the **Interface # 0** field, choose from the list of connections for which routing properties were entered. After you define an interface, the expected Linux interface name (e.g., eth0, hsi0) appears to the right of the selection list.
   - A guest **LAN** element is shown as **LAN** *lan-name*
   - A **VSWITCH** element is shown as **VSWITCH** *vswitch-name*
   - An **OSA** element is shown as **OSA** *chpid min-rdev max-rdev* where *chpid* is the channel path ID, and *min-rdev* and *max-rdev* are the lowest and highest real device numbers of the OSA.
7. (Optional) Repeat the process for the other network interfaces.
8. In the **Default gateway** field, choose the default gateway for configuring on the Linux guest system. The list includes the default gateways of the selected

connections in the **Interface#** fields. You can also type in a different default gateway, which is not included in the list.

9. Click **OK** to apply the modified settings to the Linux guest systems currently in the server complex. If you modified settings in any of the other domains (except for the Scripts domain), these are also applied to the Linux guest systems in the server complex.
10. In the Guest Selection window, choose the Linux guest systems on which to apply the reconfiguration, and click **Yes**.
11. Before applying the modified properties, the Server Complexes task runs a consistency check against the old properties. The Progress Indication window shows the progress of the action. If there is an inconsistency, you are given a choice whether to proceed or stop. Stopping saves the new properties but does not apply the configuration. After the operation is completed, click **Close** to close the Progress Indication window.

Note: At any time, with or without changing the property values, you can open this page and click the **Apply** button to directly apply the properties of this domain to Linux guest systems in the server complex. For more information, see Directly applying server complex properties in a single domain.

Defining minidisk attachments in a server complex:

By setting the minidisks properties, you can define one or more minidisk attachments for the Linux guest systems in the server complex.

A minidisk is identified by its owner (a z/VM user ID) and the owner's virtual number. You provide the desired mount point for each minidisk and the configuration occurs both on the z/VM and at the Linux level.

Complete the following steps to edit the minidisks configuration domain:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select the server complex where you want to specify minidisk attachments.
3. Click **Action** → **Edit Properties** to open the Server Complex Properties window and click the **Minidisks** tab to display the Minidisks page.

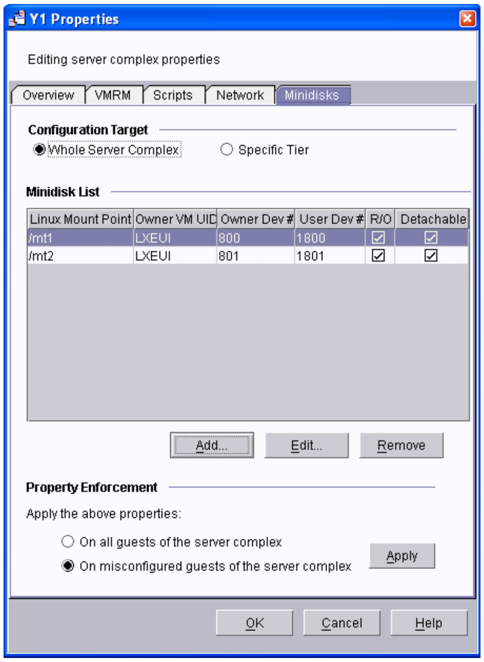

Editing server complex properties

Configuration Target

● Whole Server Complex ○ Specific Tier

Minidisk List

Linux Mount Point	Owner VM UID	Owner Dev #	User Dev #	R/O	Detachable
/mt1	LXEUI	800	1800	☑	☑
/mt2	LXEUI	801	1801	☑	☑

Property Enforcement

Apply the above properties:

○ On all guests of the server complex

● On misconfigured guests of the server complex

4. Choose whether the configuration target is the whole server complex or a specific tier.
5. If you chose the **Specific tier** radio button, choose the target tier.
6. Click **Add** to add a minidisk to the **Minidisk** list. In the Minidisk Properties window, you can enter the properties of the new minidisk.

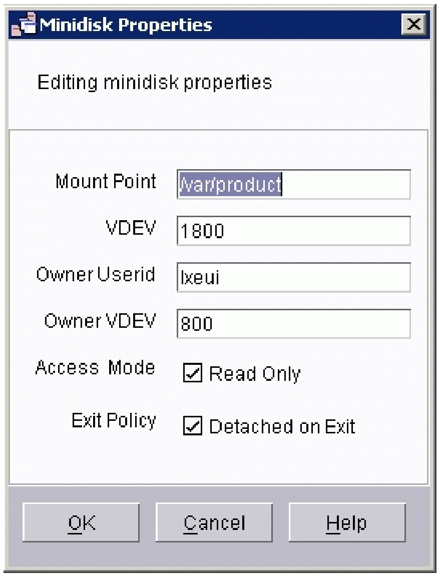

Editing minidisk properties

Mount Point	/var/product
VDEV	1800
Owner Userid	lxeui
Owner VDEV	800
Access Mode	☑ Read Only
Exit Policy	☑ Detached on Exit

a. In the **Mount point** field, type the Linux mount point for the minidisk.
b. In the **VDEV** field, type the virtual device number at the Linux guest system.
c. In the **Owner VDEV** field, type the virtual device number that the disk owner uses for the disk.
d. In the **Owner userid** field, type the VM user ID of the minidisk owner.

e. In the **Access mode** field: If you select the **Read Only** check box, a z/VM **RR** mode and a Linux **ro** mount are used. If you clear this check box, a z/VM **MR** mode and a Linux **rw** mount are used.

f. In the **Exit Policy** field: If you select the **Detached on exit** check box, the minidisk is detached when you remove the Linux guest system or the minidisk from the server complex.

g. Click **OK** to save the minidisk with these properties. The minidisk is added to the **Minidisk** list.

7. (Optional) You can select a line in the **Minidisk** list and click **Edit** (or alternatively, double-click the line) to display or edit the properties of that minidisk.
8. (Optional) Select a line in the **Minidisk** list and click **Remove** to remove a minidisk from the list.
9. Click **OK** to apply the modified settings to the Linux guest systems currently in the server complex. If you modified settings in any of the other domains (except for the Scripts domain), these are also applied to the Linux guest systems in the server complex.
10. In the Guest Selection window, choose the Linux guest systems on which to apply the reconfiguration, and click **Yes**.
11. Before applying the modified properties, the Server Complexes task runs a consistency check against the old properties. The Progress Indication window shows the progress of the action. If there is an inconsistency, you are given a choice whether to proceed or stop. Stopping saves the new properties but does not apply the configuration. After the operation is completed, click **Close** to close the Progress Indication window.

Note: At any time, with or without changing the property values, you can open this page and click the **Apply** button to directly apply the properties of this domain to Linux guest systems in the server complex. For more information, see Directly applying server complex properties in a single domain.

Cloning Linux guest systems

Cloning is the act of creating a new Linux guest system based on an existing z/VM virtual server template and operating system template.

Cloning a new Linux guest system

Using direct cloning, you can create new Linux guest systems and configure them according to server complex properties.

Complete the following steps to clone a Linux guest system:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Make sure the networking properties have been set up in the target server complex tier (either specifically for the tier or for the whole server complex). In particular, make sure each connection defined in a network interface in the network domain has available IP addresses. (For more information, see Setting up connection properties.)
3. If the templates you want to use are not listed in the Provisioning Resources pane, click the **Refresh** button to update the display of the provisioning resources. You can manage these resources in the z/VM Virtual Server Deployment task.

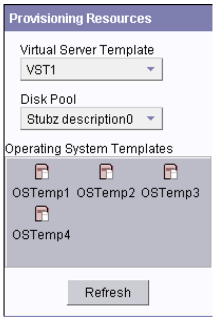

Provisioning Resources

Virtual Server Template: VST1

Disk Pool: Stubz description0

Operating System Templates

4. In the **Virtual Server Template** field, select a virtual server template.
5. In the **Disk Pool** field, choose a disk pool to get the required DASD space for the newly cloned Linux guest system(s).
6. In the **Operating System Templates** field, select an operating system template, right-click and choose **Clone into**.
7. In the Choose Destination window, select the target server complex and target tier, and click **OK**.

   Note: As an alternative method, you can drag an operating system template onto a server complex tier.
8. Define the number of clones you want to create in the target server complex tier and click **OK**. The cloning operation begins, and its progress is displayed in the Progress Indication window. During the cloning operation, the following sub-operations occur automatically:
   a. Creating a virtual server based on the virtual server template.
   b. Applying the operating system template to the created virtual server and configuring the network according to the target server complex property.
   c. Activating the virtual server.
9. When the cloning operation is completed, close the Progress Indication window

After the cloning process is completed, a virtual server icon is displayed inside the server complex that is the target of the cloning operation. The z/VM virtual server waits to be discovered by the IBM Director server. IBM Director then replaces the z/VM virtual server object with a new Linux guest system. This is reflected onscreen when the z/VM virtual server icon is replaced by a regular Linux guest system icon. The networking properties of the server complex/tier are applied in the cloning process. The other configuration properties are applied only when the IBM Director agent is discovered on the new Linux guest system.

Note: It is recommended to set the IBM Director agent on the z/VM MAP (Manageability Access Point) to unlocked status in order to enable the automatic cloning and configuration of multiple Linux guest systems. Otherwise, each Linux guest system discovered after the cloning appears locked on the IBM Director console, and you will need to use the Request Access action on the IBM Director console and provide the Linux root password to unlock it.

Destroying a cloned virtual server

You may need to delete a cloned z/VM virtual server from the z/VM, for example, if the cloning operation fails.

Complete the following steps to delete a cloned z/VM virtual server from z/VM:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Click the z/VM virtual server object inside the server complex.
3. Click **Action** → **Destroy the virtual server in the z/VM**.

   Note: As an alternative method, you can right-click the z/VM virtual server object and choose **Destroy the virtual server in the z/VM**.
4. The Progress Indication window shows the progress of the action. After the operation is completed, click **Close** to close the Progress Indication window.

Detaching a cloned virtual server

You can detach a cloned z/VM virtual server from the server complex. For example, you might create a clone and then realize that this server complex has properties that are not compatible with the intended use of the guest. In this case, you may decide to detach the clone and save it for use in a different server complex.

Since it is not a Linux guest system, the z/VM virtual server will not appear in the Free Linux Guests pane, unless the IBM Director agent on it is discovered later.

Complete the following steps to delete a z/VM virtual server from the z/VM:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Click the z/VM virtual server object inside the server complex.
3. Click **Action** → **Detach virtual server from server complex**.

   Note: As an alternative method, you can right-click the z/VM virtual server object and choose **Detach virtual server from server complex**.
4. Click **Yes** in the Warning window, to proceed.
5. The Progress Indication window shows the progress of the action. After the operation is completed, click **Close** to close the Progress Indication window.

Validating consistency in server complexes

There are situations in which the configuration of Linux guest systems is not consistent with the configuration implied by the properties of the server complex containing them. To prevent this, you can use a number of methods to check for, and ensure, consistency.

Validating the consistency of a server complex

You can validate the consistency of a server complex. To do this, there must be at least one Linux guest system in the server complex and at least one configuration domain that can be validated with defined properties.

The domain being validated may be either VMRM, minidisks, or network. The scripts domain cannot be validated.

Complete the following steps to validate the consistency of a server complex:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select a server complex.
3. Click **Action** → **Check Consistency**.

Note: As an alternative method, you can right-click the server complex and choose **Check Consistency**.

4. The actual configuration is validated against the configuration implied by the server complex properties:
   - For the minidisks and network domains, the configuration is validated for each included Linux guest system.
   - For the VMRM domain, the configuration of the VMRM is validated.
   - For the scripts domain, there is no validation.

If an inconsistency is found, it is reported in the Progress Indication window and the Linux guest system icons and consistency statuses are updated. Any VMRM inconsistency results in an inconsistent status indication for all included Linux guest systems.

Note: If the consistency check action itself fails (e.g., if a Linux guest system is offline), the Linux guest system status is not updated and the failure is reported in the Progress Indication window.

Viewing the configuration status of a Linux guest system

By displaying the Configuration Status message, you can view the consistency status that was set by the last **Check Consistency** action or by a configuration error. However, this action does not perform a new check on the Linux guest system.

Complete the following steps to view the Configuration Status message:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select a Linux guest system and click **Action** → **Guest Status**. The Configuration Status window opens and you can read the most recent configuration status of the Linux guest system.

   Note: As an alternative method, you can right-click the Linux guest system and choose **Guest Status**.
3. Click **Close** to close the Configuration Status window.

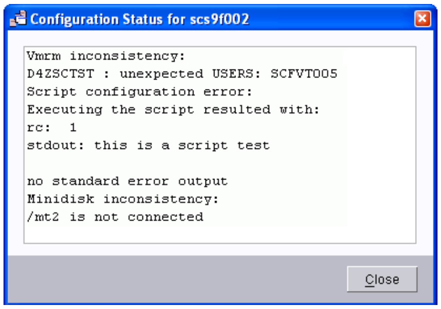

```
Vmrm inconsistency:
D4ZSCTST : unexpected USERS: SCFVT005
Script configuration error:
Executing the script resulted with:
rc:  1
stdout: this is a script test

no standard error output
Minidisk inconsistency:
/mt2 is not connected
```

Checking consistency by viewing Linux guest system icons

The Linux guest system icons in a server complex are represented as icons, with online/offline and consistency status.

You can view the icons to check the status of Linux guest systems.

▣ – Online with consistent status

▭ – Offline with consistent status

⚠ – Online with inconsistent status

⚠ – Offline with inconsistent status

Validating consistency automatically

When you modify server complex properties, the Server Complexes task runs an automatic consistency check. This check alerts you to any changes that may have been applied manually, or by another tool, or by another administrator.

This automatic check occurs after you make changes to the properties of a server complex and click **OK**, or after you move a Linux guest system to, or from, a server complex. Before the reconfiguration is applied, the new configuration is checked against the old properties. If there is an inconsistency, the Progress Indication window displays a report of the differences between the old property values and the actual state, and offers a choice to stop or to proceed.

Note: This consistency check does not occur when you click the **Apply** button in the Server Complex Properties window.

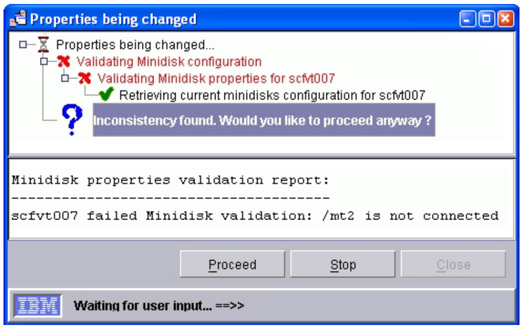

Properties being changed...
Validating Minidisk configuration
Validating Minidisk properties for scfvt007
Retrieving current minidisks configuration for scfvt007
Inconsistency found. Would you like to proceed anyway ?

```
Minidisk properties validation report:
--------------------------------------
scfvt007 failed Minidisk validation: /mt2 is not connected
```

Waiting for user input... ==>>

The screenshot shows an inconsistency being reported in the Progress Indication window.

Clearing the consistency status

Sometimes you may want to clear the consistency status of the Linux guest system icons within a server complex. This might happen if the Linux guest systems have an inconsistency status resulting from a script that could not be validated, or because a machine was offline.

Complete the following steps to clear the status:

1. If it is not started already, start the Server Complexes task for the z/VM system you want to work with.
2. Select a server complex and click **Action** → **Clear all Statuses**.

   Note: As an alternative method, you can right-click the server complex and choose **Clear all Statuses**.

Chapter 4. Troubleshooting

Use this section to troubleshoot and resolve problems with IBM Director.

Installation troubleshooting

Use this section to troubleshoot and resolve problems with IBM Director installation.

For additional troubleshooting information, see the *IBM Director Release Notes*.

Could not detect *rpm* supported Linux distribution

This problem only affects systems running Linux.

Problem

When installing IBM Director Core Services, IBM Director Agent, or IBM Director Server, the following error message might be generated:

```
Could not detect rpm supported Linux distribution
```

where *rpm* is one of the following RPMs:

- DirectorCimCore
- xSeriesCoreServices-level1
- pSeriesCoreServices-level1

Investigation

This message occurs if you install IBM Director components on a Linux distribution that is not supported by the DirectorCimCore, xSeriesCoreServices-level1, and pSeriesCoreServices-level1 RPMs.

To correct the problem, install a supported Linux distribution. See *IBM Director Installation and Configuration Guide*.

depmod: *** Unresolved symbols in /lib/modules/2.4.7-10smp/kernel/drivers/char/

This problem only affects systems running Red Hat Linux.

Problem

If you install the IBM SMbus or LM78 device driver, a warning message similar to the following might be recorded in the Bootmsg.log file:

```
depmod: *** Unresolved symbols in /lib/modules/2.4.7-10smp/kernel/drivers/
char/driver.o
```

where *driver* is either ibmsmb or ibmlm78.

Investigation

This warning message is generated if the Red Hat Linux kernel was compiled with versioned symbols enabled. Ignore this message; the device driver loaded and will operate properly.

Error 1722 is displayed

This problem affects IBM Director Server and IBM Director Console. It occurs only on systems running Windows.

Problem

When you install IBM Director, the following message is displayed:

```
Error 1722. There is a problem with this Windows Installer package. A
program run as part of the setup did not finish as expected. Contact your
support personnel or package vendor.
```

Investigation

The monitor for a system running IBM Director Server or IBM Director Console must support at least 256 colors.

To correct this problem, increase the display color palette to more than 256 colors, uninstall the partial installation, and reinstall IBM Director Server.

Installing IBM Director 4.2x over IBM Director 5.10

This problem affects IBM Virtualization Engine Systems Edition for iSeries installations.

Problem

After you have installed IBM Director 5.10 (5722-DR1 or 5722-DA1) on a system as part of IBM Virtualization Engine Systems Edition for iSeries, you must completely uninstall IBM Director 5.10 before installing the previous version, IBM Director Multiplatform 4.2x (5733-VE1 option 30 or 39).

If attempted, IBM Director 5.10 is disabled and IBM Director Multiplatform 4.2x does not work (twgstart will fail).

Investigation

To correct this problem, complete the following steps:

1. Uninstall IBM Director Multiplatform 4.2x (5733-VE1 option 30 or 39).
2. Complete the manual cleanup described in the Virtualization Engine Information Center > eServer Software Information Center at: http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/info/veicinfo/eicarmanualuninstall.htm#eicarmanualuninstall.
3. Install only one of the following releases:
   - IBM Director 5.10 (5722-DR1 or 5722-DA1)
   - IBM Director Multiplatform 4.2x (5733-VE1 option 30 or 39)

Installation package cannot be installed by Windows Installer service

This problem affects systems with a version of Windows Installer (MSI) that is earlier than version 3.0. This problem occurs only on systems running Windows 2000, Windows XP with no service packs, or Windows XP with Service Pack 1.

Problem

During a Windows Installer administrative installation, the following message is displayed:

```
This installation package cannot be installed by
Windows Installer service. You must install a Windows
service pack that contains a newer version of the
Windows Installer service.
```

Investigation

To correct this problem and perform Windows Installer administrative installation, complete the applicable procedure.

- If you are installing a Web-downloadable extension, type the following command:

  ```
  filename.exe -a admin
  ```

  where *filename* is the name of the Web-downloadable extension installation file.
- If you are installing IBM Director Server, IBM Director Console, IBM Director Agent, or IBM Director Core Services from the *IBM Director* CD, complete the following steps:
  1. On the *IBM Director* CD, change to the directory for the IBM Director component that you want to install.
  2. Type one of the following commands:

     ```
     filename.exe -a admin
     ```

     where *filename* is the name of the IBM Director component installation file.

Note: If the system that is running the administrative installation has a version of MSI that is earlier than version 3.0, these commands update MSI on that system. After MSI is updated and the administrative installation is completed, a message is displayed that you must reboot the system. Be sure to do so.

System Availability reports an unplanned outage when a system is restarted

This problem only affects systems running Red Hat Enterprise Linux AS, versions 2.1 and 3.0.

Problem

System Availability reports an unplanned outage when a Level-2 managed system is restarted.

Investigation

To avoid this problem, after installing System Availability, be sure to stop and restart IBM Director Agent before you restart the managed system.

Windows blue screen IRQL_NOT_LESS_OR_EQUAL

This problem affects IBM Director Agent. It occurs only on systems running Windows Server 2003.

Problem

During installation of IBM Director Agent, Windows might display the following blue screen trap:

```
IRQL_NOT_LESS_OR_EQUAL
```

Investigation

This problem is solved by a Microsoft update. See Microsoft Knowledge Base Article 825236 for more information.

IBM Director Server troubleshooting

Use this section to troubleshoot and resolve problems with IBM Director Server.

Some IBM Director Server problems might appear to be problems with IBM Director Console. Review all troubleshooting topics for possible solutions. For additional troubleshooting information, see the *IBM Director Release Notes*.

BladeCenter discovery does not function correctly

This problem affects BladeCenter products only.

Problem

A BladeCenter discovery does not function correctly when multiple network interface cards (NICs) are enabled.

Investigation

Determine the NICs that are connected to the BladeCenter unit network. Disable all NICs except one, which must be able to communicate with the BladeCenter management module. Perform the discovery. When the discovery is completed, re-enable the NICs that you disabled.

Note: You must do this each time you want to discover the BladeCenter unit and its components.

Communications timeout between IBM Director Server and IBM Director Console

This problem affects IBM Director Server and IBM Director Console.

Problem

A timeout occurs during communications between IBM Director Server and IBM Director Console.

Investigation

Working with large event-action plans can cause network communication errors to occur. IBM Director Server takes a long time to process large requests from IBM Director Console. During this processing period, IBM Director Console waits for a response from IBM Director Server. When no response is received after 15 seconds (or the timeout value that is configured for the IBM Director environment), a timeout error is generated. This error might occur several times for intensive operations, such as importing or exporting large event-action plans.

Despite the communication error, the event action plan works correctly.

Error in the event log: The open procedure for service PerfDisk

This problem affects IBM Director Server. The problem occurs only on servers running Windows 2000 Server.

Problem

After IBM Director Server is installed, the following error is displayed in the event log when the server is restarted:

```
The open procedure for service PerfDisk in the DLL C:\WINNT\System32\perfdisk.dll
has taken longer than the established wait time to be completed.
```

Investigation

Use the **regedit** command to modify the following key entry and change the decimal value to 30000:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services
\PerfDisk\Performance key "Open Timeout"
```

This gives the system enough time to complete the startup task before starting the PERF counters.

Errors occur during the Oracle Server database configuration

Problem

When an Oracle Server database is used, errors occur during the database configuration process.

Investigation

Configure and start the Oracle TCP/IP listener before starting the database configuration task. If a failure occurs, check the configuration of the TCP/IP listener.

Event actions fail

This problem affects IBM Director Server.

Problem

After a NIC on the management server is reconfigured, certain event actions fail.

Investigation

IBM Director Server has lost contact with the managed systems that were discovered before the configuration change. From IBM Director Console, click **Tasks** → **Discover Systems** → **System Discovery** to rediscover the managed systems.

daemon.stderr file reports an exception

This problem affects IBM Director Server and IBM Director Console. It occurs only on systems running Linux.

Problem

Either of the following components enters an error state shortly after starting:

- IBM Director Server
- IBM Director Agent

The daemon.stderr file reports the following error:

```
Exception in thread "main"
```

Investigation

To correct this problem, complete the applicable steps:

- **For IBM Director Server**: Make sure that “localhost” is an alias for the loopback address 127.0.0.1 in the /etc/hosts file on the management server. Restart IBM Director Server.
- **For IBM Director Agent**: Make sure that “localhost” is an alias for the loopback address 127.0.0.1 in the /etc/hosts file on the managed system. Restart IBM Director Agent.

IBM Director encryption fails

This problem affects IBM Director Server. It occurs only on servers running i5/OS.

Problem

IBM Director encryption fails.

Investigation

To correct this problem, complete the applicable action:

- **For IBM Director 4.20, 4.21, and 4.22:**
  - Install the Java Developer Kit 1.3, Option 5 (5722-JV1) and update it with the latest PTF packages.
  - Install the latest Licensed Internal Code (5722-999) PTF packages.

  Note: Verify that the Java Cryptography Extension (JCE) has the IBMJCE enabled in /qibm/proddata/java400/jdk13/lib/security/java.security as documented in the iSeries Information Center at: http://publib.boulder.ibm.com/infocenter/iseries/v5r3/topic/rzaha/rzahajce.htm.
- **For IBM Director 5.10:** Install the Java Developer Kit 1.4, Option 6 (5722-JV1) and the latest i5/OS Operating System (5722-SS1) program temporary fix (PTF) packages.

IBM Director Server might not discover systems and display them in IBM Director Console

This problem affects only managed systems running Linux.

Problem

When no default router is configured or a nonroutable private network is used, IBM Director might not discover systems on these networks and add them to the IBM Director Console Group Contents pane.

Investigation

To correct this problem, complete one of the following procedures:

- Seed the network in the System Discovery (IP) pane. Click **Options** → **Discovery Preferences**. Then, click **System Discovery (IP)**.
- Set a default router by issuing the following command:

  ```
  route add default gw IP_address
  ```

  where *IP_address* is your IP address. For more information, see the man page for the **route** command. Setting a default router enables the discovery of systems that are accessible using the specified router.

IBM Director Server does not discover SNMP devices

Problem

IBM Director Server does not discover SNMP devices.

Investigation

Make sure that the following conditions are met:

- The management server is running the SNMP service. If it is not, another system on the same subnet must be running an SNMP agent. In that case, remove the management server as the seed device and add the system running the SNMP agent.
- The seed devices or other devices to be discovered are running SNMP agents.
- The community names that are specified in the Discovery Preferences window allow IBM Director to read both the following tables:
  - mib-2.system table of the devices to be discovered
  - mib-2.ip.ipNetToMediaTable on the seed devices
- Correct network masks have been configured for all managed systems that must be discovered.
- Correct addresses have been entered for the seed devices. The most effective seed devices are routers and domain name servers. To configure these devices, from IBM Director Console, click **Options** → **Discovery Preferences**. SNMP discovery does not discover all SNMP devices. If a device has not communicated with other managed systems, the device might not be discovered.
- The devices are part of the local subnet and you are using Broadcast Relay or Unicast.

  Note: Broadcast and Multicast can be blocked a router.

IBM Director Server fails to start

This problem affects IBM Director Server. It occurs only on servers running i5/OS.

Problem

IBM Director Server fails to start when the Japanese coded character set identifier (CCSID) 5026 is used.

Investigation

Make sure that the job CCSID and locale match and that they are supported by Qshell. Consider using CCSID 5035 and locale JA_5035.

For more information, go to the iSeries Information Center at www.ibm.com/servers/ and search on National Language Support.

Managed system icon displays a question mark

This problem affects IBM Director Server.

Problem

After you use the Encryption Administration window to change encryption settings, certain managed systems cannot be managed and their icons display a question mark (?).

Investigation

This problem might be caused by one of the following circumstances:

- When you request a new key or a cipher algorithm, IBM Director must perform a presence check. This presence check might not be completed immediately. During the delay, IBM Director Server cannot manage the system.
- If you disable encryption on the management server, encrypted managed systems can no longer be managed. However, these systems might appear to be manageable for a certain period before they are displayed as locked.

To correct this problem, force an immediate exchange of keys on a particular managed object by requesting a presence check.

Microsoft Jet database is full

This problem occurs only on systems running Windows.

Problem

The Microsoft Jet database is full.

Investigation

Because the Jet database does not recover space when deletes and updates are done. To correct this problem, compact the database by recycling IBM Director Server. If recycling fails to correct the problem, migrate to a larger database such as IBM DB2®, Oracle Server, or Microsoft SQL Server.

RXE-100 Remote Expansion Enclosure is not discovered

Problem

After you click **Discover All Systems**, an RXE-100 Remote Expansion Enclosure is not discovered.

Investigation

To solve this problem, perform one of the following procedures:

- From IBM Director Console, click **Tasks** → **Discover Systems** → **Physical Platforms**; then, click **Discover All**.
- Right-click any blank space in the Group Contents pane and click **New** → **Physical Platforms**. The Add Physical Platforms window opens. Type the name and IP address of the Remote Supervisor Adapter that is attached to the RXE-100 Remote Expansion Enclosure; then, click **OK**.

sed errors occur when starting IBM Director Server

This problem affects IBM Director Server. It occurs only on servers running i5/OS.

Problem

When starting IBM Director Server (twgstart), the following sed errors occur:

sed: 001-2272 Error in file ″s/^\(..\).*/\1/″ on line 1: character 1 not defined in the regular expression.

sed: 001-2272 Error in file ″s/^\(..\)\(.\).*/\2/″ on line 1: character 2 not defined in the regular expression.

sed: 001-2272 Error in file ″s/^\(..\)\(..\).*/\2/″ on line 1: character 2 not defined in the regular expression.

Investigation

The job is not using valid values. To correct this problem with the valid values, see the iSeries Information Center at:
http://publib.boulder.ibm.com/infocenter/iseries/v5r3/topic/ rzahz/nls.htm.

Note: The values on this Web page do not include any user-created locales that might not be valid.

TWGServer.err reports a database initialization error

This problem affects IBM Director Server. It only occurs on servers running Linux.

Problem

When the IBM Director database is run locally on the management server and the management server is restarted, IBM Director Server fails to start. The TWGServer.err file reports a database initialization error.

Investigation

The TWGserver service might have started before the database service. Back up the etc/init.d/TWGserver script and save it to a safe location. Then, modify the etc/init.d/TWGserver script to make sure that the database service starts before the IBM Director service:

For Red Hat Linux: Locate the following section in the script:

```
# chkconfig: 35 90 10
# description: Starts and stops the IBM Director service.
```

90 is the start number and 10 is the stop number. Modify this section so that the TWGserver start number is greater than the start number for the database service, and the TWGserver stop number is greater than the stop number for the database service.

For SUSE LINUX: Locate the following section in the script:

```
### BEGIN INIT INFO
# Required-Start: $network
# Required-Stop:  $network
# Default-Start: 3 5
# Default-Stop:   0 1 6
# Description:    Starts and stops the IBM Director service.
### END INIT INFO
```

Add the database service to the Required-Start and Required-Stop lines. For example, for PostgreSQL, change the lines to read as follows:

```
# Required-Start: $network postgresql
# Required-Stop:  $network postgresql
```

Save the modified script. Run the **chkconfig** command twice, once to remove the IBM Director service and then to add it back to the list of start and stop services.

Uncertain if IBM Director Server is running

This problem affects IBM Director Server.

Problem

You are not sure if IBM Director Server is running.

Investigation

To check whether the management server is running, complete one of the following procedures:

- (i5/OS) From a Qshell command prompt, type the following command and press **Enter**:

  ```
  /QIBM/ProdData/Director/bin/twgstat
  ```

  The current status of IBM Director Server is displayed.
- (Linux) From a command prompt, type the following command and press **Enter**:

  ```
  /opt/IBM/director/bin/twgstat
  ```

  The current status of IBM Director Server is displayed.
- (Windows) Determine which of the following icons is displayed in the task bar in the lower-right corner of the screen.

- A green circle indicates that IBM Director Server is running.
- A green triangle icon indicates that IBM Director Server is in the process of starting.
- A red diamond icon indicates that IBM Director Server is not responding.

Do not attempt to start IBM Director Console until a green circle is displayed in the task bar.

User is locked out of IBM Director

This problem affects IBM Director Console and Level-2 managed systems. The problem occurs only on management servers (IBM Director Server) running Linux.

Problem

A user is locked out of IBM Director. This can happen in the following situations:

- After trying six times unsuccessfully to login to IBM Director Console
- After trying six times unsuccessfully to request access to a Level-2 managed system (IBM Director Agent)

Investigation

To correct this problem, reset the locked-out user ID by typing the following command on the management server:

```
pam_tally --user myuser --reset
```

where *myuser* is the locked-out user ID.

If you want to change the number of permitted login failures, edit the `/etc/pam.d/ibmdir` file on the management server.

IBM Director Console troubleshooting

Use this section to troubleshoot and resolve problems with IBM Director Console.

Some IBM Director Console problems might be problems with IBM Director Server or other components. Review all troubleshooting topics for possible solutions. For additional troubleshooting information, see the *IBM Director Release Notes*.

A question mark is displayed with a managed system icon

This problem affects IBM Director Console.

Problem

A question mark is displayed with the managed system icon.

Investigation

Reestablish communication between IBM Director Server and IBM Director Agent on the managed system. Click **Tasks** → **Discover Systems** → **System Discovery** to rediscover the managed system.

After a PPMO is deleted, it reappears in IBM Director Console

This problem affects IBM Director Console.

Problem

After a physical platform managed object is deleted, it reappears in IBM Director Console.

Investigation

Delete the managed system or systems that are associated with the physical platform managed object.

An event-action plan is not displayed

This problem affects viewing event-action plans.

Problem

An event-action plan is not displayed.

Investigation

When you apply an event-action plan to a group, the event-action plan is associated with *all* existing systems in the group. However, this group event-action plan is not displayed as associated with each individual managed system that is part of the group. The event-action plan is displayed as being applied to the group *only*.

To view the event action plans associated with the groups of managed systems, complete the following steps:

1. In IBM Director Console, click **Associations** → **Event Action Plans**.
2. In the Groups pane, click **All Groups**.
3. In the Group Category Contents pane, expand each group that has an event action plan applied to it and view the event action plans that are applied to the group.

BladeCenter network device displayed as offline when it is online

This problem affects BladeCenter network devices only.

Problem

An existing SNMP managed object displays a BladeCenter network device as offline when it is online.

Investigation

Make sure that the SNMP community name in the SNMP agent of the switch module is set to ″public″ and matches the community name in the SNMP managed object in IBM Director Console.

IBM Director expects the community name of the SNMP managed object of the switch module to be set to ″public.″ If the community name of the switch module is not set to ″public,″ the BladeCenter chassis discovery will not create an SNMP managed object for that switch module. Also, the community name in the SNMP agent that is installed on the switch module must match the community name in the SNMP managed object in IBM Director Console for the SNMP managed object to stay online.

If you must use an SNMP community name other than ″public″ with a switch module, complete the following steps:

1. Set the community name in the SNMP agent that is installed on the switch module to ″public.″
2. Run the BladeCenter chassis discovery. IBM Director creates an SNMP managed object for the switch module and displays its status as Online.
3. Set the community name in the SNMP agent that is installed on the switch module to the new value. IBM Director Console displays the status of the switch module as Offline.
4. Change the community name in the SNMP managed object to the new value. The status of the switch module is displayed as Online.

Blade server PPMO is not displayed in IBM Director Console

This problem affects IBM Director Console.

Problem

After a blade server is installed in a BladeCenter chassis, a physical platform managed object (PPMO) associated with the blade server is not displayed in IBM Director Console.

Investigation

Run the Inventory task on the BladeCenter chassis.

Cannot access RDM help from IBM Director Console

This problem affects the Remote Deployment Manager (RDM) task.

Problem

You cannot access the Remote Deployment Manager (RDM) help from IBM Director Console in the following situations:

- An installation that includes IBM Director, version 4.20 or 4.20.2, and RDM 4.11 (including any of the RDM 4.11 Updates 1, 2, or 3)
- An installation that includes IBM Director 4.12 and RDM 4.11 (including any of the RDM 4.11 Updates 1, 2, or 3) is upgraded to use IBM Director, version 4.20 or 4.20.2

Investigation

To access the RDM help files, complete the following steps:

1. Navigate to the *director_install*\classes\doc\ibm directory, where *director_install* is the directory where IBM Director is installed.
2. Move the rdm directory to the *director_install*\classes\doc\en\ibm directory.

Changes to a BladeCenter chassis are not displayed

This problem affects BladeCenter products only.

Problem

After you remove or insert a blade server, IBM Director Console takes several minutes to register changes to the BladeCenter chassis. This can affect the speed with which IBM Director applies chassis detect-and-deploy policies, runs event-action plans, updates the chassis association, and writes to the event log.

Investigation

If, after you insert or remove a blade server and wait several minutes, the BladeCenter chassis and blade servers are not properly represented in IBM Director Console, perform a BladeCenter chassis discovery operation.

Corrupt double-byte character set (DBCS) characters are displayed

This problem only affects systems running Japanese, Korean, Traditional Chinese, or Simplified Chinese installations of SUSE LINUX Enterprise Server 8.

Problem

IBM Director Console displays corrupt Japanese, Korean, Traditional Chinese, or Simplified Chinese characters.

Investigation

Some fonts required for double-byte character set (DBCS) languages might not be registered correctly during the installation of SUSE LINUX Enterprise Server 8. This can cause IBM Director Console to display corrupt characters. To correct this problem, complete the following steps as the root user:

1. Change to the following directory: /usr/X11R6/lib/X11/fonts/truetype
2. Run the **SuSEconfig** command.
3. Restart IBM Director Console The required fonts are now registered.

Discovery does not create an SNMP managed object for a BladeCenter network device

This problem affects BladeCenter network devices only.

Problem

The BladeCenter chassis discovery process does not create an SNMP managed object for a BladeCenter network device.

Investigation

Make sure that the SNMP community name in the SNMP agent of the switch module is set to ″public″ and matches the community name in the SNMP managed object in IBM Director Console.

IBM Director expects the community name of the SNMP managed object of the switch module to be set to ″public.″ If the community name of the switch module is not set to ″public,″ the BladeCenter chassis discovery will not create an SNMP managed object for that switch module. Also, the community name in the SNMP agent that is installed on the switch module must match the community name in the SNMP managed object in IBM Director Console for the SNMP managed object to stay online.

If you must use an SNMP community name other than ″public″ with a switch module, complete the following steps:

1. Set the community name in the SNMP agent that is installed on the switch module to ″public.″
2. Run the BladeCenter chassis discovery. IBM Director creates an SNMP managed object for the switch module and displays its status as Online.
3. Set the community name in the SNMP agent that is installed on the switch module to the new value. IBM Director Console displays the status of the switch module as Offline.
4. Change the community name in the SNMP managed object to the new value. The status of the switch module is displayed as Online.

Duplicate managed systems are displayed

This problem affects IBM Director Console.

Problem

After using imaging to deploy a system, duplicate managed systems are displayed in IBM Director Console. When using imaging, make sure that the instance of IBM Director Agent that is being cloned has not been started.

Investigation

Perform one of the following procedures on the duplicate managed system:

Linux: Complete the following steps:

1. If IBM Director Agent is running, stop it by typing `twgstop`.
2. Type the following command:
   ```
   twgreset -i
   ```
3. Start IBM Director Agent by typing `twgstart`.

Windows: Complete the following steps:

1. If IBM Director Agent is running, stop it by typing `twgstop`.
2. Remove the following registry key:
   ```
   HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\
   ComputerName\ComputerName\TWGMachineID
   ```
3. Delete the twgmach.id file. If you installed IBM Director Agent in the default location, this file is in the c:\Program Files\IBM\Director\data directory.
4. Start IBM Director Agent by typing `twgstart`.

Errors occur when logging into IBM Director Server

This problem affects IBM Director Console and IBM Director Server.

Problem

Errors occur during attempts to log on to the management server using IBM Director Console.

For example, when you try to start IBM Director Console, the following error message is displayed:

```
An IO error occurred while connecting to the IBM Director Server.
```

Investigation

To correct this problem, make sure that the following requirements are met:

- Both the management server and IBM Director Server are running before you start IBM Director Console.
  - (i5/OS) From a Qshell command prompt, type the following command and press **Enter**:

    ```
    /QIBM/ProdData/Director/bin/twgstat
    ```

    The current status of IBM Director Server is displayed.
  - (Linux) From a command prompt, type the following command and press **Enter**:

    ```
    /opt/IBM/director/bin/twgstat -r
    ```

    The current status of IBM Director Server is displayed. That status must remain Active.
  - (Windows) Determine which of the following icons is displayed in the task bar in the lower-right corner of the screen.
    - A green circle indicates that IBM Director Server is running.
    - A green triangle icon indicates that IBM Director Server is in the process of starting.
    - A red diamond icon indicates that IBM Director Server is not responding.

    Do not attempt to start IBM Director Console until a green circle is displayed in the task bar.
-
- IBM Director Console and IBM Director Server are the same version.
- The management server name, user ID, and password are valid. (For systems running Windows, you must qualify the user ID with either the domain or the local computer name of the management server.)
- The password is correct and does not require updating.
- If the Use SSL check box is not selected in IBM Director Console, make sure that IBM Director Server accepts nonsecure TCP connections. Click **Options** → **Server Preferences**. Then, click **Connections** and select the “Allow non-secure TCP console connections” check box.
- If the Use SSL check box is selected in IBM Director Console, make sure that IBM Director Server accepts SSL connections. Click **Options** → **Server Preferences**. Then, click **Connections** and select the “Allow secure SSL console connections” check box.
- (If SSL is used) Both the management console and the management server are using compatible data link connection classes and parameters in the TWGConsole.prop and TWGServer.prop files.
- The ports have not been changed in the properties files for IBM Director Server or IBM Director Console (TWGServer.prop and TWGConsole.prop files). The ports must match.

Note: It is possible that hardware or software dependencies outside of IBM Director can cause this problem, such as incorrectly configured TCP/IP or a firewall that does not permit communication.

IBM Director Server might not discover systems and display them in IBM Director Console

This problem affects only managed systems running Linux.

Problem

When no default router is configured or a nonroutable private network is used, IBM Director might not discover systems on these networks and add them to the IBM Director Console Group Contents pane.

Investigation

To correct this problem, complete one of the following procedures:

- Seed the network in the System Discovery (IP) pane. Click **Options** → **Discovery Preferences**. Then, click **System Discovery (IP)**.
- Set a default router by issuing the following command:

  ```
  route add default gw IP_address
  ```

  where *IP_address* is your IP address. For more information, see the man page for the **route** command. Setting a default router enables the discovery of systems that are accessible using the specified router.

IBM Director Console fails to open the Login window

This problem affects systems running Korean, Japanese, and Simplified and Traditional Chinese installations of Red Hat Enterprise Linux, version 4.0, for AMD64 and EM64T only.

Problem

IBM Director Console fails to open the Login window and produces a JavaCore listing with Input Method exceptions.

Investigation

To correct this problem, complete the following steps:

1. On the KDE desktop, click **System Settings** → **Input Method Switcher**.
2. In the "system-switch-im window," select **Advanced Settings**.
3. In the Input Methods list, select an input method other than iiimf, such as kinput2-canna.
4. Click **OK**.
5. Restart the X Window System.

Level-0 managed objects cannot be accessed

This problem only affects Level-0 managed systems running i5/OS.

Problem

A Level-0 managed system that is running i5/OS cannot be accessed.

Investigation

To correct this problem, make sure that Open SSH, *BASE and Option 1, Licensed Product Offering 5733-SC1 is installed and configured. Also make sure that SSH is running on the affected Level-0 managed system.

Level-1 managed system changes to Level-0 managed system

This problem affects Level-1 managed systems and management servers.

Problem

IBM Director discovers and displays a Level-1 managed system. However, after double-clicking the managed system in IBM Director Console, the managed system is displayed as a Level-0 managed system and the CIM protocol is missing from the list of specified protocols.

Investigation

The system clock on either the managed system or the management server is incorrect. To correct this problem, determine which system clock is incorrect: the managed system or the management server.

Note: If the management server clock is incorrect, that can affect many managed systems. If the managed system clock is incorrect, then only that managed system is affected.

Check the ras.log file on the managed system for any messages that indicate a certificate was not valid or expired.

- If a message states that the certificate was not valid, the system clocks are not synchronized and the managed system clock is *behind* the management server clock. For example, the managed system is set for 2004, but the validity of the management server certificate does not start until 2006.
- If a message states that the certificate has expired, the system clocks are not synchronized and the managed system clock is *ahead* of the management server clock. For example, the managed system is set for 2008, but the validity of the management server certificate ends in 2006.

If the managed system clock is incorrect, complete the following steps:

1. On the system with the incorrect clock, stop the CIMOM by typing:

For Linux	`/etc/ibm/director/diragent/stopcim`
For Windows	`net stop wmicimserver`

2. Correct the affected clock. Make sure the managed system clock is synchronized or a little ahead of the management server clock. Certificates allow up to 24 hours difference.
3. Start the CIMOM by typing:

For Linux	`/etc/ibm/director/diragent/startcim`
For Windows	`net start wmicimserver`

If the management server clock is incorrect, complete the following steps:

1. Correct the affected clock. Make sure the management server clock is synchronized or slightly behind all Level-1 managed systems. Certificates allow up to 24 hours difference.
2. Use the dircli **certmgr** command to generate a new certificate.

Level-1 managed objects are not discovered

This problem affects pSeries servers running SUSE LINUX Enterprise Server 9 for IBM POWER and xSeries servers running SUSE LINUX Enterprise Server 9 for x86.

Problem

IBM Director Server does not discover Level-1 managed objects that are running the Service Location Protocol daemon (SLPD). The Level-1 managed objects are not displayed in IBM Director Console.

Investigation

SLPD on SUSE LINUX Enterprise Server 9 for IBM POWER and SUSE LINUX Enterprise Server 9 for x86 is not SLP compliant and does not accept registrations. To correct this problem, disable SLPD on the Level-1 managed system and then restart the server.

Level-2 managed systems are not displayed

This problem affects IBM Director Agent.

Problem

Level-2 managed systems are not displayed in IBM Director Console.

Investigation

To correct this problem, make sure that the system is turned on, IBM Director Agent is running, and the network connection is reliable.

Increase the network timeout value for both IBM Director Server and IBM Director Agent:

- **Windows:** Run twgipccf.exe.
- **Linux:** Using an ASCII text editor, open the ServiceNodeLocal.properties file (located in the /opt/ibm/director/data directory), and modify the value of `ipc.timeouts`. By default, it is set to 15 seconds.

Stop and restart IBM Director Agent to ensure that the new network timeout takes effect.

Not all managed systems that meet dynamic group criteria are returned

This problem affects dynamic groups.

Problem

When a dynamic group is created using certain criteria (such as the not-equal-to operator as part of the selected criteria), not all of the managed systems that meet those criteria are returned.

Investigation

Make sure that you use the correct criteria when you create the dynamic group. Each criterion searches only the rows in the inventory database with which it is associated.

For example, when you select the following criterion:

```
SCSI Device / Device type = TAPE
```

IBM Director searches the inventory database for managed systems that have entries in the SCSI_DEVICE table. Then, IBM Director returns only the managed systems that have a value of TAPE in the DEVICE_TYPE column.

When you select the following criterion:

```
SCSI Device / Device type ^= TAPE
```

IBM Director searches the inventory database for managed systems that have entries in the SCSI_DEVICE table. Then, IBM Director returns only the managed systems that do not have a value of TAPE in the DEVICE_TYPE column.

Selecting the second criterion does not return all managed systems that do not have SCSI tape drives. It returns all managed systems that contain non-tape SCSI devices.

Request for access fails

This problem affects managed systems.

Problem

A request for access fails and the managed system remains locked.

Investigation

To correct this problem, make sure that the following conditions are met:

- You are using the correct user ID and password.
- If the managed system accepts encrypted communications only, make sure that the management server has encryption enabled also.

Request for access fails on a Level-1 managed system

This problem affects Level-1 managed systems and management servers.

Problem

A request for access on a Level-1 managed system fails.

Investigation

The system clock on either the managed system or the management server is incorrect. To correct this problem, determine which system clock is incorrect: the managed system or the management server.

Note: If the management server clock is incorrect, that can affect many managed systems. If the managed system clock is incorrect, then only that managed system is affected.

Check the ras.log file on the managed system for any messages that indicate a certificate was not valid or expired.

- If a message states that the certificate was not valid, the system clocks are not synchronized and the managed system clock is *behind* the management server clock. For example, the managed system is set for 2004, but the validity of the management server certificate does not start until 2006.
- If a message states that the certificate has expired, the system clocks are not synchronized and the managed system clock is *ahead* of the management server clock. For example, the managed system is set for 2008, but the validity of the management server certificate ends in 2006.

If the managed system clock is incorrect, complete the following steps:

1. On the system with the incorrect clock, stop the CIMOM by typing:

For Linux	`/etc/ibm/director/diragent/stopcim`
For Windows	`net stop wmicimserver`

2. Correct the affected clock. Make sure the managed system clock is synchronized or a little ahead of the management server clock. Certificates allow up to 24 hours difference.
3. Start the CIMOM by typing:

For Linux	`/etc/ibm/director/diragent/startcim`
For Windows	`net start wmicimserver`

If the management server clock is incorrect, complete the following steps:

1. Correct the affected clock. Make sure the management server clock is synchronized or slightly behind all Level-1 managed systems. Certificates allow up to 24 hours difference.
2. Use the dircli **certmgr** command to generate a new certificate.

Resetting user preferences

This problem affects IBM Director Console.

Problem

A management console user has set their user preferences to a screen resolution that displays some windows off the screen.

Investigation

To correct this problem, reset the user preferences to the default values. User preferences are stored on the management server, not the management console.

If IBM Director Server is installed on a server running Windows, complete the following steps:

1. Log off IBM Director Console on the affected management console.
2. On the management server, start the Windows Registry Editor.
3. Navigate to the following directory:

   My Computer\HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Director\User
4. Navigate to the applicable domain directory for the affected user.

5. Navigate to the applicable user directory for the affected user. The IBM Director Console user preferences are stored in this directory.
6. Delete this directory.

If IBM Director Server is installed on a server running Linux, complete the following steps:

1. Log off IBM Director Console on the affected management console.
2. On the management server, navigate to the following directory:

   ```
   /opt/ibm/director/data
   ```
3. Delete the following file:

   User.*username*

   where *username* is the user ID of the affected IBM Director Console user.

Wrong time zone is displayed

This problem affects the event viewer.

Problem

The wrong time zone is displayed.

Investigation

When the time zone setting is changed on the managed system, the time that is shown in the event viewer is not adjusted. Restart the managed system to ensure that the correct time zone is displayed.

IBM Director Agent troubleshooting

Use this section to troubleshoot and resolve problems with IBM Director Agent.

Some IBM Director Agent problems might be problems with other IBM Director components. Review all troubleshooting topics for possible solutions. For additional troubleshooting information, see the *IBM Director Release Notes*.

A PCI adapter with logical disks cannot be stopped

This problem affects IBM Director Agent. The problem occurs only on systems running Windows 2000.

Problem

A PCI adapter with logical disks cannot be stopped using the Unplug or Eject Hardware window.

Investigation

To correct this problem, install Microsoft Windows 2000 Service Pack 4.

Event ID 2003 warning message appears in the application event log

This problem affects IBM Director Agent. The problem occurs only on systems running Windows 2000 with Internet Information Services (IIS) installed.

Problem

An event ID 2003 warning message appears in the application event log when you start System Monitor and add counters.

Investigation

Microsoft has identified this as a problem. For more information, see Microsoft Knowledge Base Article 267831.

Remote Access Connection Manager service fails to start

This problem affects IBM Director Agent. The problem occurs only on systems running Windows.

Problem

The Remote Access Connection Manager service fails to start and the following error message is displayed:

```
The service cannot be started, either because it is disabled or because it has
no enabled devices associated with it.
```

Investigation

This problem is solved by a Microsoft update. See Microsoft Knowledge Base article 830459 for more information.

The event log is full

This problem affects IBM Director Agent. The problem occurs only on systems running Windows 2000.

Problem

The event log is full. This problem occurs on servers when NetBIOS is enabled and IBM Director is installed. Errors are generated until the event log is full.

Investigation

To correct this problem, uninstall and then reinstall the device driver for the NIC.

Win32_DiskDrive.Size error message

This problem affects IBM Director Agent. The problem occurs only on systems running Windows.

Problem

The following report is generated:

```
Win32_DiskDrive.Size is less than Win32_DiskPartition.Size for a removable medium
that has been formatted as a single partition.
```

Investigation

The following hard disk drives are not supported by Windows:
- Optical
- Iomega
- Jaz

Microsoft identified this as a Windows Management Instrumentation (WMI) problem.

IBM Director tasks troubleshooting

Use this section to troubleshoot and resolve problems with IBM Director tasks, except Software Distribution.

Some IBM Director task problems might be problems with other IBM Director components. Review all troubleshooting topics for possible solutions. For additional troubleshooting information, see the *IBM Director Release Notes*.

Note: See the Software Distribution troubleshooting section for information about that task.

Cannot access RDM help from IBM Director Console

This problem affects the Remote Deployment Manager (RDM) task.

Problem

You cannot access the Remote Deployment Manager (RDM) help from IBM Director Console in the following situations:
- An installation that includes IBM Director, version 4.20 or 4.20.2, and RDM 4.11 (including any of the RDM 4.11 Updates 1, 2, or 3)
- An installation that includes IBM Director 4.12 and RDM 4.11 (including any of the RDM 4.11 Updates 1, 2, or 3) is upgraded to use IBM Director, version 4.20 or 4.20.2

Investigation

To access the RDM help files, complete the following steps:
1. Navigate to the *director_install*\classes\doc\ibm directory, where *director_install* is the directory where IBM Director is installed.
2. Move the rdm directory to the *director_install*\classes\doc\en\ibm directory.

CIM Browser errors due to return of large amounts of CIM data

This problem affects the CIM Browser task. The problem occurs only on Level-2 managed systems running Windows.

Problem

When you attempt to enumerate a system running Windows, large amounts of CIM data are returned, causing errors in the CIM Browser.

Investigation

Do not attempt to enumerate the instances of the following classes:

- root/cimv2:CIM_DirectoryContainsFile
- root/cim2:Win32_Subdirectory

Those CIM classes have instances for every file and directory on every disk in your server. If you attempt to enumerate these classes, the managed system or management server might run out of memory.

error while loading shared libraries: libc.so.6:

This problem affects the Process Management task. It occurs only on managed systems running Linux.

Problem

When running the Process Management task or any action that uses Linux operating-system commands, you might receive the following error message:

```
error while loading shared libraries: libc.so.6: cannot open shared object file:
No such file or directory
```

This error is caused by Java setting the environment variable LD_ASSUME_KERNEL for its context which is incompatible with commands launched from the Java environment.

Investigation

To correct this problem, set the LD_ASSUME_KERNEL variable to an empty value as part of the action. For example, when running the ps command as a Process Management action, instead of issuing the command `ps -ef`, use `LD_ASSUME_KERNEL=;ps -ef`

Note: This problem affects at least the following commands:

- ps
- top
- pgrep
- pstree

FRU information does not appear when inventory is collected

This problem affects the Inventory task.

Problem

Field-replaceable unit (FRU) information does not appear when inventory is collected.

Investigation

If a system is not connected to the Internet when IBM Director Agent is installed, the FRU inventory might be empty. To populate the FRU inventory, run the GETFRU command.

In addition, make sure that the GETFRU command can reach the IBM Support FTP site through your firewall. For the GETFRU command to run successfully, the managed system must have firewall access through a standard FTP port.

Inventory collection fails

This problem affects IBM Director Server.

Problem

Inventory collection fails if IBM Director Server is an older version than IBM Director Agent or IBM Director Core Services.

Investigation

IBM Director Server must be the same version or a later version that IBM Director Agent and IBM Director Core Services. Update your IBM Director Server installation.

Mass Configuration fails to configure Asset ID

This problem affects the Mass Configuration and Asset ID tasks.

Problem

When you use the Mass Configuration task to configure Asset ID, the configuration fails.

Investigation

The managed system does not have sufficient data space. When the size of the configuration is larger than that of the remaining data space, the configuration fails (although there is no indication that a failure has occurred). This is a limitation of the data save area. Make sure that, for each byte of data, the managed system has the same amount of space in the data save area.

Network Configuration task displays a computer name incorrectly

This problem affects the Network Configuration task.

Problem

When you use the Network Configuration task to change the computer name of a managed system, the computer name is not displayed correctly.

Investigation

To correct this problem, restart the managed system.

Network Configuration task fails to modify the gateway address

This problem affects the Network Configuration task.

Problem

When running the Network Configuration task against a managed system to change a static IP address to an IP address generated by a DHCP server, the gateway address is not modified automatically.

Investigation

To correct this problem, manually set the gateway address to DHCP on the managed system.

Network Configuration task reverses IP addresses in the WINS pane

This problem affects the Network Configuration task. The problem occurs only on Level-2 managed systems running Windows Server 2003.

Problem

When you run the Network Configuration task and view the WINS pane, the IP addresses for the primary and secondary Windows Internet Naming Service (WINS) servers are reversed.

Investigation

This is caused by a Microsoft implementation of a CIM class. The correct IP addresses are assigned in the system Network Properties.

Non-English-language keyboard fails during a remote-control session

This problem affects the Remote Control task.

Problem

When you use a non-English-language keyboard during a remote-control session, some of the keys might not work.

Investigation

To correct this problem, make sure that the inventory has been collected before you use the Remote Control task.

Performance monitor results for CPU utilization are incorrect

This problem affects Level-2 managed systems. The problem occurs only on xSeries 450 servers running 64-bit versions of Windows.

Problem

The %Processor Utilization attribute for the _Total and CPU1 processors provides incorrect data in IBM Director and Microsoft Windows performance monitors. All other attributes are monitored correctly.

Investigation

This problem is solved by a Microsoft update. See Microsoft Knowledge Base Article 838987 for more information.

Resource Monitors tasks displays incorrect attribute names

This problem affects the Resource Monitors task. The problem occurs only on managed systems running Windows.

Problem

When you run the Resource Monitors tasks against multiple managed systems, incorrect attribute names might be displayed for the network adapters.

Investigation

The incorrect attribute names are displayed in the Available Resources pane of the Resource Monitors window, when you click **Director Agent** → **TCP/IP Monitors**.

To view the correct attribute names for the network adapters, click **Director Agent** → **Windows Performance Monitors** → **Network Interface**.

The Remote Control or Software Distribution tasks fail

This problem affects the Remote Control and Software Distribution tasks.

Problem

The Remote Control task fails when both of the following conditions are true:

- You are running the task against a managed system that is behind a firewall.
- You are simultaneously distributing a software package to that managed system.

Investigation

The Remote Control and Software Distribution tasks both use session support to increase the rate of data transmission. Session support within TCP/IP causes data to flow through a nonreserved port that is different from the one that IBM Director typically uses for communication. Most firewalls do not allow the data to be transmitted through this different port. You can disable session support by creating an INI file on the managed system. In the IBM\Director\bin directory on the managed system, create a file named tcpip.ini that contains the following command:

```
SESSION_SUPPORT=0
```

If more than one TCP/IP option is selected in the network driver configuration of the managed system, you must create an INI file for each entry. Name these files tcpip.ini, tcpip2.ini, tcpip3.ini, and so on. After creating the files, restart the managed system.

ServeRAID Manager cannot delete logical drives

This problem affects Level-2 managed systems. It occurs only on servers installed with an IBM ServeRAID-7t Serial ATA (SATA) controller or ServeRAID-8i Serial Attached SCSI (SAS) controller.

Problem

After you have restarted a managed system that has a ServeRAID-7t or ServeRAID-8i controller, you cannot delete any logical drives using the ServeRAID Manager task.

Investigation

To delete a logical drive, use one of the following programs:

ServeRAID Manager (Standalone Edition) in startable-CD mode
You can access this program by starting the server with the *ServeRAID Support* CD in the CD drive.

Adaptec RAID Configuration Utility
You can access this program during system startup by pressing Ctrl+A.

Trap destinations are missing from the SNMP agent table

This problem affects the Configure SNMP Agent task.

Problem

Trap destinations are missing from the SNMP agent table.

Investigation

A table displays only the first trap destination in the SNMP configuration interface when multiple communities and traps are associated with each community. The IBM Director inventory stores only the first value of an array-valued property, such as the SNMP trap destination.

You cannot change an attribute value for a MIB file

This problem affects the SNMP Browser task.

Problem

You cannot change an attribute value for a MIB file.

Investigation

Make sure that the following conditions are met:

- IBM Director uses a community name that allows write access to the MIB file.
- The MIB file is writable.
- The MIB file has a value that you can set to be displayed in the SNMP Browser.
- The compiled MIB file is associated with the value that you want to change.

Troubleshooting ServeRAID Manager

Could not copy the configuration from the drives: controller [number]

Explanation:

- There is no configuration on any of the physical drives that are attached to the controller.
- There are no physical drives attached to the controller.
- The controller does not support one or more features of the drive configuration.

Action:

If you have physical drives attached to the controller and the physical drives contain a valid controller configuration, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, the configuration data has been lost. Reconfigure the controller using the ServeRAID Manager.

Could not restore the configuration to the factory-default settings: controller [number]

Explanation:

The ServeRAID Manager could not restore the factory-default settings because of one of the following:

- Your configuration contains a large number of drives (10 or more).
- A hardware error occurred.

Action:

- If your configuration contains a large number of drives and all the logical drives were deleted, ignore this error. If all the logical drives were not deleted, follow the actions listed for the following list item.
- If your configuration does **not** contain a large number of drives:
  1. Verify that the controller, cables, and physical drives are installed properly.
  2. Retry the command.
  3. If the command still fails, restart the server and retry the command.

If the problem persists, complete the following steps:

1. Disconnect all the SCSI cables from controller.
2. Restore to the factory-default settings. If this does not work, contact your service representative.
3. Connect the SCSI cables to the controller.
4. If step 2 was successful, restore to the factory-default settings.

Could not unblock logical drive [number]: controller [number]

Explanation:

The specified logical drive could not be unblocked because of one of the following:

- The rebuild operation was not completed successfully.
- A hardware error occurred.

Action:

Verify that the rebuild operation was completed successfully. If it did, then complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.

3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5.

   If the problem persists, contact your service representative.

Could not create a hot-spare drive: controller [number], < drive location> Could not create a standby hot-spare drive: controller [number], < drive location>

Explanation:

The hot-spare drive or standby hot-spare drive could not be defined because a hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

Could not delete array: controller [number], array [letter]

Explanation:

The array could not be deleted because a hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not delete all of the arrays: controller [number]

Explanation:

A hardware error occurred.

Action:

Delete the arrays by using Restore to factory-default settings.

If the problem persists, contact your service representative.

Could not delete logical drive: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Logical drive was not initialized: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

Verify that the specified logical drive is not offline. If the logical drive **is** offline, replace the failed physical drives and restore the data from tape backup.

If the specified logical drive is **not** offline, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, contact your service representative.

Could not start the logical drive synchronization: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

Verify that the specified logical drive is not offline or critical (that is, one physical drive that is offline in a RAID level-1, 1E, 5, 5E, 10, 1E0, or 50 logical drive). If the logical drive **is** critical, replace the failed physical drive. If the logical drive **is** offline, replace the failed physical drives and restore the data from tape backup.

If the specified logical drive is **not** offline or critical, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, contact your service representative.

Could not set the drive to online: controller [number], < drive location>

Explanation:

The specified drive could not be brought online because a hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, replace the specified drive.

Could not remove the defunct drive: controller [number], < drive location>

Explanation:

The defunct drive could not be removed because a hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Retry the command.
3. If the command still fails, restart the server and retry the command.

Could not replace the defunct drive: controller [number], < drive location>

Explanation:

The defunct drive could not be replaced because a hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

Could not delete the hot-spare drive: controller [number], < drive location> Could not delete the standby hot-spare drive: controller [number], < drive location>

Explanation:

The hot-spare drive or standby hot-spare drive could not be deleted because a hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.

3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, replace the specified drive.

Could not set the merge-group number: controller [number], logical drive [number]

Explanation:

The specified merge-group number could not be set because a hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5.

   If the problem persists, contact your service representative.

Could not blink the device lights

Explanation:

The device lights could not flash because of one of the following:

- The physical drives are not managed by an enclosure.
- A hardware error occurred.

Action:

Verify that the device is managed by an enclosure (SAF-TE) device on the SCSI drive channel. If it is, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5.

   If the problem persists, contact your service representative.

The battery-backup cache device needs a new battery: controller [number]

Explanation:

The battery and battery-backup cache are malfunctioning because of one of the following:

- The battery and battery-backup cache device are installed improperly.
- The battery is low, and the battery-backup cache device must be replaced.

Action:

Verify that the battery and battery-backup cache device are installed properly. If they are installed properly, contact your service representative.

The battery-backup cache device is defective: controller [number]

Explanation:

The battery-backup cache device is installed improperly or is defective.

Action:

1. Verify that the battery-backup cache device is installed properly.
2. If the battery-backup cache device **is** installed properly but is defective, contact your service representative.

Background polling commands are not responding: controller [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the drives.
3. Restart the server.
4.

   If the problem persists, contact your service representative.

Commands are not responding: controller [number]

Explanation:

- If the controller status displays ″Not responding,″ a hardware error occurred.
- If the controller status displays ″Bad configuration,″ a configuration error occurred.

Action:

- If the controller status displays ″Not responding,″ complete the following steps:
  1. Verify that the controller, cables, and physical drives are installed properly.
  2. Verify that there is power to the physical drives.
  3. Restart the server.

  If the problem persists, contact your service representative.
- If the controller status displays ″Bad configuration,″ complete the following steps:
  1. In the Enterprise view, click the specified (controller).
  2. If the controller is *not* configured, click **Actions** → **Restore to factory-default settings**. If the controller *is* configured, click **Actions** → **Copy configuration from drives**. If ″Copy configuration from drives″ does not correct the problem, click **Actions** → **Restore to factory-default settings**; then, recreate the configuration using the ServeRAID Manager.

Rebuild failed: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. If the command still fails, restart the server and retry the command.
4. If the problem persists, replace the specified drive.

Synchronization failed: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

Verify that the specified logical drive is not offline or critical (that is, one physical drive that is offline in a RAID level-1, 1E, 5, 5E, 10, 1E0, or 50 logical drive). If the logical drive **is** critical, replace the failed physical drive. If the logical drive **is** offline, replace the failed physical drives and restore the data from tape backup.

If the specified logical drive is **not** offline or critical, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, contact your service representative.

Migration [logical-drive migration type] failed: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

Determine if one or more physical drives that are part of the specified logical drive have failed. If such a failure has occurred, restore the data from a tape backup. Otherwise, do the following:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

Compression failed: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

Determine if one or more physical drives that are part of the specified logical drive have failed. If such a failure has occurred, restore the data from a tape backup. Otherwise, do the following:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

Decompression failed: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

Determine if one or more physical drives that are part of the specified logical drive have failed. If such a failure has occurred, restore the data from a tape backup. Otherwise, do the following:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

defunct drive: controller [number], < drive location>

Explanation:

A hardware error occurred.

Action:

- If the specified physical drive is part of an array, refer to the event pertaining to the logical drives in that array for additional information.
- If the specified physical drive is **not** part of an array, contact your service representative.

PFA detected for drive: controller [number], < drive location>

Explanation:

The physical drive is going to fail.

Action:

Contact your service representative.

Logical drive is offline: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

Contact your service representative.

Logical drive is critical: controller [number], logical drive [number]

Explanation:

A physical drive is defunct in the specified logical drive. The data on this logical drive is at risk. If another physical drive fails, the data might be lost.

Action:

- If a rebuild operation is in progress, wait until the rebuild is complete.
- If a rebuild operation is **not** in progress, replace the failed physical drive with a new physical drive. After the physical drive is replaced, a rebuild operation will start automatically. Refer to the troubleshooting chapter of the *IBM ServeRAID User's Reference.* .

Logical drive is blocked: controller [number], controller [number] [number]

Explanation:

When the ServeRAID controller performs a rebuild operation on an array, it reconstructs the data that was stored in RAID level-1 and RAID level-5 logical drives. However, the ServeRAID controller cannot reconstruct the data that was stored in any RAID level-0 logical drives in that array. The data in the RAID level-0 logical drives is blocked when the ServeRAID controller detects that the array is valid, but the data might be damaged.

Action:

Restore the data from tape.

Could not communicate with controller: controller [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5.

   If the problem persists, contact your service representative.

User name or password is not valid

Explanation:

An incorrect or undefined user name or password was typed.

Action:

- Verify that you specified a valid user name and password for the remote system. Passwords are case sensitive.
- If you do not require security, you can disable security on the remote system. In this case, a user name and password are not required to access the system remotely.

ServeRAID Manager failed to start at port number [number]

Explanation:

The ServeRAID Manager was unable to use the specified port number on the local system. Another session of the ServeRAID Manager or another application is currently using the port.

Action:

Change the port number for the system by doing the following:

1. Click **Agent actions** → Configure.
2. Click the **General settings** tab.
3. In the Agent base port number field, type a new value for the port number. The default port number for local access is 34571. **Note:** When changing the port number, consider the following:
   - This change does not take effect until you restart the ServeRAID Manager.
   - Any system that is accessing this system remotely must have the port number changed to match the value you type. On the remote system, you can change this value in the ″ Add remote system″ window.
   - Any system that contains this system in its Notification list must change the port number to the new port number. On the remote system, you can change the port number in the Notification Manager by clicking **Modify system**.

No controllers were found in this system.

Explanation:

The ServeRAID Manager did not detect any controllers in the system.

Action:

If you know there are controllers in the system, the ServeRAID Manager might not identify the controllers because of the following:

- The device driver is not loaded.
- (Linux, UnixWare, and OpenServer only) You are not running the ServeRAID Manager as ″root.″ The ServeRAID Manager cannot access the device driver unless it is run with root authority. Without access to the device driver, the ServeRAID Manager cannot identify any controllers.

If the problem persists, contact your service representative.

Host name [ID] is unknown or the network is down

Explanation:

The ServeRAID Manager could not access the remote system.

Action:

Verify the following:

1. The remote system is on.
2. Both the local and remote systems are connected to the network.
3. TCP/IP networking support is configured on both the local and remote systems.
4. The network is functioning. Verify that you can ping the remote system.

Failed to connect to host name [ID] due to incompatible versions [Local=number Remote=number]

Explanation:

The local and remote versions of the ServeRAID Manager are incompatible versions.

Action:

Upgrade the older version of the ServeRAID Manager to the newer version.

Unable to connect to the remote system

Explanation:

The ServeRAID Manager could not communicate with the ServeRAID Manager agent on the remote system.

Action:

1. Verify that the ServeRAID Manager agent or console is running on the remote system.
2. Verify that the port number specified for the remote system in the ″ Add remote system″ window matches the value on the remote system. You can verify the port number on the remote system by going to the ServeRAID Manager running on that system and doing the following:
   a. Click **Agent actions** → **Configure**.
   b. Click the **General settings** tab.
   c. Verify the value for the port number. The default port numbers for remote access are 34571-34574.
3. Verify the following:
   a. The remote system is on.
   b. Both the local and remote systems are connected to the network.
   c. TCP/IP networking support is configured on both the local and remote systems.
   d. The host name of the remote system is defined in the Domain Name Server or a Hosts file, if you are trying to connect using a host name.
   e. The network is functioning.

Array [letter] storage space still available.

Explanation:

You have configured an array that still contains free space.

Action:

Before you apply this configuration, you can do either of the following:

- Increase the size of the new logical drives.
- If there is enough free space, create more logical drives in this array by returning to the ″Configure logical drive″ window in the Configuration wizard.

Physical drives in array [letter] contain unusable space

Explanation:

You have configured an array using physical drives of different sizes. If you configure an array using physical drives of different sizes, you cannot use all the physical drive space.

Action:

To create optimal array configurations, include only physical drives of the same size in any one array.

Hot spare is too small for use by at least one array.

None of the logical drives in the specified array support hot-spare drives.

At least one array is too large to use the hot spare drive [number]. Replace the specified drive with a larger drive.

Explanation:

This hot-spare drive will not work for any defined array because of one of the following:

- None of the existing array support hot-spare drives.
- The hot-spare drive must have the same capacity as or larger capacity than the smallest physical drive in the array.
- If any of the physical drives in your arrays fail, the specified hot-spare drive cannot replace the failed drive.

Action:

Remove the specified hot-spare drive and add a drive of the appropriate size.

Could not start logical drive migration: controller [number] logical drive [number]

Explanation:

The logical-drive migration could not start because of one of the following:

- A rebuild, synchronization, or migration operation is currently in progress on the specified controller.

- A hardware error occurred.

Action:

Verify that there is no rebuild, synchronization, or migration operation currently in progress on this controller. If you see a progress indicator in the status bar or (in animation) in the Logical devices view, one of these operations is in progress and you must wait for the operation to be completed. Otherwise, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not enable the hot-swap rebuild operation: controller [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5.
   If the problem persists, contact your service representative.

Could not create logical drive: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Verify that none of the physical drives are defunct.
4. Retry the command.
5. If the command still fails, restart the server and retry the command.
6.
   If the problem persists, contact your service representative.

Logical drive was not initialized: controller [number], logical drive [number]

Explanation:

The logical drive has not been initialized.

Action:

Do not store data on this logical drive until you initialize the logical drive. Initialize the logical drive by doing the following:

1. In the Logical devices view, click the specified logical drive.
2. Click **Actions** → **Initialize**; or, for the ServeRAID-7t or ServeRAID-8i controller, click, **Actions** → **Clear**.

Logical drive must be synchronized: controller [number], logical drive [number]

Explanation:

You must synchronize the specified logical drive before storing data on it.

Action:

Synchronize the logical drive by doing the following:

1. In the Logical devices view, click the specified logical drive.
2. Click **Actions** → **Synchronize**.

[Number] ready drives still available.

Explanation:

You have ready drives still available for configuration.

Action:

You can configure these drives as new arrays, add them to other new arrays, or define them as hot-spare drives.

Cannot communicate with the remote system

Explanation:

1. The local ServeRAID Manager has lost communication with the ServeRAID Manager on the remote system.
2. The local ServeRAID Manager has lost communication with the server on which it is installed. The ServeRAID Manager agent might fail leaving the ServeRAID Manager console available, but unable to communicate with the server.

Action:

For explanation 1, use the following actions:

- If the ServeRAID Manager is running in network mode, use the following actions:
  1. Verify that the ServeRAID Manager agent or console is running on the remote system.
  2. Verify that the port number specified for the remote system in the ″ Add remote system″ window matches the value on the remote system. You can verify the port number on the remote system by going to the ServeRAID Manager running on that system and doing the following:
     a. Click **Agent actions** → **Configure ServeRAID Manager agent**.

   b. Click the **General settings** tab.
   c. Verify the value for the port number. The default port numbers for remote access are 34571-34574.
3. Verify the following:
   a. The remote system is on.
   b. Both the local and remote systems are connected to the network.
   c. TCP/IP networking support is configured on both the local and remote systems.
   d. The host name of the remote system is defined in the Domain Name Server or a Hosts file, if you are trying to connect using a host name.
   e. The network is functioning.

- If the ServeRAID Manager is running as a plug-in to another program, use the following actions:
  1. If the connection fails but has worked before, re-initiate the ServeRAID Manager task by closing the current ServeRAID Manager window and then dragging the RAID task icon onto the appropriate system.
  2. If the connection fails and has **not** worked before, verify the following:
     a. The remote system is on.
     b. The program using the ServeRAID Manager agent as a plug-in is installed on the remote system.
     c. Both the local and remote systems are connected to the network.
     d. TCP/IP networking support is configured on both the local and remote systems.
     e. The network is functioning.

     Note:
       a. If you install a program that uses the ServeRAID Manager, the ServeRAID Manager agent is installed also. The ServeRAID Manager agent is started automatically by the program using it.
       b. When using the ServeRAID Manager agent as a plug-in to another program, the ServeRAID Manager does not recognize port numbers.

For explanation 2, use the following actions:

1. Verify that the ServeRAID Manager agent or console is running on the local system. View processes or services to verify that the agent is running.
2. Restart the ServeRAID Manager console.
3. Verify that TCP/IP networking support is configured on the local system.
4. Verify that the network is functioning.
5. Restart the server.

Error getting controller configuration.

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.

3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, complete the following steps:

1. Restore to factory-default settings.
2. Recreate the configuration.

Agent is running in local only mode.

Explanation:

- If you specified **-l** (that is, minus and a lowercase letter L) as a command-line parameter, the ServeRAID Manager starts in local-only mode. The ServeRAID Manager is not network-enabled. In this mode, you can configure and monitor the local system only. You cannot configure or monitor remote systems. If you added remote systems to your Enterprise view, they are not visible in this mode. They will be visible in the tree when you start the ServeRAID Manager in non-local-only mode.
- If the network is unreachable, the ServeRAID Manager starts in local-only mode. An additional event is displayed in the event viewer indicating that the network is unreachable.
- If the port number is in use on the local system, the ServeRAID Manager starts in local-only mode. An additional event is displayed in the event viewer indicating that the network is unreachable.

Action:

- If the network is unreachable, check the event viewer for details on any networking errors.
- If the port number is in use, exit from the application using the port number, or specify a different port number. Complete the following steps to specify a different port number:
  1. Click **Agent actions** → **Configure**.
  2. Click the **General settings** tab.
  3. In the Agent base port number field, type a new value for the port number. The default port number for local access is 34571. **Note:** When changing the port number, consider the following:
     - This change does not take effect until you restart the ServeRAID Manager.
     - Any system that is accessing this system remotely must have the port number changed to match the value you type. On the remote system, you can change this value in the ″ Add remote system″ window.
     - Any system that contains this system in its Notification list must change the port number to the new port number. On the remote system, you can change the port number in the Notification Manager by clicking **Modify system**.

Networking support is not available.

Explanation:

The network is unreachable.

Action:

Verify the following:

1. The local system is connected to the network.
2. The TCP/IP networking support is configured.
3. The network is functioning.

Could not send the event to the system.

Explanation:

An event could not be sent to the remote system.

Action:

1. Verify that the ServeRAID Manager is running on the remote system.
2. Verify that the port number specified in the Notification list matches the value for the start-up port number on the remote system.
   - You can verify the port number for this event in the Notification Manager by using the ″ System properties″ window for this system.
   - You can verify the port number on the remote system by going to the ServeRAID Manager running on that system and doing the following:
     a. Click **Agent actions** → **Configure**.
     b. Click the **General settings** tab.
     c. Verify the value for the port number.
3. Verify the following:
   a. The remote system is on.
   b. Both the local and remote systems are connected to the network.
   c. TCP/IP networking support is configured on both the local and remote systems.
   d. The network is functioning.
4. Verify that the local and remote systems are using compatible versions of the ServeRAID Manager. If they are not, upgrade to the latest version of the ServeRAID Manager.

Failed to connect to host name [ID] at port number [number].

Explanation:

The ServeRAID Manager could not communicate with the ServeRAID Manager client on the remote system.

Action:

1. Verify that the ServeRAID Manager agent or console is running on the remote system.
2. Verify that the port number specified for the remote system in the ″ Add remote system″ window matches the value on the remote system. You can verify the port number on the remote system by going to the ServeRAID Manager running on that system and doing the following:
   a. Click **Agent actions** → **Configure**.
   b. Click the **General settings** tab.
   c. Verify the value for the port number. The default port numbers for remote access are 34571-34574.

3. Verify the following:
   a. The remote system is on.
   b. Both the local and remote systems are connected to the network.
   c. TCP/IP networking support is configured on both the local and remote systems.
   d. The host name of the remote system is defined in the Domain Name Server or a Hosts file, if you are trying to connect using a host name.
   e. The network is functioning.

Note: If the local ServeRAID Manager is trying to connect to a remote system that has multiple network adapters, the local ServeRAID Manager must use a host name to connect to the remote system.

Failed to connect to host name [ID] due to incompatible versions [Local=id Remote=id].

Explanation:

The local and remote versions of the ServeRAID Manager are incompatible versions.

Action:

Upgrade the older version of the ServeRAID Manager to the newer version.

ServeRAID Manager failed to start at port number [number].

Explanation:

The ServeRAID Manager was unable to use the specified port number. The port is currently in use.

Action:

Change the port number for the system by doing the following:

1. Click **Agent actions** → **Configure**.
2. Click the **General settings** tab.
3. In the Agent base port number field, type a new value for the port number. The default port number for local access is 34571. **Note:** When changing the port number, consider the following:
   - This change does not take effect until you restart the ServeRAID Manager.
   - Any system that is accessing this system remotely must have the port number changed to match the value you type. On the remote system, you can change this value in the ″ Add remote system″ window.
   - Any system that contains this system in its Notification list must change the port number to the new port number. On the remote system, you can change the port number in the Notification Manager by clicking **Modify system**.

Refused connection from [remote system ID].

Explanation:

Security is enabled on this system, and the ServeRAID Manager received a remote log-in request containing an incorrect user ID or password.

Action:

The remote system requesting a log-in to this system must specify a valid user ID and password. Be sure the remote system issuing the log-in request specifies a valid user ID and password.

FlashCopy with backup failed: controller [number], logical drive [number]

Explanation:

The FlashCopy failed because a hardware error occurred. The specified logical drive might be offline.

Action:

If the source logical drive is offline, replace the failed physical drives and restore the data from tape backup. If the target logical drive is offline, replace the failed physical drives. FlashCopy will not work when the source or target logical drives are offline.

If the source or target logical drives are *not* offline, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5.

   If the problem persists, contact your service representative.

Could not import configuration with more than eight drives into a ServeRAID-3L controller.

Explanation:

The ServeRAID-3L controller does not support more than eight drives in an array when the stripe-unit size is 32 KB or 64 KB.

Action:

Do **one** of the following:

- Replace the ServeRAID-3L controller with a ServeRAID-3H controller. The configuration stored on the physical drives is supported only on a ServeRAID-3H controller or later hardware. *or*
- Create a new configuration that does not contain more than eight drives in an array with the stripe-unit size being 32 KB or 64 KB.

  Attention: ⊙ When creating a new configuration, you will destroy the data currently on the physical drives.

Could not copy the configuration from the drives: controller [number], < drive location>

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

Possible non-warranted physical drive found: controller [number], < device location>

Explanation:

The ServeRAID Manager has detected that the specified physical drive might not be a qualified and warranted part.

Action:

This message is an alert only. The physical drive will continue to function and the ServeRAID Manager will not make any changes or modifications to the configuration.

For further assistance, contact the source from which you obtained the specified physical drive.

Could not set the host name: controller [number]

Explanation:

A hardware error occurred.

Action:

Verify that the cluster partner system is turned off. If it is, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

Could not set the partner host name: controller [number]

Explanation:

A hardware error occurred.

Action:

Verify that the cluster partner system is turned off. If it is, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not change the rebuild rate: controller [number], < device location>

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not change the stripe-unit size: controller [number], < device location>

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

Could not change the write-cache mode: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not change the SCSI transfer speed: controller [number], channel [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.

4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not enable unattended mode: controller [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not disable unattended mode: controller [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not enable read cache: controller [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not disable read cache: controller [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

5. If the problem persists, contact your service representative.

Could not set the SCSI initiator ID: controller [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Note: (Cluster and failover environment only) Verify that the cable from the second system (for clustering) or controller (for failover) is disconnected from the SCSI backplane.

Could not switch the active and passive controllers.

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not clear the controller event logs for system [number].

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not copy the configuration from the non-shared logical drives (merge group [number]): controller [number]

Explanation:

A hardware error occurred.

Action:

1. Delete the arrays by using Restore to factory-default settings.
2. Set the controller name, the partner controller name, and the SCSI initiator IDs.
3. Retry the command.

If the problem persists, contact your service representative.

Could not change the BIOS-compatibility mapping to [Extended or Limited]: controller [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Could not change the write-cache mode: controller [number], channel [number], SCSI ID [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, replace the specified drive.

Enclosure device is not responding: controller [number], channel [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Enclosure fan [number] is malfunctioning: controller [number], channel [number]

Explanation:

A hardware error occurred.

Action:

Verify that the fan in the enclosure device is installed properly. If it is, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, replace the specified fan.

Enclosure power supply [number] is malfunctioning: controller [number], channel [number]

Explanation:

A hardware error occurred.

Action:

Verify that the power supply in the enclosure device is installed properly. If it is, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, replace the specified power supply.

Enclosure temperature is out of normal range: controller [number], channel [number]

Explanation:

A hardware error occurred.

Action:

Verify that the fans in the enclosure device are installed properly and working. If they are, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Retry the command.
3. If the command still fails, restart the server and retry the command.
4. If the problem persists, contact your service representative.

Could not save the event logs: controller [number]

Explanation:

A hardware error occurred.

Action:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.
5. If the problem persists, contact your service representative.

Version mismatch detected: controller [number]

Explanation:

The version of the BIOS, firmware, and device driver are not compatible.

Action:

Install a compatible version of the BIOS, firmware, and driver for this controller. You can download the latest version from the

IBM Support Web site

Compaction failed: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

Determine if one or more physical drives that are part of the specified logical drive have failed. If such a failure has occurred, restore the data from a tape backup. Otherwise, do the following:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

Expansion failed: controller [number], logical drive [number]

Explanation:

A hardware error occurred.

Action:

Determine if one or more physical drives that are part of the specified logical drive have failed. If such a failure has occurred, restore the data from a tape backup. Otherwise, do the following:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.

3. Retry the command.
4. If the command still fails, restart the server and retry the command.

Battery has exceeded normal operating temperature: battery controller [number]

Explanation:

Battery temperature has exceeded 50 degrees Celsius.

Action:

Check operating environment: verify that the controller is installed properly, that the server has adequate ventilation, and so on. If the problem persists, the battery may be bad or the server may have a problem. Contact your service representative.

Network connection was not found and/or host name was not resolved.

Explanation:

The ServeRAID Manager could not communicate with the network or establish a network connection.

Action:

- If your workstation is not connected to the network, change your workstation name to "localhost" or add an entry for your workstation to your local host file.
- If your workstation is connected to the network, make sure that your host name is resolvable over DNS; for more information, contact your network administrator.

One or more logical drives contain a bad stripe: controller [number], logical drive [number]

Explanation:

The Bad Stripe Table (BST) provides a means of recovering most data on a logical drive after multiple hardware errors prevent access to a logical drive stripe. An entry in the BST indicates that the data contained in a stripe has been lost.

While many conditions can produce a Bad Stripe Table entry, the most common cause is an error accessing one of the stripe units within a stripe of a critical logical drive. A single stripe unit failure is correctable and recoverable but two or more failures within the same redundant RAID stripe are not.

For example, in a critical RAID-5 array, in which one of the drives in the array is defunct, a stripe will be marked bad with an entry in the BST if a non-recoverable media error occurs when accessing one of the other drives of the array.

Once an entry is logged in the BST, the controller will return an error code to the driver whenever the host system tries to access a Logical Block Address (LBA) within the affected stripe. This is one immediate indication that some part of the logical drive is unusable.

Note: It is not possible to correlate the bad stripe with a specific file in the operating system.

Action:

- Check the ServeRAID Manager event logs to identify the affected logical drive(s).
- Because the data has been lost, the only way to recover from this condition is to do the following:
  1. Delete the array.
  2. Recreate the array and its logical drives.
  3. Restore the data from backup media.
- The alternative is to take the entire logical drive offline, thus resulting in the loss of all data contained on that logical drive.
- To minimize the risk of lost data, be sure to schedule frequent periodic backups.

Exception removing timer from active queue

Explanation:

- There is no configuration on any of the physical drives that are attached to the controller.
- There are no physical drives attached to the controller.
- The controller does not support one or more features of the drive configuration.

Action:

If you have physical drives attached to the controller and the physical drives contain a valid controller configuration, complete the following steps:

1. Verify that the controller, cables, and physical drives are installed properly.
2. Verify that there is power to the physical drives.
3. Retry the command.
4. If the command still fails, restart the server and retry the command.

If the problem persists, the configuration data has been lost. Reconfigure the controller using the ServeRAID Manager.

Set drive to defunct: controller [number], < device location>

Explanation:

The specified physical drive reported an unrecoverable media error.

Action:

Replace the specified physical drive.

Set drive to defunct: controller [number], < device location>

Explanation:

There was an error in communication between the physical drive and the controller.

Action:

1. Verify that there is power to external enclosures and connections.
2. Verify that cables are properly seated.
3. Check for damaged, bent, or pushed pins on the following:
   - Termination (for non-backplane systems)

  - Backplanes
  - Physical drives
  - Controller

If you cannot identify a connection problem, the replacement order for parts is the following:

1. Termination (for non-backplane systems)
2. Cables
3. Physical drives
4. Backplanes
5. Controller

Before replacing a physical drive or controller, verify that it is updated with the latest firmware (microcode) and that they still fail.

Note: If multiple physical drives are marked defunct within the same time period, check the cables and backplanes.

Set drive to defunct: controller [number], < device location>

Explanation:

One of the following occurred:

- The user has removed or rearranged the specified physical drive.
- There was an error in communication between the physical drive and the controller.

Action:

1. Verify that there is power to external enclosures and connections.
2. Verify that cables are properly seated.
3. Check for damaged, bent, or pushed pins on the following:
   - Termination (for non-backplane systems)
   - Backplanes
   - Physical drives
   - Controller

If you cannot identify a connection problem, the replacement order for parts is the following:

1. Termination (for non-backplane systems)
2. Cables
3. Physical drives
4. Backplanes
5. Controller

Before replacing a physical drive or controller, verify that it is updated with the latest firmware (microcode) and that they still fail.

Note: If multiple physical drives are marked defunct within the same time period, check the cables and backplanes.

Set drive to defunct: controller [number], < device location>

Explanation:

One of the following occurred:

- The user has removed or rearranged the specified physical drive.
- A controller might be configured with unattended mode. Changes to the RAID configuration are accepted automatically in unattended mode.
- There was an error in communication between the physical drive and the controller.

Note: (Clustering and failover environments only) If you move a logical drive to another controller, these entries might be normal.

Action:

1. Verify that there is power to external enclosures and connections.
2. Verify that cables are properly seated.
3. Check for damaged, bent, or pushed pins on the following:
   - Termination (for non-backplane systems)
   - Backplanes
   - Physical drives
   - Controller

If you cannot identify a connection problem, the replacement order for parts is the following:

1. Termination (for non-backplane systems)
2. Cables
3. Physical drives
4. Backplanes
5. Controller

Before replacing a physical drive or controller, verify that it is updated with the latest firmware (microcode) and that they still fail.

Note: If multiple physical drives are marked defunct within the same time period, check the cables and backplanes.

Set drive to defunct: controller [number], < device location>

Explanation:

The specified physical drive reported an unrecoverable hardware error.

Action:

Replace the specified physical drive.

Set drive to defunct: controller [number], < device location>

Explanation:

The controller failed.

Action:

Update the controller with the latest firmware (microcode). If the problem persists, replace the controller.

Set drive to defunct: controller [number], < device location>

Explanation:

- There was an error in communication between the physical drive and the controller.
- The controller failed.

Action:

1. Verify that there is power to external enclosures and connections.
2. Verify that cables are properly seated.
3. Check for damaged, bent, or pushed pins on the following:
   - Termination (for non-backplane systems)
   - Backplanes
   - Physical drives
   - Controller

If you cannot identify a connection problem, the replacement order for parts is the following:

1. Termination (for non-backplane systems)
2. Cables
3. Physical drives
4. Backplanes
5. Controller

Before replacing a physical drive or controller, verify that it is updated with the latest firmware (microcode) and that they still fail.

Note: If multiple physical drives are marked defunct within the same time period, check the cables and backplanes.

Set drive to defunct: controller [number], < device location>

Explanation:

The specified physical drive does not match the valid configuration signature after the configuration was imported from the physical drives.

Action:

Do one of the following:

- Rebuild the critical array.
- Remove the specified physical drive.

Set drive to defunct: controller [number], < device location>

Explanation:

There was an error in communication between the physical drive and the controller while the controller was attempting to validate the RAID configuration of the physical drive.

Action:

1. Verify that there is power to external enclosures and connections.
2. Verify that cables are properly seated.
3. Check for damaged, bent, or pushed pins on the following:
   - Termination (for non-backplane systems)
   - Backplanes
   - Physical drives
   - Controller

If you cannot identify a connection problem, the replacement order for parts is the following:

1. Termination (for non-backplane systems)
2. Cables
3. Physical drives
4. Backplanes
5. Controller

Before replacing a physical drive or controller, verify that it is updated with the latest firmware (microcode) and that they still fail.

Note: If multiple physical drives are marked defunct within the same time period, check the cables and backplanes.

Set drive to defunct: controller [number], < device location>

Explanation:

A user used the operating-system utilities to mark the specified physical drive defunct.

Action:

Rebuild the critical array.

Set drive to defunct: controller [number], < device location>

Explanation:

The specified physical drive is not part of a valid configuration but was found in a cluster setup.

Action:

Rebuild the specified physical drive into an array. If you cannot rebuild the physical drive, replace or remove the physical drive.

Set drive to defunct: controller [number], < device location>

Explanation:

The specified physical drive is not part of a valid configuration but was found in a cluster setup.

Action:

Rebuild the specified physical drive into an array. If you cannot rebuild the physical drive, replace or remove the physical drive.

Set drive to defunct: controller [number], < device location>

Explanation:

The specified physical drive is configured as part of another system.

Action:

Rebuild the specified physical drive into an array on this system. If you cannot rebuild the physical drive, replace the physical drive.

Set drive to defunct: controller [number], < device location>

Explanation:

The specified channel or port is not working.

Action:

1. Verify that there is power to external enclosures and connections.
2. Verify that cables are properly seated.
3. Check for damaged, bent, or pushed pins on the following:
   - Termination (for non-backplane systems)
   - Backplanes
   - Physical drives
   - Controller

If you cannot identify a connection problem, the replacement order for parts is the following:

1. Termination (for non-backplane systems)
2. Cables
3. Physical drives
4. Backplanes
5. Controller

Before replacing a physical drive or controller, verify that it is updated with the latest firmware (microcode) and that they still fail.

Note: If multiple physical drives are marked defunct within the same time period, check the cables and backplanes.

This event log entry is informational.

Explanation:

None required.

Action:

None required.

Software Distribution troubleshooting

Use this section to troubleshoot and resolve problems with the Software Distribution task.

Some Software Distribution problems might be problems with other IBM Director components. Review all troubleshooting topics for possible solutions. For additional troubleshooting information, see the *IBM Director Release Notes*.

Backslash (\) symbol is displayed instead of the won symbol

This problem only affects managed systems running Korean installations of Windows.

Problem

In the Distribution Preferences window, the **Share Name** field is filled in with the following example share name by default:

```
WWsystemWshare
```

where `W` represents the won symbol.

However, when you press the won key, the **Share Name** field incorrectly displays the backslash (\) symbol.

Investigation

To correct this problem, complete the following steps:

1. Do not overtype or delete the example share name.
2. Retain the won symbols in the example and replace only `system` and `share` with the system name and share name that you want to use.

   Note: If you press the won key, do not use the backslashes; the backslashes cause redirected distribution to fail.
3. Close the Distribution Preferences window; then, reenter this window, and retain the won symbols in the **Share Name** field example.

Backslash (\) symbol is displayed instead of the yen symbol

This problem only affects managed systems running Japanese installations of Windows.

Problem

In the Distribution Preferences window, the **Share Name** field is filled in by default with the following example share name:

```
¥¥system¥share
```

However, when you press the yen key, the **Share Name** field incorrectly displays the backslash (\) symbol.

Investigation

To correct this problem, complete the following steps:

1. Do not overtype or delete the example share name.
2. Retain the yen symbols in the example and replace only `system` and `share` with the system name and share name that you want to use.

   Note: If you press the yen key, do not use the backslashes; the backslashes cause redirected distribution to fail.
3. Close the Distribution Preferences window; then, reenter this window, and retain the yen symbols in the **Share Name** field example.

Importing an SPB package causes an error message to be displayed

This problem occurs only on systems running Linux.

Problem

If you export a software distribution package to Software Package Bundle (SPB) format and then re-import the package, an error message is displayed.

Investigation

To correct this problem, change the permission levels. From the local command prompt, type the following command:

```
chmod 644 filename.spb
```

IO error, file (\\server\share)\ (package name)not found on managed system

This problem affects the Software Distribution task.

Problem

When a software package is distributed using a redirector share, the following error message is displayed:

```
IO error, file (\\server\share)\ (package name)not found on managed system
(system name)
```

Investigation

This problem occurs if you manually delete a software package from the redirector share. To delete packages from the share, you must use the File Distribution Servers Manager window. Right-click the **Software Distribution** task and click **File Distribution Servers Manager**.

Packages are not streamed from the file-distribution server

This problem only affects systems running Windows.

Problem

Software packages are streamed from the management server, although a file-distribution server is configured for use by the managed systems.

Investigation

To correct this problem, make sure that one of the following conditions is met:

- The file-distribution server is a member of the same domain as the management server.
- The file-distribution server has a trust relationship with the domain where the management server is located.

Packages created with IBM Update Assistant cannot be exported

This problem affects Software Distribution (Premium Edition).

Problem

After you upgrade to Software Distribution (Premium Edition), you cannot export a package that was created with IBM Update Assistant.

Investigation

Delete the software package that was created with Software Distribution (Standard Edition). Re-import the package using IBM Update Assistant in Software Distribution (Premium Edition).

Software package creation fails

This problem affects the Software Distribution task.

Problem

The software package creation fails.

Investigation

If you selected **Get files from the local system**, the Software Distribution task attempted to create the package on the management console. Check the available disk space on the management console. If the management-console disk space is insufficient, the package creation fails.

If you selected **Get files from IBM Director Server**, the Software Distribution task attempted to create the package on the management server. Check the available disk space on the management server. If the management-server disk space is insufficient, the package creation fails.

The Remote Control or Software Distribution tasks fail

This problem affects the Remote Control and Software Distribution tasks.

Problem

The Remote Control task fails when both of the following conditions are true:

- You are running the task against a managed system that is behind a firewall.
- You are simultaneously distributing a software package to that managed system.

Investigation

The Remote Control and Software Distribution tasks both use session support to increase the rate of data transmission. Session support within TCP/IP causes data to flow through a nonreserved port that is different from the one that IBM Director typically uses for communication. Most firewalls do not allow the data to be transmitted through this different port. You can disable session support by creating an INI file on the managed system. In the IBM\Director\bin directory on the managed system, create a file named tcpip.ini that contains the following command:

```
SESSION_SUPPORT=0
```

If more than one TCP/IP option is selected in the network driver configuration of the managed system, you must create an INI file for each entry. Name these files tcpip.ini, tcpip2.ini, tcpip3.ini, and so on. After creating the files, restart the managed system.

Unable to export package error message

This problem affects the Software Distribution task.

Problem

When you attempt to export a software distribution package to a network share, the following error message is displayed:

```
Unable to export package.
```

Investigation

The Software Distribution task does not support exporting packages to a network share. Modify the operation to export the package to a local drive.

Upward Integration Module troubleshooting

Use this section to troubleshoot and resolve problems with IBM Director Upward Integration Modules (UIMs).

Some upward integration module problems might be problems with IBM Director Agent. Review all troubleshooting topics for possible solutions. For additional troubleshooting information, see the *IBM Director Release Notes*.

The Mozilla browser does not respond

This problem affects the IBM Director UIM for HP OpenView on Linux only.

Problem

The Mozilla browser stops responding when you access the IBM Director help system.

Investigation

The Mozilla browser requires that the system run the Red Hat Linux Advanced Server, version 2.1 operating system with update 6 or later. You can download the latest Red Hat updates from www.redhat.com/security/updates/notes/

z/VM Center troubleshooting

Use this section to troubleshoot and resolve problems with the z/VM Center task.

For additional troubleshooting information, see the *IBM Director Release Notes*.

SLP does not work

This problem affects the discovery infrastructure for z/VM systems.

Problem

A z/VM managed object could not be discovered automatically using Service Location Protocol (SLP).

Investigation

Possible reasons for the problem are:

- Network communication problems between IBM Director Server and the z/VM manageability access point.
- The SLP service agent on the z/VM manageability access point is not running.
- The z/VM management SLP service could not be registered.

Network

If IBM Director Server and the z/VM manageability access point reside in the same network, ensure that both broadcast and multicast are enabled.

If IBM Director Server and the z/VM manageability access point reside in different networks, SLP multicast packets are not blocked by routers or firewalls.

SLP service agent

Find out if the service agent on your z/VM manageability access is running. From a terminal session on your z/VM manageability access point issue

```
ps aux |grep slp_srvreg
```

To start the SLP service agent issue

```
/etc/init.d/ibmsa start
```

To stop the agent issue

```
/etc/init.d/ibmsa stop
```

Ensure that only one SLP service agent is running on the z/VM manageability access point, issue

```
netstat –anp |grep 427
```

and look if another process is listening.

If an OpenSLP service agent or directory agent is running then stop it by issuing:

```
/etc/init.d/slpd stop
```

z/VM management SLP service

Issue the following commands to discover the z/VM management SLP service:

```
cd /opt/ibm/director/bin
./slp_query --type=management-software.IBM:zvm-management-agent-https
```

If no service is found, check the z/VM management SLP daemon. You can start and stop the daemon with the following commands:

```
/etc/init.d/zvm-mgmtslp start
/etc/init.d/zvm-mgmtslp stop
```

Check /var/log/zvm-mgmtslpd.log to find out whether the SLP service could be registered with an SLP service agent.

z/VM system is not discovered automatically

This problem affects both the Virtual Server Deployment task and the Server Complexes task.

Problem

A z/VM managed object could not be discovered automatically

Investigation

Check the Discovery Preferences. Ensure that Broadcast and Multicast are enabled in the General Discovery settings.

Ensure that your z/VM System discovery settings are suitable for your network environment.

If the problem persists, check your SLP setup (see "SLP does not work" on page 517), or use the "Add z/VM systems" task to create the z/VM managed object manually.

IBM Director server cannot access the IP address for a z/VM managed object

This problem affects both the Virtual Server Deployment task and the Server Complexes task. The problem can occur if the z/VM manageability access point runs on z/VM 5.2 and has multiple network interfaces.

Problem

Discovery for a z/VM managed object yields an IP address for the z/VM manageability access point that IBM Director Server cannot access.

Investigation

Multiple network interfaces result in multiple SLP services for the manageability access point. The z/VM managed object uses the last SLP service that is discovered. This might not be the SLP service with the correct IP address. To resolve the problem modify the SLP setup of the manageability access point. Assure that the z/VM management SLP service is registered only for the network interface that can be accessed by IBM Director Server.

- Check which network interface and IP address can be reached by IBM Director Server. You have assigned the IP address for the connection to the IBM Director Server when setting up communications for the manageability access point. You can verify that an IP address is the correct address by issuing a ping command for it from IBM Director Server.
- In the Linux instance on the manageability access point, open /etc/zvm-mgmtslp/daemon.conf, locate the slp.service.interfaces section, and add a line of this form:

  ```
  slp.service.interfaces=ip-address
  ```

  in the command *ip-address* represents the correct IP address that IBM Director Server can access.

  For example, after adding the line, the slp.service.interfaces section might look like this:

  ```
  # slp.service.interfaces = <ipaddress[:port][,ipaddress[:port][, ...]]>
  #
  # This setting contains a list of ipaddress and port information.
  # These address information is used to compose the SLP service: URL
  # host part. For each ipaddress:[port] entry a SLP service will be
  # registered.
  #
  # Default:
  # slp.service.interfaces
  slp.service.interfaces=9.152.24.149
  ```
- Restart the z/VM management SLP service:

  ```
  /etc/init.d/zvm-mgmtslp restart
  ```
- Optional: Verify your new setup:

  ```
  zvm-mgmtslpd -s -f /etc/zvm-mgmtslp/daemon.conf
  ```

  The output shows the new setting as illustrated in the following example:

  ```
  z/VM Management SLP Configuration
                  cimom.cli.cmd => /opt/ibm/director/cimom/bin/CLI
              cimom.cli.cmdparm => -niq ei IBM_ZvmOperatingSystem
               cimom.port.https => 5989
                     configfile => /etc/zvm-mgmtslp/daemon.conf
                daemon.interval => 3200
                 daemon.logfile => /var/log/zvm-mgmtslpd.log
                 daemon.pidfile => /var/run/zvm-mgmtslpd.pid
                  env.ldlibpath => /opt/ibm/director/lib:/opt/ibm/director/
                                   cimom/lib
         slp.service.interfaces => 9.152.24.149:5989
  ```

z/VM systems management API problems

Problems for the Virtual Server Deployment task and the Server Complexes task can result from z/VM systems management API problems.

Problem

Error messages point to problems with the z/VM systems management API

Investigation

The z/VM Center extension task logs z/VM systems management API information to the syslog daemon. For SUSE LINUX Enterprise Server 9 for IBM System z9, zSeries and S/390 the log file to examine is the system log. For Red Hat Enterprise Linux AS, version 4.0, for IBM System z9, zSeries and S/390 the log to examine resides in /var/log/messages.

Tasks time out

This problem affects both the Virtual Server Deployment task and the Server Complexes task.

Problem

Tasks time out on z/VM.

Investigation

A timeout can be caused by a network failure or by failed components on z/VM. If most tasks run but creating operating system templates and creating new instances of operating systems time out, there might be a problem with the DATAMOVE service machine. The DATAMOVE service machine might not be operational or there might be an active task that did not complete.

1. Assure that the network is operational. Ensure that the connections between the following components are working:
   - The z/VM manageability access point
   - The administrator ID of the z/VM manageability access point
   - The Systems management API (VSMSERVE service machine)
   - The directory manager and its subcomponents
2. Assure that the VSMSERVE, the z/VM manageability access point, and the directory manager environment are all operational. Restart any failed components.
3. Check if the DATAMOVE service machine is blocked by an active task that did not complete. Table 16 lists some commands you can use if DirMaint™ is your directory manager. If you have multiple DATAMOVE service machines, check each of them.

   Table 16. Commands for investigating the status of a DATAMOVE service machine

Command	Purpose
`dirm datamove display errlog`	Check the DATAMOVE error log
`dirm DATAMOVE getcons`	Access the console for DATAMOVE
`dirm getcons`	Access the console for DIRMAINT
`dirm status datamove all`	Show the DATAMOVE status
`dirm q unassigned`	List tasks that have not been assigned (waiting tasks)
`dirm status locked both`	Display locked files
`dirm stat w`	Display available work units
`dirm stat w` *workunit*	Display details of the specified work unit, *workunit*

   For more information, refer to your directory manager documentation. For example, refer to *z/VM Directory Maintenance Facility Command Reference*, SC24-6025.

Linux guest operating system is not associated with a z/VM

This problem affects the Server Complexes task and the Virtual Server Deployment task.

Problem

In the z/VM Server Complexes Association or Linux on System z9 and zSeries Platform Membership association, a Linux guest operating system is not shown under the z/VM system (not associated). Server Complexes task window does not show the corresponding Linux guest system.

Investigation

- Ensure that the Linux guest operating system can be accessed by SSH.
- Ensure that CPINT is installed on the Linux guest operating system.

  From a command prompt on the Linux guest operating system, issue `modprobe cpint` (or `modprobe cpint_mod` or `cpint_load`) followed by `hcp q userid`. If CPINT is installed and working correctly, the command output provides a z/VM user ID and the z/VM system name.
- If it is a Level-1 system, ensure that the Level-1 system supports the SSH protocol.
- If it is a level-2 system, ensure that the zSeriesAgentExt.s390.rpm is installed.

VMRM configurations cannot be applied

This problem affects the Server Complexes task.

Problem

A VMRM configuration cannot be applied.

Investigation

Check that all the VMRM setup steps were performed. In particular ensure:

- That the administrator ID of the z/VM manageability access point is authorized for the VMRM file pool and configuration file.
- If a custom VMRM setup is used:
  - That the VMSERVE service machine is authorized for the VMRM file pool and configuration file
  - That the custom configuration file is specified in CimVm.properties on the z/VM manageability access point.

Not enough space for new z/VM virtual server

This problem affects both the Virtual Server Deployment task and the Server Complexes task.

Problem

There appears to be sufficient disk space but creating a new z/VM virtual server fails because of insufficient space.

Investigation

Check your z/VM disk pool. There might be enough disk space overall but not sufficient contiguous space.

z/VM virtual server creation fails

This problem relates to the Virtual Server Deployment task.

Problem

An attempt to create a z/VM virtual server returns errors DNZZVS323E, DNZZSM365E, and DNZZMS402E with rc=400 and reason=8.

Investigation

In the “Create z/VM Virtual Server” wizard, you might have specified a password for the z/VM virtual server that does not conform to the password policies at your installation. Retry creating the z/VM virtual server with a password that adheres to the password policies.

z/VM virtual server has no network

This problem affects both the Virtual Server Deployment task.

Problem

The Linux instance on a newly created z/VM virtual server cannot be reached within the expected network or ping is not working.

Investigation

Check your network specification for the new Linux instance. A Linux instance that is created with the Virtual Server Deployment task must be in the same subnet as the master Linux system on which the new Linux instance is based.

Check your master Linux system. A correct network specification must be included in the master system directory entry. Ensure that:

- The device numbers for any OSA devices are included.
- Where applicable, NICDEF definitions are included.
- If the connection uses a VSWITCH, that access to the switch has been granted to the new user.

Multiple cloning fails

This problem affects the Server Complexes task.

Problem

The second cloning step of a multiple cloning operation fails because of an existing user ID.

Investigation

Check that the used virtual server template used for the cloning allows for generating multiple user IDs. Display the template in the Virtual Server

Deployment task. The **User ID Pattern** must contain a pattern that ends with an asterisk (*). The asterisk acts as a wildcard. Ensure that the user IDs according to the pattern are not all used up.

The more digits are covered by the wildcard character, the greater is the number of unique user IDs that can be provided. If the specified user ID pattern, including wildcard, is less than eight characters, the generated names are padded with zeros.

Examples:

LINUXVS*
would provide for the 10 user IDs from LINUXVS0 through LINUXVS9

LINUX*
would provide for the 1000 user IDs from LINUX000 through LINUX999

Change the user ID pattern if necessary.

z/VM virtual server is not removed after failed cloning

This problem affects the Server Complexes task.

Problem

An error occurs while a Linux guest system is being cloned with the Server Complexes task. The z/VM virtual server that is created in the failed cloning process is not cleaned up.

Investigation

Ensure that communications with the z/VM manageability access point are not interrupted. Check that the cimserver process is still running on the z/VM manageability access point. If the cimserver process is not running, restart it by issuing:

```
/etc/init.d/dacimom start
```

Check if the Linux instance has been installed successfully on the z/VM virtual server, despite the error message. A successfully installed Linux instance can be discovered by IBM Director. In the Server Complexes window, the corresponding Linux guest system appears in the Free Linux Guests pane instead of the target server complex. You can drag the Linux guest system into the target server complex to complete the configuration.

If the Linux instance has not been installed successfully, you can delete the z/VM virtual server with the Virtual Server Deployment task.

z/VM virtual server cannot be activated

This problem affects both the Virtual Server Deployment task and the Server Complexes task.

Problem

A newly created z/VM virtual server cannot be activated.

Investigation

If you are using a security manager (for example, RACF), ensure that you have defined the new z/VM virtual server to the security manager.

Check /var/log/personalization.log to find out if the personalization in the Linux guest operating system was performed correctly.

Minidisk attachment fails

This problem affects both the Virtual Server Deployment task and the Server Complexes task.

Problem

A minidisk cannot be attached to a z/VM virtual server.

Investigation

The error messages that report the failed attempt to attach the disk usually tells the reason. For example the VDEV might be used or the minidisk might not exist. If you are using a security manager beyond the security built into z/VM (for example, RACF), you might have to grant access to the minidisk with the security manager.

Minidisk detached after reboot

This problem affects the Server Complexes task.

Problem

A minidisks was attached to a Linux guest system by applying server complex minidisk properties, but is no longer attached after Linux is rebooted.

Investigation

Ensure that the zVMPersonalization.s390.rpm is installed correctly.

- Issue `rpm -q zVMPersonalization`
- Check that /etc/init.d/personalize exists and that 'chkconfig --list personalize' is 'on' for runlevel 2,3,4,5)

Check that /etc/sc_fstab includes a line of this format for the minidisk:

```
devno partNum mountPoint fsType mountOptions
```

where *devno* is the virtual device number of the minidisk and the other variables represent information on the disk and its mount parameters.

Network interface missing or mapped to the wrong port

This problem affects both the Virtual Server Deployment task and the Server Complexes task.

Problem

A newly created operating system instance does not have all network interfaces or some of the network interfaces are not mapped to the expected ports.

Investigation

This problem can be caused by incomplete network interface data in the operating system template from which the operating system instance has been created.

The network interface information of the master operating system on which the operating system template was based might not have been registered correctly or might not be up-to-date. To ensure that newly created operating system instances assign all required network interfaces to the correct port, the operating system template needs to have a complete list of all required network interfaces and ports.

In the Virtual Server Deployment task window, select the operating system template from which the new operating system instance has been created. In the right pane, select the Network Ports tab. If the information is incomplete, select the Relationships tab to find out from which master operating system the operating system template had been created. Re-register that master operating system and ensure that the network interface information is complete and correct.

You can now create a new operating system template and from this new template a new operating system instance.

New Linux instance has incorrect IP address or host name

This problem relates to Linux instances that are created with the Virtual Server Deployment "Create Operating System" wizard.

Problem

A newly created Linux instances does not have the host name and IP addresses you specified in the "Create Operating System" wizard.

Investigation

Examine /var/log/personalization.log for possible reasons. If this file does not exist, the Linux instance on which the new Linux instance was based might not have been prepared as a master Linux instance.

In the "Create Operating System" wizard you specified an operating system template. To find out the source of the template, select the operating system template in the Provisioning Resources tree and then click the Relationship table. The source of the template must have been set up as described in "Preparing a master Linux system" on page 428.

OSA information is not displayed

This problem affects both the Virtual Server Deployment task and the Server Complexes task.

Problem

The information on OSA devices is missing in the network information shown in z/VM Center interfaces.

Investigation

Ensure that the administrator ID of the z/VM manageability access point has the required privileges to query OSA information.See the information on how to set up z/VM for z/VM Center in the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html

Contacting customer support

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your xSeries or IntelliStation system, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system. Information about diagnostic tools is in the *Hardware Maintenance Manual and Troubleshooting Guide* on the IBM *xSeries Documentation* CD or in the IntelliStation *Hardware Maintenance Manual* at the IBM Support Web site.
- Go to the IBM Support Web site at www.ibm.com/pc/support/ to check for technical information, hints, tips, and new device drivers or to submit a request for information.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the online help or in the publications that are provided with your system and software. The information that comes with your system also describes the diagnostic tests that you can perform. Most xSeries and IntelliStation systems, operating systems, and programs come with information that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the information for the operating system or program.

Using the documentation

Information about your IBM xSeries or IntelliStation system and preinstalled software, if any, is available in the documentation that is included with your system. That documentation includes printed books, online books, readme files, and help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to www.ibm.com/pc/support/ and follow the instructions. Also, you can order publications through the IBM Publications Ordering System at www.elink.ibmlink.ibm.com/public/applications/publications/cgibin/pbi.cgi.

Getting help and information from the World Wide Web

On the World Wide Web, the IBM Web site has up-to-date information about IBM xSeries and IntelliStation products, services, and support. The address for IBM xSeries information is http://www.ibm.com/eserver/xseries/. The address for IBM IntelliStation information is http://www.ibm.com/pc/intellistation/.

You can find service information for your IBM products, including supported options, at http://www.ibm.com/pc/support/.

Software service and support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with xSeries servers, IntelliStation workstations, and appliances. For information about which products are supported by Support Line in your country or region, go to http://www.ibm.com/services/sl/products/.

For more information about Support Line and other IBM services, go to http://www.ibm.com/services/, or go to http://www.ibm.com/planetwide/ for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

Appendix A. IBM Director commands

This topic provides reference information about the IBM Director commands.

Syntax diagram conventions

This topic describes the syntax-diagram conventions that are used in each command description.

To read syntax diagrams, follow the path of the line. Read from left to right and from top to bottom.

Symbol	Meaning
►►—	Indicates the beginning of the syntax diagram.
—►	Indicates that the syntax diagram continues on the next line.
►—	Indicates that the syntax diagram continues from the previous line.
—►◄	Indicates the end of the syntax diagram.

Required items

When a keyword, variable, or operand appears on the main line, you must specify that item. In the following example, you must choose A, B, and C.

```
►►──A──B──C────────────────────────────────────────────►◄
```

When two or more items are in a stack and one of them is on the main line, you must specify one item. In the following example, you must choose A, B, or C.

```
►►──┬─A─┬──────────────────────────────────────────────►◄
    ├─B─┤
    └─C─┘
```

Optional items

When one or more items appear above or below the main line, the items are optional. In the following example, you can choose A, B, C, or nothing at all.

```
►►──┬───┬──────────────────────────────────────────────►◄
    ├─A─┤
    ├─B─┤
    └─C─┘
```

When an optional item appears above the main line, the item above the line is the default value when no optional item is specified in the command. In the following example, the user has the same choices as above (A, B, C, or nothing at all), but if nothing is selected, the default value will be A.

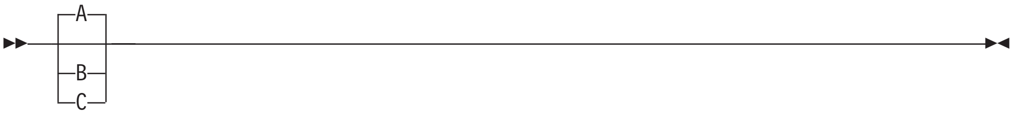

Repeatable items

An arrow returning to the left means you can repeat the item, for example:

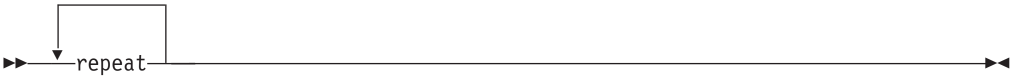

If one or more characters appear in the arrow's line, those characters are required to separate repeated items.

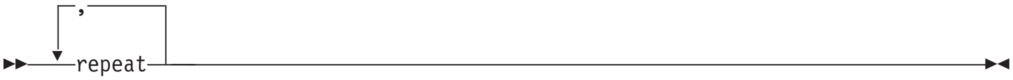

If you can choose from two or more items, they are displayed vertically in a stack. A stack of items followed by an arrow returning to the left means that you can select more than one item or, in some cases, repeat a single item. In the following example, you can choose any combination of A, B, or C.

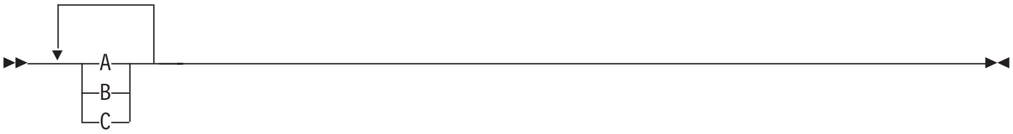

Syntax fragments

Sometimes, a portion of the syntax is used multiple times in the syntax diagram or would make the syntax diagram difficult to read if presented in its entirety at the position where it is used. A syntax fragment may be used to make the syntax diagram easier to read. When a syntax fragment is used, the syntax diagram changes in two ways:

- In the main diagram, the syntax fragment is represented by a name appearing between vertical bars. The fragment name appears at each point where the syntax is used.
- After the main diagram, the name and complete syntax of the fragment is displayed, with vertical bars appearing as start and stop terminators for the syntax.

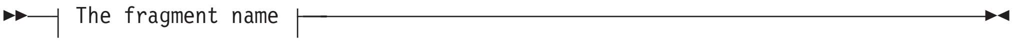

The fragment name:

```
├──┬─A─┬──────────────────────────────────────────────────────────┤
   ├─B─┤
   └─C─┘
```

Variables

Italicized, lowercase elements denote variables. In the following example, you must specify a variable value when you enter the keyword command:

```
►►──keyword──variable─────────────────────────────────────────────►◄
```

IBM Director CLI (dircli)

This topic lists the commands available in the IBM Director command-line interface, **dircli**.

CLI

Command	Description
“lsbundle” on page 549	Lists the **dircli** bundles and commands.

Configuration Manager

Command	Description
“lscmcfg” on page 552	Retrieves configuration information from BladeCenter configuration profiles or BladeCenter chassis.
“mkcmprof” on page 575	Creates BladeCenter configuration profiles.
“rmcmprof” on page 585	Deletes BladeCenter configuration profiles.

Groups

Command	Description
“chgp” on page 541	Modifies groups.
“lsgp” on page 556	Retrieves information about groups.
“mkgp” on page 577	Creates groups.
“rmgp” on page 587	Deletes groups.

Managed objects

Command	Description
“accessmo” on page 532	Requests access to secured managed objects.
“chmo” on page 543	Changes writable attribute values for managed objects
“discover” on page 546	Discovers managed objects.

Command	Description
“lsmo” on page 561	Retrieves managed-object attribute information.
“mkmo” on page 580	Creates managed objects.
“pingmo” on page 583	Pings managed objects.
“rmmo” on page 589	Deletes managed objects.

Power management

Command	Description
“rpower” on page 592	Performs power-management operations.

Security

Command	Description
“certmgr” on page 537	Generates, imports, distributes, and revokes security certificates for Level-1 managed systems.

Tasks

Command	Description
“lstask” on page 569	Lists IBM Director tasks.
“runtask” on page 595	Runs non-interactive IBM Director tasks.

Troubleshooting

Command	Description
“lsmethods” on page 560	Used to help identify problems. Normally not used except when instructed to by IBM support personnel.

Update Assistant

Command	Description
“impuapkg” on page 548	Imports IBM Director Update Assistant packages.

accessmo

This topic provides information about the **accessmo** command. This command is used to request access to secured managed objects.

```
►►──dircli──accessmo──┬──────────────┬──┬────┬──┬─────────┬──┬────────────┬──────────►
                      └─-L──language─┘  └─-v─┘  └─-t──type─┘  └─-u──user_id─┘
```

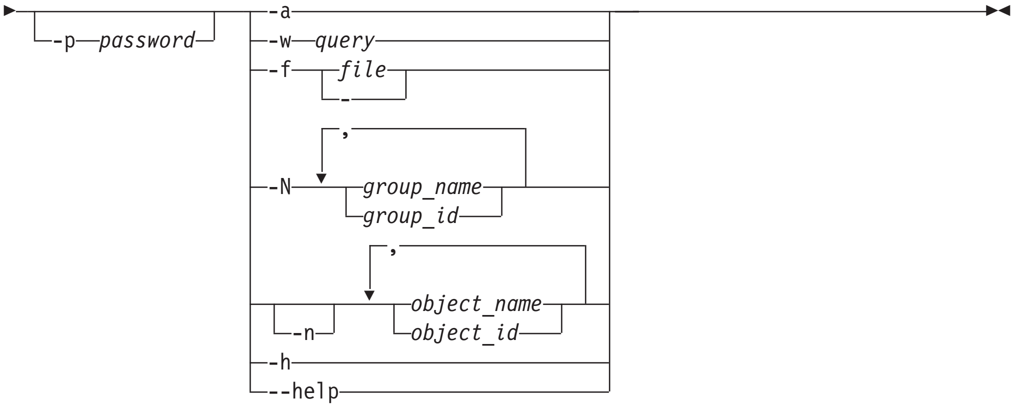

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 17. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose

Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-t | --type *type*

In conjunction with a primary targeting option (explicit or default), specifies the type of managed object targeted. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

Notes: If you specify a parent managed-object type (for example, ″Chassis″), its children (in this case, ″BladeCenter Chassis″) are also targeted.

Use the following command to list all managed-object types for your installation:

```
dircli lsmo -i
```

-u | --user *user_id*

Specifies a valid user ID with administrative privileges on the managed object. If this option is not issued, IBM Director will prompt the user for the user ID.

Notes: The value for *user_id* will be used for all managed objects targeted by the command.

Use of the **-u** and **-p** options presents a potential security risk, since the user name and password might be recorded in the shell or other operating-system areas.

-p | --password *password*

Specifies the password associated with a valid user ID with administrative privileges on the managed object. If this option is not provided, IBM Director will prompt the user for the password.

Notes: The value for *password* will be used for all managed objects targeted by the command.

Use of the **-u** and **-p** options presents a potential security risk, since the user name and password might be recorded in the shell or other operating-system areas.

-a | --all

Targets all managed objects.

-w | --where *query*

Targets managed objects based on managed-object attribute values specified in *query*. The *query* operand is a quote-delimited string defining a simple SELECT query that uses the following format:

```
key_1=value_1 [AND|OR] key_2=value_2 ... [AND|OR] key_n=value_n
```

Enclose the statement in quotation marks. If any managed-object attributes are strings that include spaces, enclose the strings in single quotation marks. Use parentheses to nest logical constructions.

-f | --file {*file* | -}

Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]
Targets the managed objects that are in the specified groups. If a managed object belongs to more than one group, the managed object is targeted only once.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]
Targets the managed objects specified by name or ID.

object_id
Unique ID of the managed object.

object_name
Name of the managed object.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

Request access to all Level-1 managed objects
The following example illustrates use of the **-t** and **-a** options. The **-a** option instructs **accessmo** to request access to all managed objects; the **-t** option limits that selection based on the specified managed-object type.

```
dircli accessmo -t "Level 1: Core Services" -a
```

Because a user ID and password are not specified in the command, the user is prompted to enter them. The entered user ID and password are used in the access request to each Level-1 managed object. The user ID and password information is not, however, recorded in the shell or other operating-system areas.

Request access to all Level-1 managed objects using the -w option
The following example illustrates use of the **-w** option to specify a SELECT statement identifying managed objects by an attribute value. The result of the following example is the same as for the previous example.

```
dircli accessmo -w "MO.MOFID='Level 1: Core Services'"
```

Note that the value for the key is enclosed in single quotes because it contains three space characters. Proper quoting is essential for accurate processing of query statements using the **-w** option.

Request access by managed-object name, specifying user ID and password
The following command requests access to a managed object named webserver, using user ID admin95 and password 987sixfive.

```
dircli accessmo -u admin95 -p 987sixfive -n webserver
```

Note: Use of the **-u** and **-p** options presents a potential security risk, since the user name and password might be recorded in the shell or other operating-system areas.

Request access to managed objects listed in a file

The following command requests access to the managed objects listed in the file `building7`.

```
dircli accessmo -f building7
```

The file must contain the names or IDs of managed objects, separated by commas or line breaks.

Request access to managed objects listed by another command

The following command line uses the **lsmo** command to list the managed objects in the group ″Level 2: IBM Director Agents″, then uses that list as input for the **accessmo** command.

```
dircli lsmo --groups "Level 2: IBM Director Agents" | dircli accessmo -f -
```

Request access to managed objects in a managed-object group

The following command requests access to the same managed objects as the previous example. Instead of using the **lsmo** command to list the managed objects in the group, the group name is specified directly in the **accessmo** command.

```
dircli accessmo -N "Level 2: IBM Director Agents"
```

Return codes

The following table contains the codes returned by the **accessmo** command.

Code	Meaning
0	The managed object was successfully accessed.
1	A usage error occurred.
10	The file was not found.
20	A managed object was not found.
21	A group was not found.
25	A number-formatting error occurred.
26	A invalid attribute type was specified
27	An invalid attribute was specified.
32	An invalid managed-object ID was specified.
51	A managed object was not locked.
60	a request for access is not supported. The managed object cannot be locked or unlocked.
61	A managed object is not available.
62	A managed object can not be accessed.
63	The action was not permitted by the target system.
64	The request-access operation failed.
65	An encrypted managed object was not accessed.
66	The access request is not finished.

certmgr

This topic provides information about the **certmgr** command. This command generates, imports, distributes, and revokes security certificates for Level-1 managed systems.

Note: Although any managed object may be specified as a target for the **certmgr** command, only managed objects with IBM Director Core Services (Level-1 managed systems) are actually targeted by the command. All other managed objects are ignored.

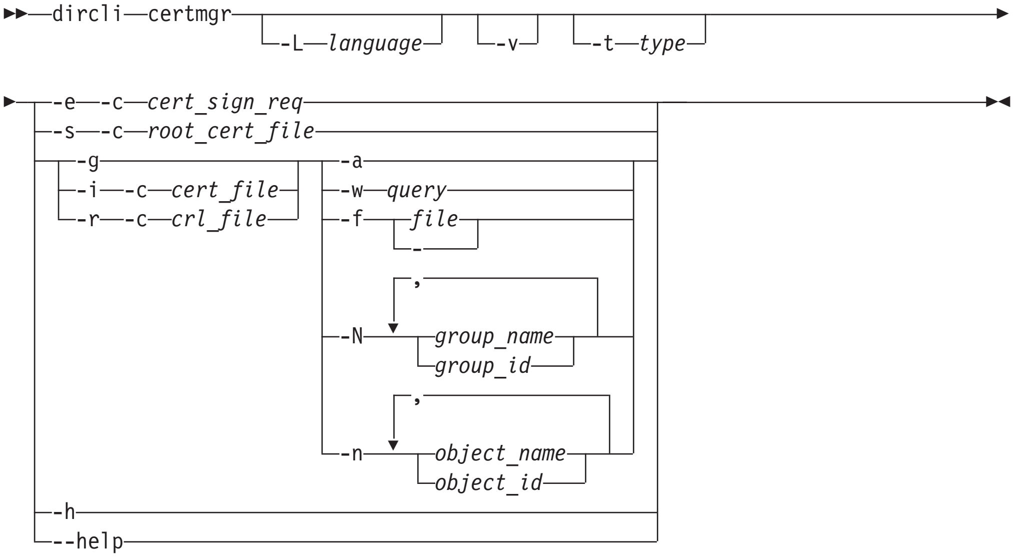

Options and operands

-L | --lang *language*
Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 18. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose
Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-t | --type *type*

In conjunction with a primary targeting option (explicit or default), specifies the type of managed object targeted. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

Notes: If you specify a parent managed-object type (for example, ″Chassis″), its children (in this case, ″BladeCenter Chassis″) are also targeted.

Use the following command to list all managed-object types for your installation:

```
dircli lsmo -i
```

-e *cert_sign_req*

Generates a certificate-signing request file using the absolute path and filename specified by *cert_sign_req*. The generated certificate-signing request file is a standard base64-encoded file which can be sent to a third-party signer of security certificates.

-s *root_cert_file*

Imports the base64-encoded root certificate file using the absolute path and filename specified by *root_cert_file* to establish a trust relationship with the signer. This action should typically only be required once for each signer of security certificates.

-g Generates a new base64-encoded security certificate and distributes it to the targeted Level-1 managed systems. IBM Director-generated security certificates are valid for 365 days. System administrators are notified when certificates are nearing expiration.

-i *cert_file*

Imports a signed base64-encoded security certificate using the absolute path and filename specified by *cert_file* and distributes it to the targeted Level-1 managed systems.

-r *crl_file*

Imports the certificate-revocation list (CRL) file using the absolute path and filename specified by *crl_file* and revokes the security certificates listed in the file on the targeted Level-1 managed systems.

-a | --all

Targets all managed objects with IBM Director Core Services installed (Level-1 managed systems).

-w | --where *query*

Targets managed objects based on managed-object attribute values specified in *query*. The *query* operand is a quote-delimited string defining a simple SELECT query that uses the following format:

```
key_1=value_1 [AND|OR] key_2=value_2 ... [AND|OR] key_n=value_n
```

Enclose the statement in quotation marks. If any managed-object attributes are strings that include spaces, enclose the strings in single quotation marks. Use parentheses to nest logical constructions.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]

Targets the managed objects that are in the specified groups. If a managed object belongs to more than one group, the managed object is targeted only once.

group_id

Unique ID of the IBM Director managed-object group.

group_name

Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-f | --file {*file* | -}

Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]

Targets the managed objects specified by name or ID.

object_id

Unique ID of the managed object.

object_name

Name of the managed object.

-h | -?

Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help

Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

The following example illustrates the command usage for networks using IBM Director-generated security certificates.

Generate a security certificate and distribute it to all Level-1 managed systems

The following command distributes a new IBM Director-generated security certificate to all Level-1 managed systems:

```
dircli certmgr -g -a
```

The following examples illustrate the command usage for networks using security certificates from third parties such as Verisign. The examples are listed in the sequence in which they are typically performed.

Create a certificate-signing request file

The following command creates the certificate-signing request file /opt/ibm/director/ibmd.csr:

```
dircli certmgr -e -c /opt/ibm/director/ibmd.csr
```

The system administrator sends the created certificate-signing request file to the organization which will provide security certificates for Level-1 managed systems.

Import a root certificate

The following command imports the root certificate /opt/ibm/director/rootca.cert provided by a third-party issuer of security certificates:

```
dircli certmgr -s -c /opt/ibm/director/rootca.cert
```

The root certificate should only have to be imported once for each organization issuing security certificates.

Import a signed security certificate

The following command imports the signed security certificate /opt/ibm/director/ibmd.cert provided by a third-party issuer of security certificates and distributes the certificate to the Level-1 managed systems belonging to the group WestCampus:

```
dircli certmgr -i -c /opt/ibm/director/ibmd.cert -N WestCampus
```

Third-party security certificates may expire at different intervals. When the certificate is nearing expiration, the system administrator will be notified.

Revoke security certificates identified in a certificate-revocation list file

The following command is used infrequently, but revokes the security certificates listed in the certificate-revocation list file /opt/ibm/director/ibmd.crl:

```
dircli certmgr -i -c /opt/ibm/director/ibmd.crl -a
```

Certificate-revocation list files are not generated by IBM Director, but are used by third-party issuers of security certificates.

Return codes

The following table lists the codes returned by the **certmgr** command.

Return code	Meaning
0	The operation was successful.
1	A usage error occurred.
10	The file was not found.
20	A managed object was not found.
21	A group was not found.
26	An invalid object type was specified.
27	An invalid attribute was specified.

chgp

This topic provides information about the **chgp** command. Use this command to modify groups. The groups affected may be groups of managed objects, groups of tasks, or groups containing other groups.

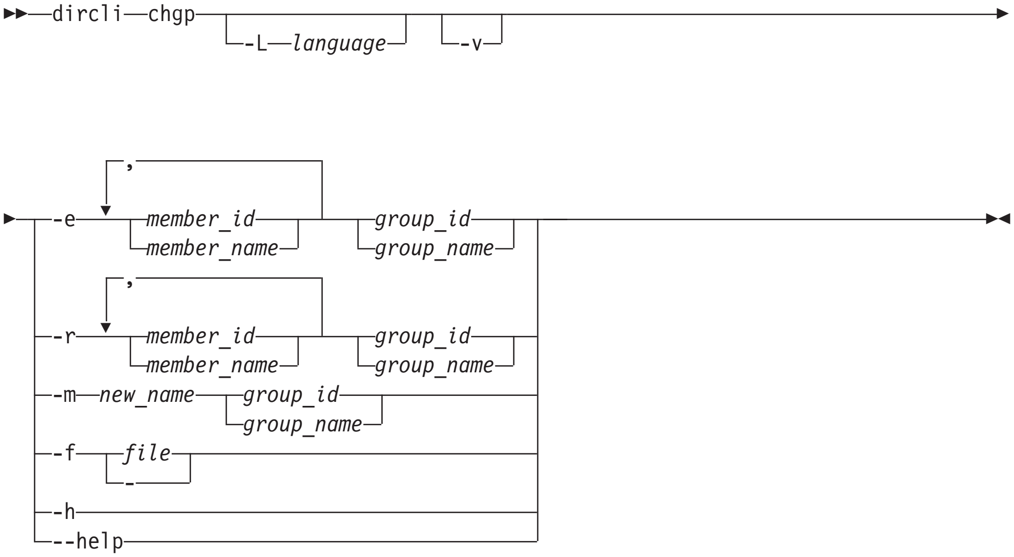

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 19. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose

Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-e | --extend {*member_id* | *member_name*} [, {*member_id* | *member_name*} ...] {*group_id* | *group_name*}
Adds objects to a group.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

member_id
Unique ID of a group, managed object, or task which is a member of the specified parent group.

member_name
Name of a group, managed object, or task which is a member of the specified parent group. If the member name contains space characters, it should be quoted.

-r | --remove {*member_id* | *member_name*} [, {*member_id* | *member_name*} ...] {*group_id* | *group_name*}
Removes the specified members from a group.

-m | --move *new_name* {*group_id* | *group_name*}
Renames the group specified by *group_id* or *group_name* to a new name specified by *new_name*.

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Each line of the input data represents a **chgp** instruction and uses one of the following three formats:

```
{group_name|group_id}:extend:{member_id|member_name}[,{member_id|member_name}...]
{group_name|group_id}:remove:{member_id|member_name}[,{member_id|member_name}...]
{group_name|group_id}:move:new_name
```

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

Add members to a task group
The following command adds the tasks ″Remote Control″ and ″Remote Session″ to the task group named RemoteGroup:

```
dircli chgp -e "Remote Control","Remote Session" RemoteGroup
```

Remove an object from a group
The following command removes a single managed object (WebServer1) from the static group named WebSystems:

```
dircli chgp -r WebServer1 WebSystems
```

Rename a group
The following command renames RemoteGroup to RemoteTaskGroup:

```
dircli chgp -m RemoteTaskGroup RemoteGroup
```

Use an input file to modify several groups

The following command modifies four groups, using the data contained in the c:\temp\modify_groups.txt file:

```
dircli chgp -f c:\temp\modify_groups.txt
```

The modify_groups.txt file contains the following text:

```
group1:extend:mo_id1,mo_id2
group2:remove:task_id1,task_id2
0x300A:extend:group_id1,group_id2
OldName:move:NewName
```

Return codes

The following table contains the codes returned by the **chgp** command.

Code	Meaning
0	The group was modified successfully.
1	A usage error occurred.
10	The file was not found.
20	A managed object was not found.
21	A group was not found.
22	A task was not found.
25	A number-formatting error occurred.
27	An invalid attribute was specified.
50	A group is read-only.

chmo

This topic provides information about the **chmo** command. This command changes writable attribute values for managed objects.

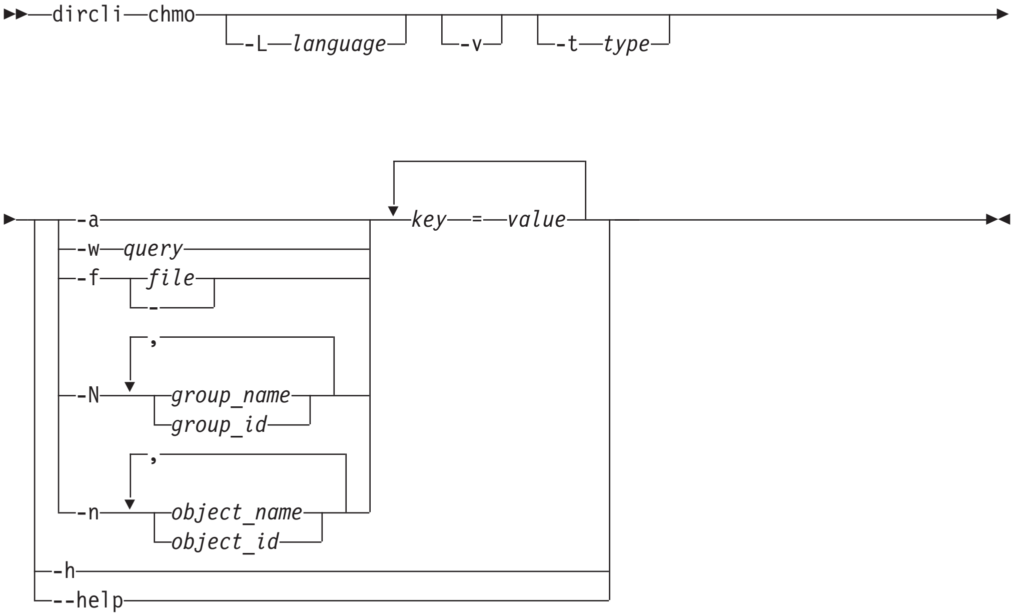

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 20. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose

Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-t | --type *type*

In conjunction with a primary targeting option (explicit or default), specifies the type of managed object targeted. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

Notes: If you specify a parent managed-object type (for example, ″Chassis″), its children (in this case, ″BladeCenter Chassis″) are also targeted.

Use the following command to list all managed-object types for your installation:

```
dircli lsmo -i
```

-a | --all

Targets all managed objects.

-w | --where *query*

Targets managed objects based on managed-object attribute values specified in *query*. The *query* operand is a quote-delimited string defining a simple SELECT query that uses the following format:

```
key_1=value_1 [AND|OR] key_2=value_2 ... [AND|OR] key_n=value_n
```

Enclose the statement in quotation marks. If any managed-object attributes are strings that include spaces, enclose the strings in single quotation marks. Use parentheses to nest logical constructions.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]

Targets the managed objects that are in the specified groups. If a managed object belongs to more than one group, the managed object is targeted only once.

group_id

Unique ID of the IBM Director managed-object group.

group_name

Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-f | --file {*file* | -}

Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]

Targets the managed objects specified by name or ID.

object_id

Unique ID of the managed object.

object_name

Name of the managed object.

*key***=***value*

Sets the indicated managed-object *key* to the specified *value*. Only writable managed-object attributes can be changed. If *value* is a string that contains spaces, enclose the string in quotation marks.

Use the **lsmo** command to list managed-object definitions, including type statements and *key=value* pairs.

-h | -?

Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help

Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

Set the ping interval for managed objects

The following command targets all managed objects with a ping interval of 15 minutes and changes the ping interval to 30 minutes:

```
dircli chmo -w MO.ping=15 MO.ping=30
```

Change the name

The following command targets a single managed system (webserver) and changes its name to SMTPserver:

```
dircli chmo -n webserver MO.name=SMTPserver
```

Set the ping interval for managed objects listed in a file

The following command targets all managed objects listed in the /tmp/modef file and changes the ping interval to 120 minutes:

```
dircli chmo -f /tmp/modef MO.ping=120
```

Return codes

The following table lists the codes returned by the **chmo** command.

Return code	Meaning
0	The attribute was successfully changed.
1	A usage error occurred.
10	The file was not found.
20	A managed object was not found.
21	A group was not found.
25	A number-formatting error occurred.
26	An invalid object type was specified.
27	An invalid attribute was specified.
60	The attribute is not supported.
61	The attribute is read-only.
62	No attribute was specified.
63	The managed object cannot be renamed.

discover

This topic provides information about the **discover** command. This command discovers managed objects.

```
►►──dircli──discover──┬───────────────┬──┬────┬──┬───────────┬──┬─────────┬────────►◄
                      └─-L──language──┘  └─-v─┘  └─-t──type──┘  ├─-h──────┤
                                                                └─--help──┘
```

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 21. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose
Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-t |--type *type*
Specifies the type of managed object to be discovered. If this option is not specified, IBM Director discovers all systems and devices by default. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

Discover all managed objects
The following command discovers all managed objects:

```
dircli discover
```

Discover all BladeCenter chassis
The following command discovers all BladeCenter chassis:

```
dircli discover -t "BladeCenter Chassis"
```

Return codes

The following table contains the codes returned by the **rmmo** command.

Code	Meaning
0	The discovery operation started successfully.
1	A usage error occurred.
26	An invalid object type was specified.
50	The specified object type does not support discovery.

impuapkg

This topic provides information about the **impuapkg** command. This command imports IBM Director Update Assistant packages. An IBM Director Update Assistant package is either an xSeries Update package or a Solution Install (SI) package.

```
►►──dircli──impuapkg──┬─────────────┬──┬────┬──────────────────────────►
                      └─-L──language─┘  └─-v─┘

►──┬─-x──xml_file──┬──────────────────┬─┬───────────────────────────►◄
   │               └─-r──response_file─┘ │
   ├─-h──────────────────────────────────┤
   └─--help──────────────────────────────┘
```

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 22. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose

Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-x | --XML *filename*
Specifies the update package to import. The *filename* operand may specify the fully-qualified or relative path and filename of one of the following files:

- the xSeries Update package XML file
- the Deployment Descriptor of a Solution Install (SI) package
- the jar/zip file that comprises an SI package

-r | --response *response_file*
Specifies the fully-qualified or relative path and filename of the alternate response file that will be used when importing a package.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Example

Import a package in the current directory named update85.xml
The following command imports the package named update85.xml:

```
dircli impuapkg -x update85.xml
```

Return Codes

The following table contains the codes returned by the **impuapkg** command.

Code	Meaning
0	The operation started successfully.
1	A usage error occurred.
50	Invalid package.
51	Package failed security check(s).
52	Invalid input file.
53	Failure(s) importing package.

lsbundle

This topic provides information about the **lsbundle** command. This command lists the **dircli** bundles and commands.

```
►►──dircli──lsbundle──┬─────────────┬──┬────────┬──────────────────►◄
                      └─-L──language─┘  ├─-h─────┤
                                        └─--help─┘
```

Options and operands

-L | --lang *language*
Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore

character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 23. Common values for the language code

Language/country	*language*
Chinese/China	zh_CN
English/United States	en_US
French/France	fr_FR
German/Germany	de_DE
Japanese/Japan	ja_JP
Korean/South Korea	ko_KR
Spanish/Mexico	es_MX

-h | -?

Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help

Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

List the dircli commands

The following command lists all the bundles and commands defined for **dircli**:

```
dircli lsbundle
```

This command returns the following output (actual commands listed may vary, depending on the command bundles installed):

```
BladeCenterChassis/addbcchassis
BladeCenterChassis/discoverbcchassis
BladeCenterChassis/help
BladeCenterChassis/list
BladeCenterChassis/listbcchassis
Chassis/chassislist
Chassis/chassissubsystemlist
Chassis/chassissubsystemtypelist
Chassis/help
Chassis/list
cli/help
cli/listbundle
event/applyeventactionplan
event/createeventactionplan
event/help
event/list
event/listeventactionplans
event/listeventactions
event/listevents
event/listeventtypes
event/listfilters
monitor/applythreshold
monitor/help
monitor/list
monitor/listthresholds
mpa/help
mpa/list
mpa/listobjectattributes
mpa/listobjectattributevalues
```

```
mpa/listobjectsbyattribute
mpa/mpcli
mpa/setcredentials
mpa/setsysinterconnectconnection
native/addsystem
native/help
native/list
native/listsystems
native/startdiscovery
procmon/applypmtask
procmon/createpmtask
procmon/help
procmon/list
procmon/listpmtasks
scheduler/canceljobactivation
scheduler/cancelrdmtaskactivation
scheduler/getjobactivationlog
scheduler/getjobstatus
scheduler/help
scheduler/list
scheduler/listjobactivations
scheduler/listjobactivationsbysystem
scheduler/listjobs
server/accessobjects
server/addtostaticgroup
server/createdynamicgroup
server/createstaticgroup
server/deletegroups
server/deleteobjects
server/discoverall
server/doconstraintdump
server/help
server/list
server/listdynamicgroupcriteria
server/listgroupattributes
server/listgroupmembers
server/listgroups
server/listgroupsbyattribute
server/listinventoryvalues
server/listnoninteractivetasks
server/listobjectattributes
server/listobjects
server/listobjectsbyattribute
server/listtaskactivationstatus
server/pingobjects
server/removefromstaticgroup
server/renameobject
server/runtask
snmp/addsystem
snmp/get
snmp/getbulk
snmp/getnext
snmp/help
snmp/inform
snmp/list
snmp/listsystems
snmp/set
snmp/startdiscovery
snmp/trap
snmp/walk
user/addgroupaccess
user/addtaskaccess
user/help
user/list
user/listgroups
user/listprivilegetokens
user/listtasks
```

```
user/listuserattributes
user/listusers
user/modifyuserattributes
user/removegroupaccess
user/removetaskaccess

certmgr/certmgr
cli/lsbundle
configmgr/lscmcfg
configmgr/mkcmprof
configmgr/rmcmprof
group/chgp
group/lsgp
group/mkgp
group/rmgp
hwcontrol/rpower
managedobject/accessmo
managedobject/chmo
managedobject/discover
managedobject/lsmo
managedobject/mkmo
managedobject/pingmo
managedobject/rmmo
remotelib/lsmethods
swd/impuapkg
task/lstask
task/runtask
```

Return codes

The following table contains the codes returned by the **lsbundle** command.

Code	Meaning
0	The bundles and commands were listed successfully.
1	A usage error occurred.

lscmcfg

This topic provides information about the **lscmcfg** command. This command retrieves configuration information from BladeCenter configuration profiles or BladeCenter chassis.

```
►►──dircli──lscmcfg──┬──────────────┬──┬────┬──┬────────────┬──────────────────►
                     └─-L──language─┘  └─-v─┘  └─-d──symbol─┘
```

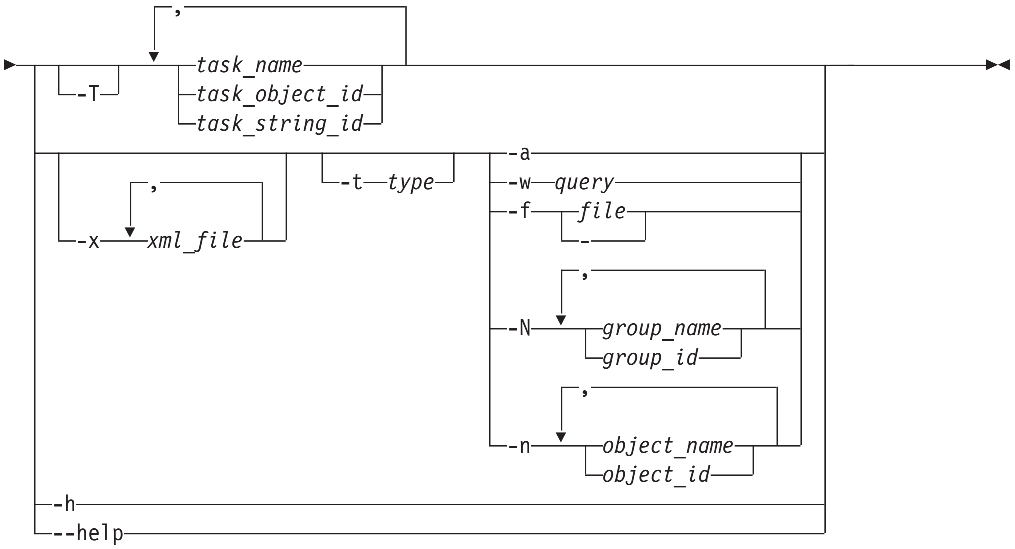

Options and operands

-L | --lang *language*
Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 24. Common values for the language code

Language/country	*language*
Chinese/China	zh_CN
English/United States	en_US
French/France	fr_FR
German/Germany	de_DE
Japanese/Japan	ja_JP
Korean/South Korea	ko_KR
Spanish/Mexico	es_MX

-v | --verbose
Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-d | --delimiter *symbol*
Specifies the character or set of characters that separates data.

-T | --tasks {*task_name* | *task_object_id* | *task_string_id*} [, {*task_name* | *task_object_id* | *task_string_id*} ...]
Specifies the tasks against which the command is targeted.

task_name
Name of the task.

task_object_id
Unique object ID of the task.

task_string_id
String ID of the task.

-x | --xmlfile *xml_file*[,*xml_file* ...]
Specifies the name of the XML file that contains the BladeCenter configuration information. *xml_file* is the fully qualified name of the XML file. If you list more than one XML file, separate the file names with commas.

-t | --type *type*
In conjunction with a primary targeting option (explicit or default), specifies the type of managed object targeted. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

Notes: If you specify a parent managed-object type (for example, "Chassis"), its children (in this case, "BladeCenter Chassis") are also targeted.

Use the following command to list all managed-object types for your installation:

```
dircli lsmo -i
```

-a | --all
Targets all managed objects.

-w | --where *query*
Targets managed objects based on managed-object attribute values specified in *query*. The *query* operand is a quote-delimited string defining a simple SELECT query that uses the following format:

```
key_1=value_1 [AND|OR] key_2=value_2 ... [AND|OR] key_n=value_n
```

Enclose the statement in quotation marks. If any managed-object attributes are strings that include spaces, enclose the strings in single quotation marks. Use parentheses to nest logical constructions.

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]
Targets the managed objects that are in the specified groups. If a managed object belongs to more than one group, the managed object is targeted only once.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]
Targets the managed objects specified by name or ID.

object_id
Unique ID of the managed object.

object_name
Name of the managed object.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

List configuration information by task object id
The following command lists configuration information for profiles with task object ids 0xA and 0xB:

```
dircli lscmcfg -T 0xA,0xB
```

List configuration information described in an XML file
The following command lists configuration information described in the file get.xml for managed objects rsa1 and rsa2:

```
dircli lscmcfg -x get.xml -n rsa1,rsa2
```

List configuration information using a query statement
The following command lists configuration information for the managed object with IP address 192.168.1.100:

```
dircli lscmcfg -w MO.ipaddr=192.168.1.100
```

Return codes

The following table lists the codes returned by the **lscmcfg** command.

Code	Meaning
0	The configuration information was listed successfully.
1	A usage error occurred.
10	The profile was not found.
20	The managed object was not found.
25	A number-formatting error occurred.
27	An invalid attribute was specified.
50	″Some failed.″
51	″All failed but with different reasons.″
53	An invalid task was specified.
54	″Unknown Configuration Manager plug-in.″

Code	Meaning
55	″Configuration Manager plug-in not registered.″
56	″Configuration Manager plug-in connection failed.″
57	″Configuration Manager plug-in lost connection.″

lsgp

This topic provides information about the **lsgp** command. This command retrieves information about groups.

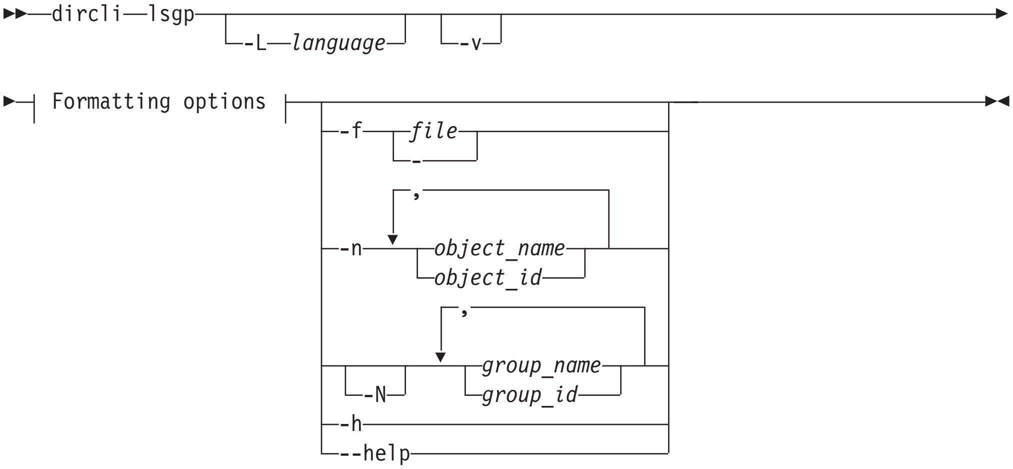

Formatting options:

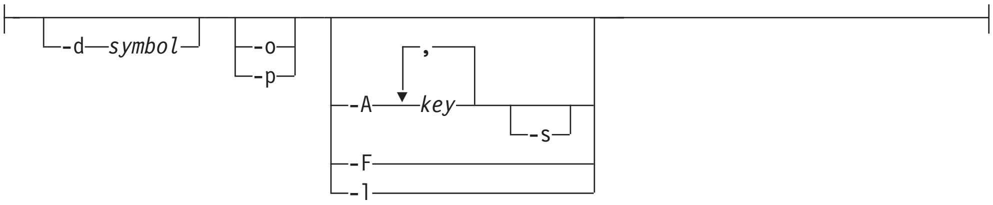

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 25. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`

Table 25. Common values for the language code (continued)

Language/country	*language*
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose
Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. The input data must specify the groups using either group names or group IDs. Items in the input file must be separated by commas or line breaks.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]
Targets the groups that contain the specified managed objects.

object_id
Unique ID of the managed object.

object_name
Name of the managed object.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]
Targets the groups specified by name or ID.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

-d | --delimiter *symbol*
Specifies the character or set of characters that separates data. The behavior of this option depends on the use of other options in the command:

Options used with -d |--delimiter	Result
--A | --attribute	Data fields in a record are separated by the specified delimiter *symbol*. Data records are separated by line breaks.
-l | --long or **-F | --format**	The **-d | --delimiter** option is ignored.
other options	Data fields in a record are separated by a comma followed by a space. Data records are separated by the specified delimiter *symbol*.

For examples of this option, refer to “Examples” on page 565.

-o | --oid
Specifies that the object ID (OID) is displayed together with the managed-object name. The **-o** option may also be combined with the **-l** and **-A** options. For an example of this option, refer to “Examples” on page 565.

-p | --pipe
Specifies that the object ID (OID) is displayed instead of the managed-object name. When used alone, this option enables the output to be piped to other **dircli** commands. The **-p** option may also be combined with the **-l** and **-A** options. For an example of this option, refer to “Examples” on page 565.

-A | --attribute *key* [, *key* ...] **-s | --sort**
Specifies the group-attribute values that are displayed. *key* is the attribute key. The following table lists the group-attribute keys and describes their values.

Key	Value
GP.name	The group name
GP.type	The group type
GP.readonly	True or false
GP.deletable	True or false
GP.members	All managed objects that belong to the group
GP.definition	Information varies depending on the value of GP.type: • (if GP.type = Static) All managed objects that belong to the group • (if GP.type = Task) The tasks defined for the group • (If GP.type = ManagedObject) The object type for the group • If GP.type = GroupCategory) The groups that belong to the group category Other group types are not supported.

If you issue the **-s | --sort** option, the output is sorted based on the value of the first group-attribute key specified.

-F | --format
Specifies that the output is displayed in the following format:

```
GP.name:GP.type:GP.definition
```

This output can be saved as a group-definition file, or it can be used as input for the **mkgp** command. The values for the group-attribute keys are the same as those listed for the **-A** *attribute_key* parameter, with the exception of GP.definition, which is the task string ID.

Note: IBM Director generates errors if you use this option in the following situations:

- With a command that targets dynamic or managed-object groups
- With static or group-category groups that contain objects with non-unique names

-l | --long

Specifies that the output is displayed in the following format:

```
<group_name>:
    GP.name = string
    GP.type = ManagedObject | Static | GroupCategory | Task
    GP.readonly = true | false
    GP.deletable = true | false
    GP.members = {'mo_name1', 'mo_name2'...}
    GP.definition = {'obj_name1','obj_name2',...}
```

Examples

List the names of all groups

The following command lists the names of all IBM Director groups:

```
dircli lsgp
```

For example, the output might read as follows:

```
All Groups
All Systems and Devices
Chassis and Chassis Members
Clusters and Cluster Members
Group Categories
Hardware Status Critical
Hardware Status Information
Hardware Status Warning
IBM Director Systems
Platforms and Platform Members
Racks with Members
Scalable Systems and Members
Systems with Windows XP
```

List the names and object IDs of all groups

The following command lists the names and object IDs of all IBM Director groups:

```
dircli lsgp -o
```

List the attributes for groups specified by name

The following command lists the attributes of the specified IBM Director groups:

```
dircli lsgp -l -N "Racks with Members",MyRackGroup
```

List the attributes for groups specified in a file

The following command lists the attributes of the IBM Director groups specified in /tmp/groups:

```
dircli lsgp -l -f /tmp/groups
```

Return codes

The following table contains the codes returned by the **lsgp** command.

Code	Meaning
0	The groups were listed successfully.
1	A usage error occurred.
10	The file was not found.
20	A managed object was not found.
21	A group was not found.
25	A number-formatting error occurred.

lsmethods

This topic provides information about the **lsmethods** command. This command lists available SDK/API methods. Normally, users should not use the **lsmethods** command except when instructed to by IBM support personnel.

```
►►──dircli──lsmethods──┬─All─────────────────┬──────────────────────────────►◄
                       ├─remote_sdk_api_name─┤
                       ├─-h──────────────────┤
                       └─--help──────────────┘
```

Options and operands

All
Specifies that **lsmethods** lists all methods for all APIs.

remote_sdk_api_name
Specifies that **lsmethods** lists all methods for the specified API.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

List methods for the Inventory API

The following command lists methods for the Inventory API:

```
dircli lsmethods Inventory
```

This command returns the following output:

```
Methods for uri: Inventory
java.util.Vector GetTopLevelResourceSets ()
java.util.Vector GetChildResourceSets (java.lang.String)
java.util.Vector GetAllQueries ()
java.util.Vector GetQueriesForResourceSet (java.lang.String)
java.util.Vector ExecuteQuery (java.util.Vector)
```

List methods for all APIs

The following command lists methods all APIs:

```
dircli lsmethods All
```

This command returns the following output:

```
Methods:
Methods for uri: Inventory
java.util.Vector GetTopLevelResourceSets ()
java.util.Vector GetChildResourceSets (java.lang.String)
java.util.Vector GetAllQueries ()
java.util.Vector GetQueriesForResourceSet (java.lang.String)
java.util.Vector ExecuteQuery (java.util.Vector)

Methods for uri: BladeCenterProvider

Methods for uri: AsynchNotification
java.lang.Boolean Subscribe (java.util.Vector)
java.util.Vector GetUpdates (java.util.Vector)
void GetUpdatesNow ()

Methods for uri: Task
java.lang.Object GetNonInteractiveTasksForGroup (com.ibm.sysmgt.sdk.remote.ManagedGroup)
java.lang.Object GetInteractiveTasksForGroup (com.ibm.sysmgt.sdk.remote.ManagedGroup)
java.lang.Object GetNonInteractiveTasksForSystem (com.ibm.sysmgt.sdk.remote.ManagedSystem)
java.lang.Object GetInteractiveTasksForSystem (com.ibm.sysmgt.sdk.remote.ManagedSystem)
java.lang.Object GetAllTaskActivationsByTask (com.ibm.sysmgt.sdk.remote.Task)
java.lang.Object GetTaskActivationsStatus (com.ibm.sysmgt.sdk.remote.TaskActivation)
java.lang.Object GetTaskActivationsStatusForMoid (com.ibm.sysmgt.sdk.remote.
TaskActivation,java.lang.Long)

...
```

Return codes

The following table contains the codes returned by the **lsmethods** command.

Code	Meaning
0	The methods were listed successfully.
1	A usage error occurred or the API was not found.

lsmo

This topic provides information about the **lsmo** command. This command retrieves managed-object attribute information.

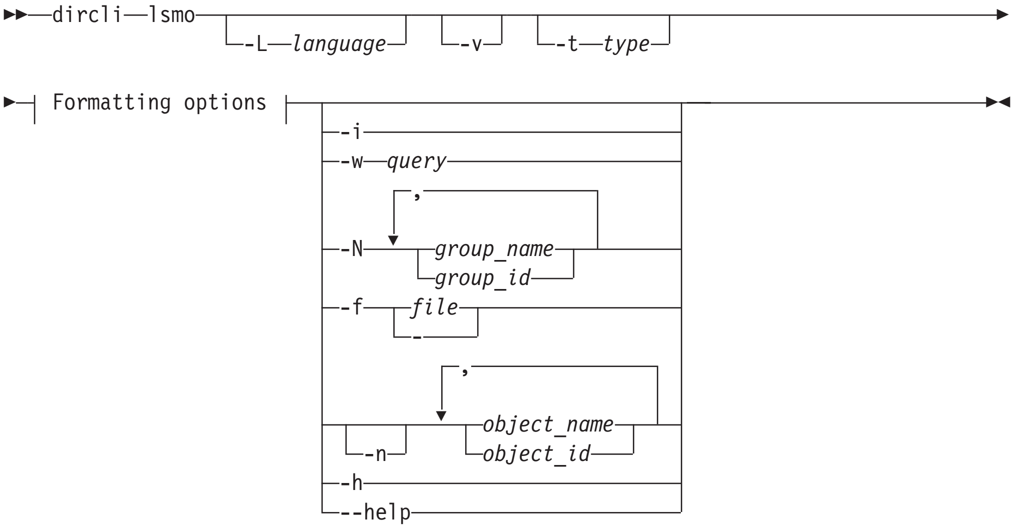

Formatting options:

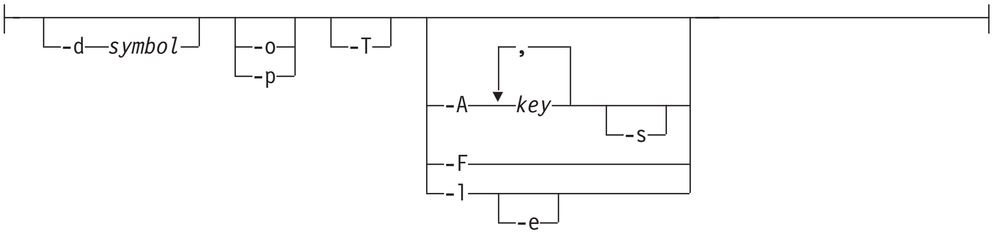

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 26. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose

Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-t | --type *type*

In conjunction with a primary targeting option (explicit or default), specifies the type of managed object targeted. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

Notes: If you specify a parent managed-object type (for example, "Chassis"), its children (in this case, "BladeCenter Chassis") are also targeted.

Use the following command to list all managed-object types for your installation:

```
dircli lsmo -i
```

-i | --listtype
Lists the names of all types of managed objects supported by IBM Director. Any additional options are ignored.

-w | --where *query*
Targets managed objects based on managed-object attribute values specified in *query*. The *query* operand is a quote-delimited string defining a simple SELECT query that uses the following format:

```
key_1=value_1 [AND|OR] key_2=value_2 ... [AND|OR] key_n=value_n
```

Enclose the statement in quotation marks. If any managed-object attributes are strings that include spaces, enclose the strings in single quotation marks. Use parentheses to nest logical constructions.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]
Targets the groups specified by name or ID.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]
Targets the managed objects specified by name or ID.

object_id
Unique ID of the managed object.

object_name
Name of the managed object.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

-d | --delimiter *symbol*
Specifies the character or set of characters that separates data. The behavior of this option depends on the use of other options in the command:

Options used with -d |--delimiter	**Result**
--A | --attribute	Data fields in a record are separated by the specified delimiter *symbol*. Data records are separated by line breaks.
-l | --long or **-F | --format**	The **-d | --delimiter** option is ignored.
other options	Data fields in a record are separated by a comma followed by a space. Data records are separated by the specified delimiter *symbol*.

For examples of this option, refer to "Examples" on page 565.

-o | --oid
Specifies that the object ID (OID) is displayed together with the managed-object name. The **-o** option may also be combined with the **-l** and **-A** options. For an example of this option, refer to "Examples" on page 565.

-p | --pipe
Specifies that the object ID (OID) is displayed instead of the managed-object name. When used alone, this option enables the output to be piped to other **dircli** commands. The **-p** option may also be combined with the **-l** and **-A** options. For an example of this option, refer to "Examples" on page 565.

-T | --showtype
Lists the managed-object type after the managed-object name. By default, the managed-object name and type are separated by a comma followed by a space.

-A | --attribute *key* [, *key* ...] **-s | --sort**
Specifies the managed-object attributes that are displayed. *key* is the attribute key. If you issue the **-s | --sort** option, the output is sorted based on the value of the first attribute key specified.

-F | --format
Specifies that the output is displayed in a format which may be saved as a managed-object definition file, or may be used as input for the **mkmo** command. Although the actual attributes displayed may vary with different managed-object types, the format is often like the following lines:

```
type = value
ip = value
name = value
network = value
```

Note: If you specify either the **-o | --oid** or **-p | --pipe** options, they are ignored.

-l | --long
Specifies that all managed-object attributes are displayed in a long format. The information is displayed in the following format:

```
Managed-object name
    key = value
    key = value
    key = value
    ...
```

For an example of this option, refer to "Examples" on page 565.

-e Used with the **-l** option, expands the long listing format to also display message strings for the attribute key and attribute value. The information is displayed in the following format:

```
<Managed-object name>
    key (key_string) = value (value_string)
    key (key_string) = value (value_string)
    key (key_string) = value (value_string)
    ...
```

Examples

List the name of all managed objects

The following command lists the names of all managed objects:

```
dircli lsmo
```

The command results in the following output:

```
_YUMA 36-
IDWORLD
ibmdirector.ibm.com
...
```

List managed objects based on attribute criteria

The following command lists the names of all managed systems that do not have a license for IBM Director:

```
dircli lsmo -w "MO.hasLicense=False"
```

The following command lists the names of all managed systems with POWER™ or 32-bit architecture:

```
dircli lsmo -w "MO.ArchType=POWER OR MO.ArchType=IA32"
```

The following command lists all Level-2 managed systems that are offline and have either POWER or 32-bit architecture:

```
dircli lsmo -w "MO.MOFID='Level 2: IBM Director Agents' AND
(Mo.ArchType=POWER OR Mo.ArchType=IA32) AND MO.state=Offline"
```

List specific attributes for managed objects

The following command lists the presence check interval and power status for all managed objects:

```
dircli lsmo -sA MO.ping,MO.status
```

List managed-object types

The following command lists the managed-object types supported by IBM Director:

```
dircli lsmo -i
```

The command results in the following output:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
RMON Devices
Racks
Remote I/O Enclosures
```

```
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Scalable Partition
Scalable Systems
Storage Devices
Windows Clusters
```

List the object ID with the -o option

```
dircli lsmo -o -N TestGroup728
```

The command results in the following output:

```
_YUMA 36-, 0x264
IDWORLD, 0x219
ibmdirector.ibm.com, 0x30e
```

The following command also uses the **-d** option to specify a record delimiter:

```
dircli lsmo -d " || " -o -N TestGroup728
```

The command results in the following output:

```
_YUMA 36-, 0x264 || IDWORLD, 0x219 || ibmdirector.ibm.com, 0x30e
```

List the object ID with the -p option

```
dircli lsmo -p -N TestGroup728
```

The command results in the following output:

```
0x264
0x219
0x30e
```

The following command also uses the **-d** option to specify a record delimiter:

```
dircli lsmo -d " || " -p -N TestGroup728
```

The command results in the following output:

```
0x264 || 0x219 || 0x30e
```

List specified attributes for managed objects in a group

The following command lists three managed-object attributes for the objects in group TestGroup728:

```
dircli lsmo -A MO.AgentType,MO.state,MO.OpSys -N TestGroup728
```

The command results in the following output:

```
_YUMA 36-: Unsupported, Online, WINDOWS_NT
IDWORLD: Director_Server, Online, WINDOWS_NT
ibmdirector.ibm.com: Agentless, Online,
```

List specified attributes for managed objects in a group (showing type)

The following command uses the **-T** option to display the managed-object type in addition to the specified attributes:

```
dircli lsmo -A MO.AgentType,MO.state,MO.OpSys -T -N TestGroup728
```

The command results in the following output:

```
_YUMA 36-, Level 2: IBM Director Agents: Unsupported, Online, WINDOWS_NT
IDWORLD, Level 2: IBM Director Agents: Director_Server, Online, WINDOWS_NT
ibmdirector.ibm.com, Level 0: Agentless Systems: Agentless, Online,
```

Specify a delimiter when listing specified attributes

The following command uses the **-d** option to set a delimiter between the specified attributes:

```
dircli lsmo -d " || " -A MO.AgentType,MO.state,MO.OpSys -N TestGroup728
```

The command results in the following output:

```
_YUMA 36-: Unsupported || Online || WINDOWS_NT
IDWORLD: Director_Server || Online || WINDOWS_NT
ibmdirector.ibm.com: Agentless || Online ||
```

Specify a delimiter when listing specified attributes and object ID

The following command uses the **-d** option to set a delimiter between the specified attributes:

```
dircli lsmo -d " || " -o -A MO.AgentType,MO.state,MO.OpSys -N TestGroup728
```

The command results in the following output:

```
_YUMA 36-, 0x264: Unsupported || Online || WINDOWS_NT
IDWORLD, 0x219: Director_Server || Online || WINDOWS_NT
ibmdirector.ibm.com, 0x30e: Agentless || Online ||
```

Note: The managed-object name and object ID are separated by a comma and space, *not* the delimiter characters.

Specify a delimiter when using the -p pipe option

The following command uses the **-d** option to set a delimiter between the specified attributes:

```
dircli lsmo -d " || " -p -A MO.AgentType,MO.state,MO.OpSys -N TestGroup728
```

The command results in the following output:

```
0x219: Director_Server || Online || WINDOWS_NT || 0x264: Unsupported ||
 Online || WINDOWS_NT || 0x30e: Agentless || Online ||
```

List attributes with the -F format option

The following command uses the **-F** option to specify both the listed attributes and their format:

```
dircli lsmo -F -N TestGroup728
```

The command results in the following output:

```
type=Systems
ip=9.42.221.37
name=_YUMA 36-
network=TCPIP

type=Systems
ip=192.168.1.12
name=IDWORLD
network=TCPIP

type=Systems
ip=192.168.1.12
name=ibmdirector.ibm.com
```

List attributes using the -l long format option

The following command uses the **-l** option to list all attributes in a long format:

```
dircli lsmo -l -N TestGroup728
```

The command results in the following output:

```
_YUMA 36-:
    MO.name = _YUMA 36-
    MO.MOFID = Level 2: IBM Director Agents
    MO.state = Online
    MO.ping = 15
    MO.hasLicense = true
    MO.secunsecsupport = false
    MO.accessdenied = true
    MO.encryptionenabled = true
    MO.agenttimezoneoff = -240
    NativeMO.path = TCPIP::9.42.221.37
    NativeMO.uniqueid = 83c3872a4455e994
    NativeMO.allpaths = {'TCPIP::9.42.221.37'}
    MO.IPaddrs = {'9.42.221.37'}
    MO.IPhosts = {'yuma.raleigh.ibm.com'}
    MO.OpSys = WINDOWS_NT
    MO.OpSysMajVer = 5
    MO.OpSysMinVer = 1
    MO.UUID = 0E97DCA7-1451-118D-ADEC-0E0A4EC936CB
    MO.MACAddress = 00096BA7F8F6
    MO.MACAddrList = {'00096BA7F8F6'}
    MO.ArchType = IA32

IDWORLD:
    MO.name = IDWORLD
    MO.MOFID = Level 2: IBM Director Agents
    MO.state = Online
    MO.ping = 15
    MO.hasLicense = true
    MO.secunsecsupport = true
    MO.unsecureclient = false
    MO.accessdenied = false
    MO.encryptionenabled = true
    MO.agenttimezoneoff = -240
    NativeMO.path = TCPIP::192.168.1.12
    NativeMO.uniqueid = 2e29e0074d56eb43
    NativeMO.allpaths = {'TCPIP::192.168.1.12'}
    MO.IPaddrs = {'192.168.1.12'}
    MO.IPhosts = {'ibmdirector.ibm.com'}
    MO.OpSys = WINDOWS_NT
    MO.OpSysMajVer = 5
    MO.OpSysMinVer = 0
    MO.AgentType = Director_Server
    NativeMO.agentProdVersion = 5.10
    MO.AgentDate = Fri Jul 22 08:00:00 EDT 2005
    MO.UUID = 05027C40-B440-124A-8E75-3890A12EC67A
    MO.MACAddress = 000255AC026C
    MO.ComputerName = IDWORLD
    MO.MACAddrList = {'000255AC026C'}
    MO.ArchType = IA32

ibmdirector.ibm.com:
    MO.name = ibmdirector.ibm.com
    MO.MOFID = Level 0: Agentless Systems
    MO.state = Online
    MO.ping = 15
    MO.secunsecsupport = false
    MO.accessdenied = true
    MO.encryptionenabled = false
    MO.IPaddrs = {'192.168.1.12'}
    MO.IPhosts = {'ibmdirector.ibm.com'}
    MO.AgentType = Agentless
    MO.Protocols = { 1 }
```

List expanded attributes using the -l -e expanded long format option

The following command uses the **-l -e** option to list attributes in an expanded long format:

```
dircli lsmo -l -e -n "_YUMA 36-"
```

The command results in the following output:

```
_YUMA 36-:
    MO.name (System Name) = _YUMA 36- (_YUMA 36-)
    MO.MOFID (System Factory ID) = Level 2: IBM Director Agents
      (Level 2: IBM Director Agents)
    MO.state (System State) = Online (Online)
    MO.ping (System Presence Check Setting (minutes)) = 15 (15)
    MO.hasLicense (Granted License) = true (true)
    ...
```

Return codes

The following table lists the codes returned by the **lsmo command**.

Code	Meaning
0	The managed objects were listed successfully.
1	A usage error occurred.
10	The file was not found.
20	A managed object was not found.
21	A group was not found.
25	A number-formatting error occurred.
26	An invalid object type was specified.
27	An invalid attribute was specified.

lstask

This topic provides information about the **lstask** command. This command lists IBM Director tasks.

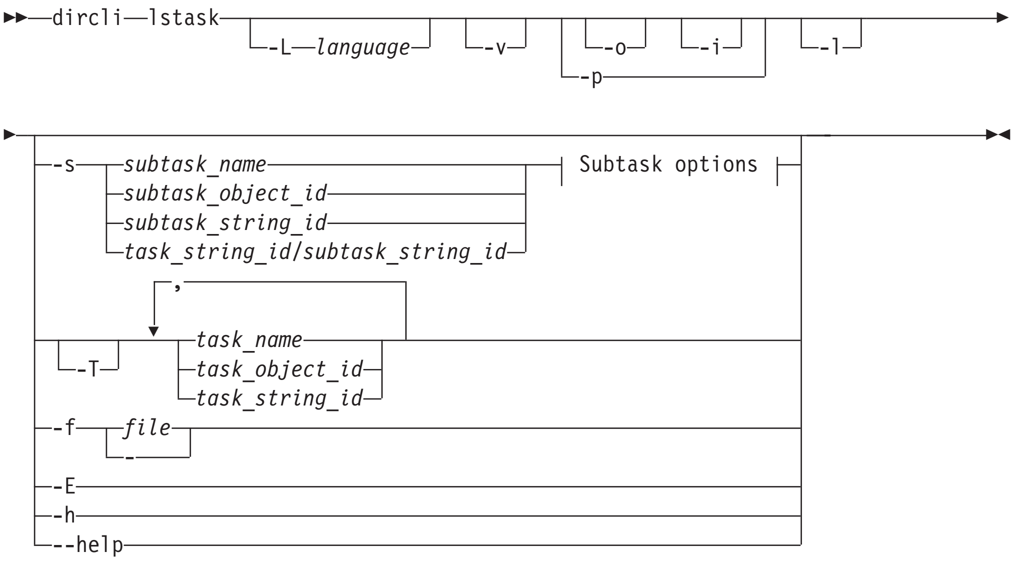

Subtask options:

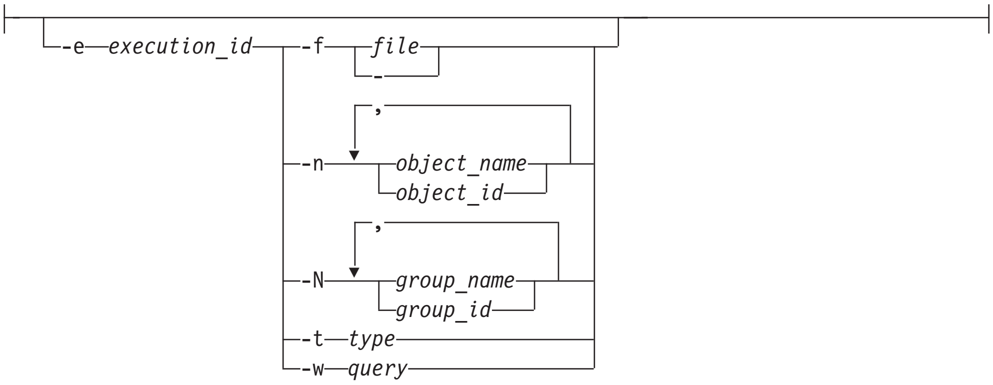

Options and operands

Note: If no options or operands are specified, the **lstask** command lists all IBM Director tasks.

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 27. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose

Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-o | --oid

Lists the object ID (OID) for tasks and the job ID for subtasks. A job ID of `0x0` indicates an interactive task.

-i | --id

Lists the unique task or subtask ID string.

-p | --pipe
Specifies that the task object ID (OID) and subtask job ID are displayed instead of the task name. When used alone, this option enables the output to be piped to other **dircli** commands.

-l | --long
Specifies that all task and subtask information is displayed. The information is displayed in the following format:

```
task_name
task_name/subtask_name
   ST.1.name = subtask_name
   ST.1.targeted = value
   ST.1.interactive = value
   ...
```

-s | --subtask {*subtask_name* | *subtask_object_id* | *subtask_string_id* | *task_string_id/subtask_string_id*}
Lists the status for a specific subtask.

subtask_name
Name of the subtask. Subtask names are locale-specific.

subtask_object_id
Object ID (OID) of the subtask. This is applicable to non-interactive subtasks only.

subtask_string_id
String ID of the subtask, starting with a % character. Subtask strings might not be unique.

task_string_id/subtask_string_id
Unique ID of the subtask, comprised of the string ID of the task (starting with a % character), a forward slash, and the string ID of the subtask.

The **-l** option cannot be combined with the **-s** option.

-T | --tasks {*task_name* | *task_object_id* | *task_string_id*} [, {*task_name* | *task_object_id* | *task_string_id*} ...]
Specifies the tasks against which the command is targeted.

task_name
Name of the task.

task_object_id
Unique object ID of the task.

task_string_id
String ID of the task.

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

-E | --executable
Lists the non-interactive subtasks that can be scheduled. The **-p** and **-l** options cannot be used with this option.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

-e | --exec *execution_id*
Lists the execution status for the specified managed objects and subtasks. *execution_id* is an integer. It is incremented by one every time the subtask is run.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]
Targets the managed objects specified by name or ID.

object_id
Unique ID of the managed object.

object_name
Name of the managed object.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]
Targets the managed objects that are in the specified groups. If a managed object belongs to more than one group, the managed object is targeted only once.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-t | --type *type*
In conjunction with a primary targeting option (explicit or default), specifies the type of managed object targeted. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

Notes: If you specify a parent managed-object type (for example, "Chassis"), its children (in this case, "BladeCenter Chassis") are also targeted.

Use the following command to list all managed-object types for your installation:

```
dircli lsmo -i
```

-w | --where *query*
Targets managed objects based on managed-object attribute values specified in *query*. The *query* operand is a quote-delimited string defining a simple SELECT query that uses the following format:

```
key_1=value_1 [AND|OR] key_2=value_2 ... [AND|OR] key_n=value_n
```

Enclose the statement in quotation marks. If any managed-object attributes are strings that include spaces, enclose the strings in single quotation marks. Use parentheses to nest logical constructions.

Examples

List all IBM Director tasks

The following command lists the names of all IBM Director tasks:

```
dircli lstask
```

When this command is run on a management server before any user-defined tasks are created, the following output is returned:

```
Active Console Viewer
Asset ID
BladeCenter Assistant
BladeCenter Assistant/BladeCenter Configuration Manager
BladeCenter Assistant/Switch Management launch pad
CIM Browser
Cluster Functions
Configure Alert Standard Format
Configure SNMP Agent
Discover All Systems and Devices
DMI Browser
Event Action Plan Builder
Event Action Plans
Event Action Plans/Event Action Plan Wizard
Event Action Plans/Log All Events
...
```

List the object IDs for all IBM Director tasks

The following command lists the object IDs for all tasks:

```
dircli lstask -p
```

This command returns the following output:

```
0x104
0x105
0x10b
0x10e
0x115
0x124
0x157
0x159
0x15b
0x15d
...
```

List names and object IDs for all IBM Director tasks

The following command lists the names and object IDs for all tasks:

```
dircli lstask -o
```

This command returns the following output:

```
Active Console Viewer, 0x1a4
Asset ID, 0x1a0
BladeCenter Management, 0x196
BladeCenter Management/BladeCenter Configuration Manager, 0x210
BladeCenter Management/Launch Web Browser, 0x20e
BladeCenter Management/Network Device Manager, 0x1e7
CIM Browser, 0x18b
Cluster Functions, 0x186
Configure Alert Standard Format, 0x18d
Configure SNMP Agent, 0x189
```

```
Discover All Managed Objects, 0x1f5
Event Action Plan Builder, 0x16a
Event Action Plans, 0x16b
Event Action Plans/cjaction, 0x225
```

List names and unique string IDs for all IBM Director tasks

The following command lists the names and unique string IDs for all tasks:

```
dircli lstask -i
```

```
Active Console Viewer, %Tivoli|TWGActiveConsoleViewer
Asset ID, %itdcim|AssetIDDirect
BladeCenter Management, %ibm|sysmgt|mpa|BladeTask
BladeCenter Management/BladeCenter Configuration Manager, %ibm|sysmgt|Con
figurationManager|BladeCenter
BladeCenter Management/Launch Web Browser, %BladeLaunchPadTask
BladeCenter Management/Network Device Manager, %SwitchMgtLauncher
CIM Browser, %Tivoli|CIMTasks
Cluster Functions, %Tivoli|ClusterTask
Configure Alert Standard Format, %itdcim|CimASFDirect
Configure SNMP Agent, %itdcim|CimSNMPDirect
...
```

List all task information

The following command lists all task and subtask information:

```
dircli lstask -l
```

For example, the output would include the following information about Event Action Plans:

```
Event Action Plan Builder:
   ST.1.name = Event Action Plan Builder
   ST.1.targeted = false
   ST.1.interactive = true

Event Action Plans:
   ST.1.name = Event Action Plans/Event Action Plan Wizard
   ST.1.targeted = false
   ST.1.interactive = true
   ST.2.name = Event Action Plans/Event Action Plan Builder
   ST.2.targeted = false
   ST.2.interactive = true

Event Action Plans/Event Action Plan Wizard:
   ST.1.name = Event Action Plans/Event Action Plan Wizard
   ST.1.targeted = false
   ST.1.interactive = true
```

List the task ID for the Asset ID task

The following command lists the task ID for the `Asset ID` task:

```
dircli lstask -T "Asset ID" -o
```

This command returns the following output:

```
Asset ID, 0x1d8
```

List information for a specific subtask

The following command lists information for the `LED On` subtask:

```
dircli lstask -s "LED On"
```

This command returns the following output:

```
System Identification/LED On:
   ST.targeted = true
   ST.interactive = false
```

List all information for a subtask

The following command lists all task and subtask information for the Power Management/Shutdown task:

```
dircli lstask -T 0x40 -o -l
```

This command returns the following output:

```
Power management/ShutDown, 0x40:
  ST.1.name = Power management/ShutDown, 0x4A
  ST.1.targeted = true
  ST.1.interactive = false
  ST.1.exec.1.stat = Active
  ST.1.exec.1.stat.client.0 = {Sys1, 0x1A} Pending
  ST.1.exec.1.stat.client.1 = {Sys2, 0x2B} Failed
  ST.1.exec.2.stat = Complete
  ST.1.exec.2.stat.client.0 = {ServerA, 0x5F}, Complete
  ST.1.exec.2.stat.client.1 = {ServerB, 0x6B}, Failed
```

List the job-execution status for a subtask

The following command lists the job-execution status for the task associated with OID 0x15 and execution ID 1:

```
dircli lstask -s 0x15 -e 1
```

This command returns the following output:

```
ST.1.exec.1.stat = Active
  ST.1.exec.1.stat.client.0 = Sys1, Pending
  ST.1.exec.1.stat.client.1 = Sys2, Failed
```

Return codes

The following table contains the codes returned by the **lstask** command.

Code	Meaning
0	The tasks were listed successfully.
1	A usage error occurred.
10	The file was not found.
20	The managed object was not found.
22	The task was not found.
25	A number-formatting error occurred.
50	The subtask was not found.
51	The subtask is not unique.
52	The subtask is interactive.
53	The execution ID was not found.

mkcmprof

This topic provides information about the **mkcmprof** command. This command creates BladeCenter configuration profiles. Once a profile is created, it can be run against a managed object(s) using the **runtask** command.

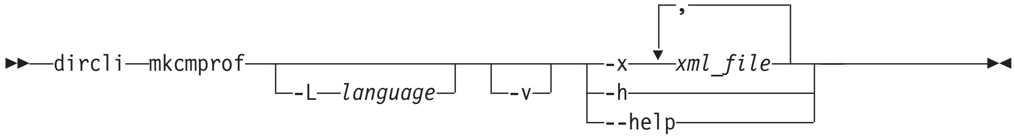

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 28. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose

Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-x | --xmlfile *xml_file*[,*xml_file* ...]

Specifies the name of the XML file that contains the BladeCenter configuration information. *xml_file* is the fully qualified name of the XML file. If you list more than one XML file, separate the file names with commas.

-h | -?

Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help

Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

Create a BladeCenter configuration profile

The following command creates a BladeCenter configuration profile using the information stored in the configuration.xml file:

```
dircli mkcmprof -x c:\temp\configuration.xml
```

Create two BladeCenter configuration profiles

The following command creates two BladeCenter configuration profiles:

```
dircli mkcmprof -x c:\temp\profile1.xml,c:\temp\profile2.xml
```

Return codes

The following table contains the codes returned by the **mkcmprof** command.

Code	Meaning
0	The profile was successfully created.
1	A usage error occurred.
50	Some failed.
51	All failed but with different reasons.
52	Failed with details.

mkgp

This topic provides information about the **mkgp** command. This command creates groups.

```
►►──dircli──mkgp──┬───────────────┬──┬─────┬──────────────────────────────────►
                  └─-L──language──┘  └─-v──┘

►──┬─-f──┬─file─┬───────────────────────────────────────────┬────────────►◄
   │     └─-────┘                                           │
   │         ┌─,─────────────────┐                          │
   │         ▼                   │                          │
   ├─-n──────┬─object_name─┬─────┴──new_group_name──────────┤
   │         └─object_id───┘                                │
   │               ┌─,───────────────────┐                  │
   │               ▼                     │                  │
   ├─┬─────┬───────┬─task_name─────────┬─┴──new_group_name──┤
   │ └─-T──┘       ├─task_object_id────┤                    │
   │               └─task_string_id────┘                    │
   │         ┌─,─────────────────┐                          │
   │         ▼                   │                          │
   ├─-N──────┬─group_name─┬──────┴──new_group_name──────────┤
   │         └─group_id───┘                                 │
   ├─-h─────────────────────────────────────────────────────┤
   └─--help─────────────────────────────────────────────────┘
```

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 29. Common values for the language code

Language/country	*language*
Chinese/China	zh_CN

Table 29. Common values for the language code (continued)

Language/country	*language*
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose
Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Each line of the input data represents a **mkgp** instruction to create a static group, task group, or group category, and uses one of the following three formats:

```
new_group_name:Static:{object_id|object_name}[,{object_id|object_name}...]
new_group_name:Task:{task_name|task_object_id|task_string_id}[,{task_name|task_object_id|task_string_id}...]
new_group_name:GroupCategory:{group_id|group_name}[,{group_id|group_name}...]
```

Each group definition in the input file must be separated by a line break.

new_group_name
Specifies the name of the group being created.

object_id
Unique ID of the managed object.

object_name
Name of the managed object.

task_name
Name of the task.

task_object_id
Unique object ID of the task.

task_string_id
String ID of the task.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...] *new_group_name*

Creates a static group named *new_group_name* and comprised of the listed managed objects.

new_group_name
Specifies the name of the group being created.

object_id
Unique ID of the managed object.

object_name
Name of the managed object.

-T | --tasks {*task_name* | *task_object_id* | *task_string_id*} [, {*task_name* | *task_object_id* | *task_string_id*} ...] *new_group_name*

Creates a task-based group named *new_group_name* and comprised of the listed tasks.

new_group_name
Specifies the name of the group being created.

task_name
Name of the task.

task_object_id
Unique object ID of the task.

task_string_id
String ID of the task.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...] *new_group_name*

Creates a group category named *new_group_name* and comprised of the listed groups.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

new_group_name
Specifies the name of the group being created.

-h | -?

Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help

Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

Create a static group

The following command creates a static group (NewGroup) that contains two managed objects:

```
dircli mkgrp -n node1,node2 NewGroup
```

Create a group that supports the CIM Browser task

The following command creates a dynamic, task-based group named CIM:

```
dircli mkgp -T "CIM Browser" CIM
```

Create a group using a definition file

The following command uses the data in the c:\temp\MyNewGroup.txt file to create three new groups:

```
dircli mkgp -f c:\temp\MyNewGroup.txt
```

The MyNew.txt file contains the following code:

```
group1:Static:mo_id1,mo_id2,mo_id3
group2:Task:task_id1,task_id2,task_id3
group3:GroupCategory:group_id1,group_id2,group_id3
```

Return codes

The following table contains the codes returned by the **mkgp** command.

Code	Meaning
0	The group was created successfully.
1	A usage error occurred.
10	The file was not found.
20	A managed object was not found.
21	A group was not found.
22	A task was not found.
25	A number-formatting error occurred.
50	An invalid group-definition format was specified.
51	No valid group definitions were specified.

mkmo

This topic provides information about the **mkmo** command. This command creates managed objects and lists required managed-object attributes for object creation.

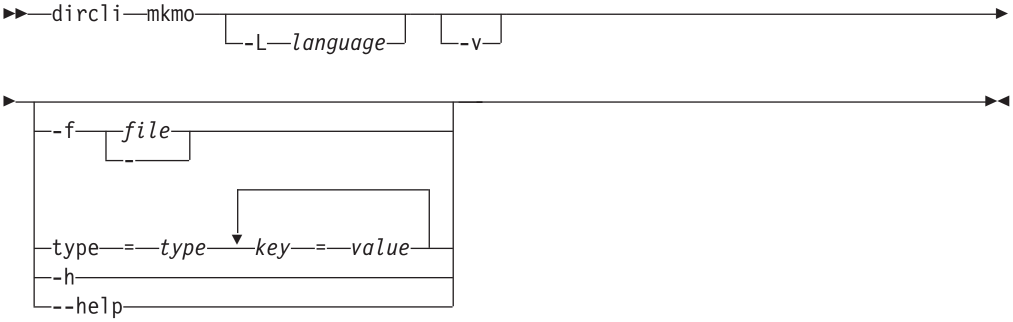

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 30. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose
Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

type=*type*
Specifies the type of managed object being created. The *type* variable may be one of the following values:

```
BladeCenter Chassis
Physical Platforms
SMI-S Storage Devices
SNMP Devices
Systems
Windows Clusters
```

Depending on the value of *type*, one or more additional attribute-key values may need to be specified.

Issue the **mkmo** command without any options to list the managed-object types and required attribute-key value pairs.

key=value
Specifies attribute values for the managed object being created. Depending on the value of *type*, one or more attribute-key values may need to be specified. Refer to “Examples” on page 582, or issue the **mkmo** command without any options to list the managed-object types and required attribute-key value pairs.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

List the managed-object definition information

The following command lists the information required to create a managed object:

```
dircli mkmo
```

It returns the following text:

```
BladeCenter Chassis:
type=BladeCenter Chassis
name=<Specify Name>
ip=<Specify IP Address>
userid=<Specify User ID>
password=<Specify Password>

Physical Platforms:
type=Physical Platforms
name=<Specify Name>  (Optional)
ip=<Specify IP Address>

SMI-S Storage Devices:
type=SMI-S Storage Devices
name=<Specify Name>
ip=<Specify IP Address>

SNMP Devices:
type=SNMP Devices
ip=<Specify IP Address>
version=<Specify SNMP Version>
community=<Specify Community Name>
seed=<Specify "true" if this should be used as a discovery seed>

Systems:
type=Systems
name=<Specify Name>  (Optional)
ip=<Specify IP Address>
network=<Specify Network Protocol>  (Optional)
Available Protocols: TCPIP)

Windows Clusters:
type=Windows Clusters
name=<Specify Name>
label=<Specify Label>
```

Create a Level-2 managed system

The following command creates a Level-2 managed system using the information specified in the c:\temp\managed_object.txt file:

```
dircli mkmo -f c:\temp\managed_object.txt
```

Create a Level-2 managed system

The following command creates a Level-2 managed system using the information specified:

```
dircli mkmo type=Systems name=WebServer network=TCPIP ip=192.168.1.100
```

Return codes

The following table contains the codes returned by the **mkmo** command:

Code	Meaning
0	The managed object was created successfully.
1	A usage error occurred.

Code	Meaning
10	The definition file was not found.
26	An invalid platform was specified.
27	An invalid attribute was specified.
50	An invalid object definition was specified.

pingmo

This topic provides information about the **pingmo** command. This command pings managed objects. When the ping operation is completed, the results are reflected in the managed-object state attribute (MO.state). This command performs an IBM Director presence check and not an ICMP ping.

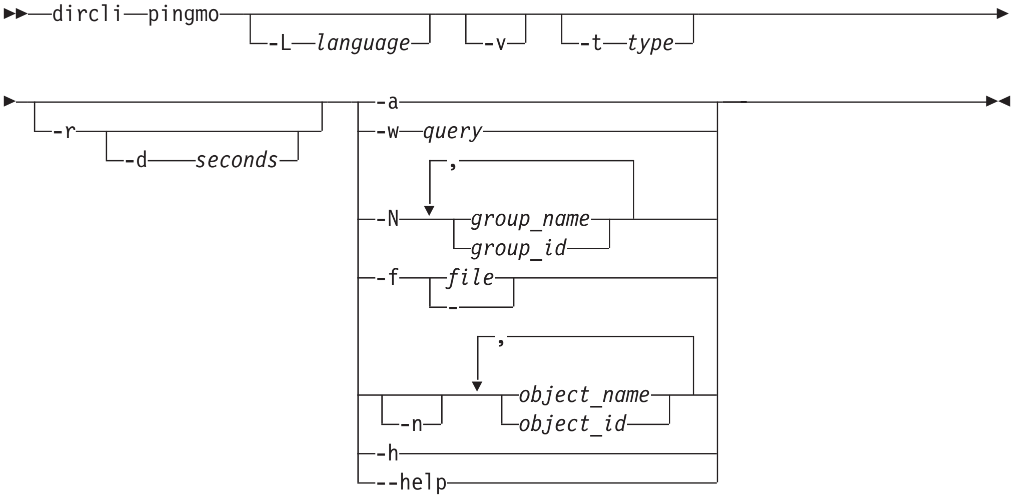

Options and operands

-L | --lang *language*
Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 31. Common values for the language code

Language/country	*language*
Chinese/China	zh_CN
English/United States	en_US
French/France	fr_FR
German/Germany	de_DE
Japanese/Japan	ja_JP
Korean/South Korea	ko_KR
Spanish/Mexico	es_MX

-v | --verbose
Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-t | --type *type*
In conjunction with a primary targeting option (explicit or default), specifies the type of managed object targeted. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

Notes: If you specify a parent managed-object type (for example, ″Chassis″), its children (in this case, ″BladeCenter Chassis″) are also targeted.

Use the following command to list all managed-object types for your installation:

```
dircli lsmo -i
```

-r | --return -d | --delay *seconds*
Returns the value of the MO.state attribute for all target objects. If you do not use the **-d | --delay** *seconds* option, by default, the value is returned after 10 seconds. To change the time interval, issue the **-d | --delay** *seconds* option. *seconds* is the number of seconds.

-a | --all
Targets all managed objects.

-w | --where *query*
Targets managed objects based on managed-object attribute values specified in *query*. The *query* operand is a quote-delimited string defining a simple SELECT query that uses the following format:

```
key_1=value_1 [AND|OR] key_2=value_2 ... [AND|OR] key_n=value_n
```

Enclose the statement in quotation marks. If any managed-object attributes are strings that include spaces, enclose the strings in single quotation marks. Use parentheses to nest logical constructions.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]
Targets the managed objects that are in the specified groups. If a managed object belongs to more than one group, the managed object is targeted only once.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]
Targets the managed objects specified by name or ID.

object_id
Unique ID of the managed object.

object_name
Name of the managed object.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

Ping a managed object
The following command pings the managed system named `webserver`:

```
dircli pingmo webserver
```

Ping several managed objects
The following command pings the managed objects listed in the c:\temp\mobject.txt file:

```
dircli pingmo -f c:\temp\mobject.txt
```

Return codes

The following table contains the codes returned by the **pingmo** command.

Code	Meaning
0	The ping operation was started successfully.
1	A usage error occurred
10	The file was not found.
20	The managed object was not found.
21	The group was not found
25	A number-formatting error occurred.
26	A invalid object type was specified.
27	An invalid attribute type was specified.

rmcmprof

This topic provides information about the **rmcmprof** command. This command deletes BladeCenter configuration profiles.

```
►►──dircli──rmcmprof──┬─────────────┬──┬────┬──────────────────────────►
                      └─-L──language─┘  └─-v─┘


               ┌─,───────────────┐
               ▼                 │
►──┬──┬────┬──┬─task_name──────┬─┴──┬──────────────────────────────────►◄
   │  └─-T─┘  ├─task_object_id─┤    │
   │          └─task_string_id─┘    │
   ├─-h─────────────────────────────┤
   └─--help─────────────────────────┘
```

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 32. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose

Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-T | --tasks {*task_name* | *task_object_id* | *task_string_id*} [, {*task_name* | *task_object_id* | *task_string_id*} ...]

Specifies the tasks against which the command is targeted.

task_name

Name of the task.

task_object_id

Unique object ID of the task.

task_string_id

String ID of the task.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Example

Remove a profile using a task string ID
The following command removes the BladeCenter configuration profile by specifying the task string ID:

```
dircli rmcmprof -T ibm|sysmgt|ConfigurationManager|BladeCenter|1112195540453
```

Remove profiles using a task object ID
The following command removes the BladeCenter configuration profiles associated with the task object ID 0xA and 0xB:

```
dircli rmcmprof -T 0xA,0xB
```

Return codes

The following table lists the codes returned by the **rmcmprof** command.

Code	Meaning
1	The BladeCenter configuration profile was deleted successfully.
25	A usage error occurred.
50	Some failed.
51	All failed but with different reasons.
52	Failed with details.

rmgp

This topic provides information about the **rmgp** command. This command deletes groups.

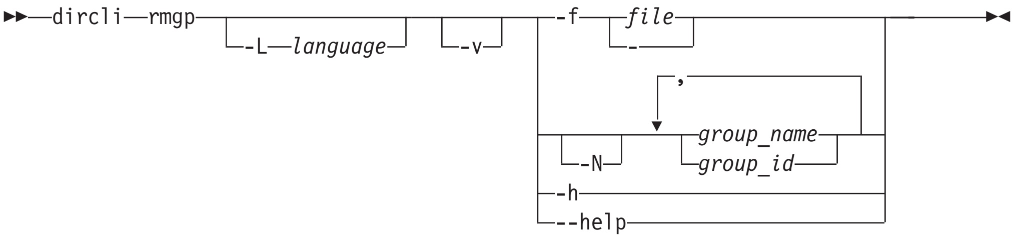

Options and operands

-L | --lang *language*
Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 33. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose
Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. The input data must specify the groups using either group names or group IDs. Items in the input file must be separated by commas or line breaks.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]
Targets the groups specified by name or ID.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

Delete two groups
The following command deletes two groups:

```
dircli rmgrp "ID systems","Test systems"
```

Delete a group using a definition file
The following command uses the data in the c:\temp\MyNewGroup.txt file to delete several groups:

```
dircli rmgp -f c:\temp\MyNewGroup.txt
```

Return codes

The following table contains the codes returned by the **rmgp** command.

Code	Meaning
0	The group was removed successfully.
1	A usage error occurred.
10	The file was not found.
21	A group was not found.
25	A number-formatting error occurred.
52	A group cannot be deleted.

rmmo

This topic provides information about the **rmmo** command. This command deletes managed objects.

```
►►──dircli──rmmo──┬───────────────┬──┬────┬──┬─────────┬──────────────────────►
                  └─-L──language──┘  └─-v─┘  └─-t──type─┘

►──┬─-w──query──────────────────────────┬──────────────────────────────────►◄
   │         ┌─,──────────────┐          │
   │         ▼                │          │
   ├─-N──────┬─group_name─┬───┴──────────┤
   │         └─group_id───┘              │
   ├─-f──┬─file─┬────────────────────────┤
   │     └─-────┘                        │
   │               ┌─,────────────────┐  │
   │               ▼                  │  │
   ├──┬────┬───────┬─object_name─┬────┴──┤
   │  └─-n─┘       └─object_id───┘       │
   ├─-h──────────────────────────────────┤
   └─--help──────────────────────────────┘
```

Options and operands

-L | --lang *language*
Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 34. Common values for the language code

Language/country	*language*
Chinese/China	zh_CN
English/United States	en_US
French/France	fr_FR
German/Germany	de_DE
Japanese/Japan	ja_JP
Korean/South Korea	ko_KR

Table 34. Common values for the language code (continued)

Language/country	*language*
Spanish/Mexico	es_MX

-v | --verbose
Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-t | --type *type*
In conjunction with a primary targeting option (explicit or default), specifies the type of managed object targeted. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

Notes: If you specify a parent managed-object type (for example, ″Chassis″), its children (in this case, ″BladeCenter Chassis″) are also targeted.

Use the following command to list all managed-object types for your installation:

```
dircli lsmo -i
```

-w | --where *query*
Targets managed objects based on managed-object attribute values specified in *query*. The *query* operand is a quote-delimited string defining a simple SELECT query that uses the following format:

```
key_1=value_1 [AND|OR] key_2=value_2 ... [AND|OR] key_n=value_n
```

Enclose the statement in quotation marks. If any managed-object attributes are strings that include spaces, enclose the strings in single quotation marks. Use parentheses to nest logical constructions.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]
Targets the managed objects that are in the specified groups. If a managed object belongs to more than one group, the managed object is targeted only once.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]
Targets the managed objects specified by name or ID.

object_id
Unique ID of the managed object.

object_name
Name of the managed object.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Examples

Remove managed objects
The following command deletes the specified managed objects:

```
dircli rmmo -n webserver,0x9A5
```

Remove managed objects
The following command deletes the managed objects that are specified in the /tmp/mobjects file:

```
dircli rmmo -f /tmp/mobjects
```

Return codes

The following table contains the codes returned by the **rmmo** command.

Code	Meaning
0	The managed object was successfully deleted.
1	A usage error occurred.
10	A file was not found.
20	The managed object was not found.
21	The group was not found.
25	A number-formatting error occurred.
26	An invalid object type was specified.
27	An invalid attribute type was specified.
50	The managed object cannot be removed.
51	The operation failed.

rpower

This topic provides information about the **rpower** command. This command performs power-management operations.

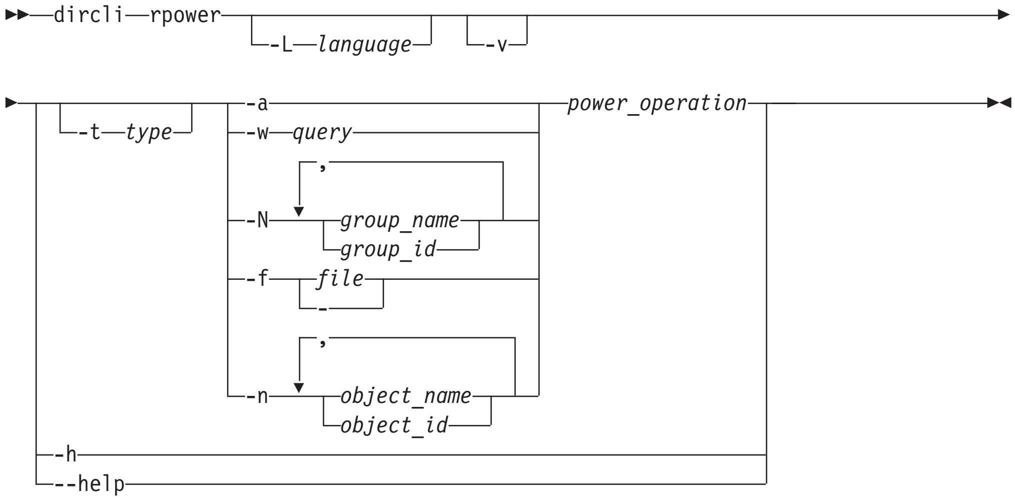

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 35. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose

Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-a | --all

Targets all managed objects.

-w | --where *query*

Targets managed objects based on managed-object attribute values specified in *query*. The *query* operand is a quote-delimited string defining a simple SELECT query that uses the following format:

```
key_1=value_1 [AND|OR] key_2=value_2 ... [AND|OR] key_n=value_n
```

Enclose the statement in quotation marks. If any managed-object attributes are strings that include spaces, enclose the strings in single quotation marks. Use parentheses to nest logical constructions.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]

Targets the groups specified by name or ID.

group_id

Unique ID of the IBM Director managed-object group.

group_name

Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-f | --file {*file* | -}

Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]

Targets the managed objects specified by name or ID.

object_id

Unique ID of the managed object.

object_name

Name of the managed object.

-t | --type *type*

In conjunction with a primary targeting option (explicit or default), specifies the type of managed object targeted. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

Notes: If you specify a parent managed-object type (for example, "Chassis"), its children (in this case, "BladeCenter Chassis") are also targeted.

Use the following command to list all managed-object types for your installation:

```
dircli lsmo -i
```

power_operation

The *power_operation* operand is a string which specifies the power-management operation to be performed. The following table lists the possible values for *power_operation* and describes the power-management operation performed.

Note: Not all operations are supported on all managed objects.

Value	Operation performed
`PowerOffNow`	Powers off the system
`PowerOn`	Powers on the system
`PowerOnHold`	Powers on but does not continue with boot operations
`PowerOnRelease`	Resumes boot operations after a `PowerOnHold` operation
`Restart`	Restarts the operating system after a `Suspend` operation
`RestartNow`	Forces a restart of the operating system
`Resume`	Resumes operating system activity for a suspended system
`ShutDown`	Shuts down the operating system
`ShutDownAndPowerOff`	Shuts down the operating system and then powers off the system
`Support`	Lists the power operations supported for the targeted systems
`Suspend`	Suspends the operating system

These strings are not case sensitive.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Example

Restart a group of managed systems
The following command restarts all the managed systems in the group named `bladecenter1`:

```
dircli rpower -N bladecenter1 restart
```

The command returns the names and object IDs for the targeted systems, the power operation performed, and the results of the operation.

Return codes

The following table contains the codes returned by the **rpower** command.

Code	Meaning
0	The power-management operation started successfully.
1	A usage error occurred.
10	The file was not found.
20	A managed object was not found.
21	A group was not found.
25	A number-formatting error occurred.
26	A invalid attribute type was specified

Code	Meaning
27	An invalid attribute was specified.
50	An invalid power-management operation was specified.
51	The power-management operation failed.
52	One or more managed systems were locked.

runtask

This topic provides information about the **runtask** command. This command runs non-interactive IBM Director tasks or profiles created with the **mkcmprof** command.

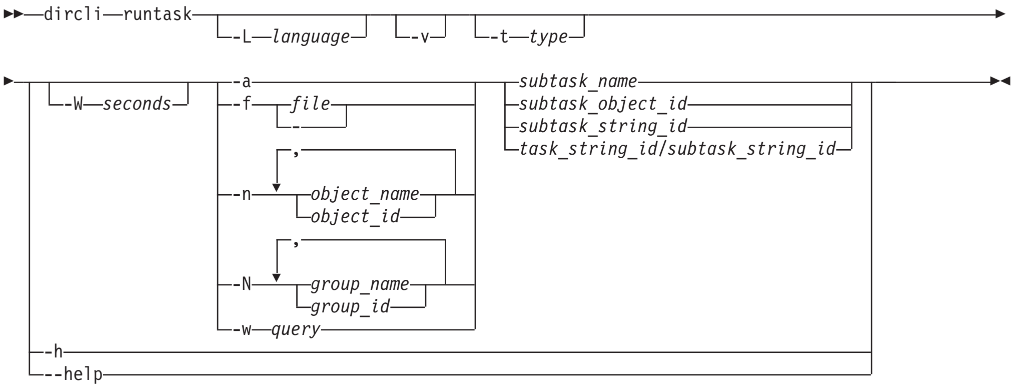

Options and operands

-L | --lang *language*

Specifies the locale for the command. This option overrides all mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 36. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

-v | --verbose

Writes verbose messages to standard output. Unless verbose messaging is selected, the **dircli** command suppresses non-critical messages to the user.

For example, when the **accessmo** command is run against a managed object that is already unlocked, normally no message is displayed. With verbose messaging, the following message appears:

```
The following managed objects already are unlocked: 'IBM-KLGGN3P', 0x254.
```

-t | --type *type*
In conjunction with a primary targeting option (explicit or default), specifies the type of managed object targeted. The variable *type* is one of the following strings:

```
All Managed Objects
BladeCenter Chassis
Chassis
Clusters
HMC
Level 0: Agentless Systems
Level 1: Core Services
Level 2: IBM Director Agents
Logical Platforms
Physical Platforms
Platforms
Racks
Remote I/O Enclosures
RMON Devices
Scalable Partitions
Scalable Systems
SMI-S Storage Devices
SNMP Devices
SNMP Printers
Storage Devices
Windows Clusters
```

Notes: If you specify a parent managed-object type (for example, ″Chassis″), its children (in this case, ″BladeCenter Chassis″) are also targeted.

Use the following command to list all managed-object types for your installation:

```
dircli lsmo -i
```

-W | --wait *seconds*
Specifies that the results of the command are displayed after waiting an interval specified in *seconds*, regardless of whether or not the task has been completed. If *seconds* is set to `0`, the command waits indefinitely for the task to be completed before listing the results.

-a | --all
Targets all managed objects.

-f | --file {*file* | -}
Targets objects based on information that is provided from either the specified input *file* or from the standard input pipe. To receive piped input, use a hyphen instead of a filename. Items in the input file must be separated by commas or line breaks.

-n | --names {*object_id* | *object_name*} [, {*object_id* | *object_name*} ...]
Targets the managed objects specified by name or ID.

object_id
Unique ID of the managed object.

object_name
Name of the managed object.

-N | --groups {*group_id* | *group_name*} [, {*group_id* | *group_name*} ...]
Targets the managed objects that are in the specified groups. If a managed object belongs to more than one group, the managed object is targeted only once.

group_id
Unique ID of the IBM Director managed-object group.

group_name
Name of the IBM Director managed-object group. If the group name contains space characters, it should be quoted.

-w | --where *query*
Targets managed objects based on managed-object attribute values specified in *query*. The *query* operand is a quote-delimited string defining a simple SELECT query that uses the following format:

```
key_1=value_1 [AND|OR] key_2=value_2 ... [AND|OR] key_n=value_n
```

Enclose the statement in quotation marks. If any managed-object attributes are strings that include spaces, enclose the strings in single quotation marks. Use parentheses to nest logical constructions.

subtask_name
Name of the subtask. Subtask names are locale-specific.

subtask_object_id
Object ID (OID) of the subtask. This is applicable to non-interactive subtasks only.

subtask_string_id
String ID of the subtask, starting with a % character. Subtask strings might not be unique.

task_string_id/subtask_string_id
Unique ID of the subtask, comprised of the string ID of the task (starting with a % character), a forward slash, and the string ID of the subtask.

-h | -?
Lists brief information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

--help
Lists complete information about the command. If any additional options other than **-L | lang** are issued, they are ignored.

Example

Run a task on a managed object
The following command runs a task with ID string `%bluelight_ON` on a managed object with the name `WebServer`:

```
dircli runtask -n WebServer %bluelight_ON
```

Return codes

The following table contains the codes returned by the **runtask** command.

Code	Meaning
0	The task started successfully.
1	A usage error occurred.
10	The file was not found.
20	A managed object was not found.
21	A group was not found.
25	A number-formatting error occurred.
26	A invalid attribute type was specified
27	An invalid attribute was specified.

Code	Meaning
50	The task was not found.
51	The task is not unique.
52	The task is interactive.
53	The task failed. One or more managed objects are not available.
54	The task failed to start.

Locale specification for dircli commands

This topic describes the various locale-specification methods for dircli commands, and the precedence in which these methods are applied.

Although the command locale for dircli commands may be specified with the **-L | –lang** *language* option in the command syntax, the default locale for dircli commands may be set in two other ways: by setting the DIR_LANG environment variable, or by retrieving locale information from system environment variables or the operating system itself. The locale-specification methods are applied in the following order:

-L | --lang *language* **option in dircli command syntax**

Specifying the locale explicitly for a dircli command using the **-L | --lang** *language* option overrides all other mechanisms for specifying the locale, including the DIR_LANG environment variable and operating-system settings. The *language* code is a text string consisting of five characters: a two-character lowercase ISO 639-1 language code, an underscore character, and a two-character uppercase ISO 3166 country code. Following are some common examples of the *language* code.

Table 37. Common values for the language code

Language/country	*language*
Chinese/China	`zh_CN`
English/United States	`en_US`
French/France	`fr_FR`
German/Germany	`de_DE`
Japanese/Japan	`ja_JP`
Korean/South Korea	`ko_KR`
Spanish/Mexico	`es_MX`

DIR_LANG environment variable

If the dircli command does not specify the locale explicitly using the **-L | --lang** *language* option, the value of the DIR_LANG environment variable (if set) is used to specify the locale. Acceptable values for DIR_LANG are the same as for the **-L | --lang** *language* option.

Operating System locale specification

If the dircli command does not specify a locale and the DIR_LANG environment variable is not set, dircli uses the following operating-system information to determine the locale:

Operating system	Locale information used
UNIX	The values of three environment variables are used in the following order: 1. LC_ALL 2. LANG 3. LC_LANG
Windows	The locale specified in Regional and Language Options is used.

IBM Director dircmd-command bundles

This topic lists command bundles which can be executed using either the **dircmd** or **dircli** command-line interface.

Both implementations are documented here; however, use of the **dircli** command-line interface is strongly recommended for four reasons:

- **dircli** has enhanced security compared with **dircmd**
- **dircli** is a high-performance client and runs approximately 10-100 times faster than **dircmd**, depending on configuration
- **dircli** is available on all platforms, while **dircmd** is limited to IBM xSeries servers
- Support for the **dircmd** command-line interface is unlikely to continue beyond version 5.10 of IBM Director

Bundle	Description
"BladeCenter-chassis bundle" on page 600	BladeCenter-chassis bundle commands facilitate management and configuration of managed BladeCenter-chassis devices.
"Chassis bundle" on page 602	Chassis bundle commands facilitate management and configuration of chassis managed objects.
"Command-line interface bundle" on page 604	Command-line interface bundle commands enable users to list available commands and bundles.
"Event-management bundle" on page 607	Event-management bundle commands facilitate management of events and event action plans.
"Level 2 managed-system bundle" on page 611	Managed-system bundle commands facilitate adding and listing Level-2 managed systems with IBM Director Agent.
"Management Processor Assistant bundle" on page 614	Management Processor Assistant bundle commands allow you to list information about managed objects or set credentials for communicating with a service processor.
"Process-monitor bundle" on page 617	Process-monitor bundle commands are used to manage process-monitor tasks.
"Resource-monitor bundle" on page 619	Resource-monitor bundle commands apply or list resource-monitor threshold tasks.
"Scheduler bundle" on page 622	Scheduler bundle commands are used to retrieve information about scheduled jobs and to cancel activation of jobs.
"Server-management bundle" on page 624	Server-management bundle commands provide the ability to list inventory values, and to create or list dynamic groups.
"SNMP-device bundle" on page 629	SNMP-device bundle commands facilitate management and configuration of managed SNMP devices.
"User-administration bundle" on page 634	User-administration bundle commands facilitate reporting and modifying user access privileges.

In addition to the bundle topics, two other reference topics provide additional information relating to the **dircmd** command-line interface.

Topic	Description
“IBM Director dircmd usage” on page 641	This topic describes the syntax and options specific to **dircmd** but not associated with any command bundles.
“Deprecated dircmd commands” on page 643	This topic lists commands formerly available through the **dircmd** command-line interface that are now deprecated. The equivalent functionality is now implemented through new commands in the **dircli** command-line interface.

BladeCenter-chassis bundle

This topic describes BladeCenter-chassis bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. BladeCenter-chassis bundle commands facilitate management and configuration of managed BladeCenter-chassis devices.

General syntax (dircli)

```
►►──dircli──bladecenterchassis──┬─/─┬──┤ Command-argument syntax ├──────────────►◄
                                └───┘
```

General syntax (dircmd)

```
►►──dircmd──-s──server──-u──user──-p──password──┤ options ├──────────────────►

►──bladecenterchassis──┤ Command-argument syntax ├───────────────────────────►◄
```

-s *server*
: Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
: Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
: Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to “IBM Director dircmd usage” on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

```
├──┬─addbcchassis─chassis_name─net_address─userid─password─┬──────────────┤
   ├─discoverbcchassis──────────────────────────────────────┤
   ├─help───────────────────────────────────────────────────┤
   ├─list───────────────────────────────────────────────────┤
   └─listbcchassis──┬──────┬─────────────────────────────────┘
                    ├─-r─┤
                    └─-t─┘
```

addbcchassis

```
├──addbcchassis─chassis_name─net_address─userid─password──────────────┤
```

Adds a BladeCenter-chassis managed object.

discoverbcchassis

```
├──discoverbcchassis──────────────────────────────────────────────────┤
```

Starts a BladeCenter chassis discovery.

help

```
├──help───────────────────────────────────────────────────────────────┤
```

Displays general help for the bundle usage.

list

```
├──list───────────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

listbcchassis

```
├──listbcchassis──┬──────┬────────────────────────────────────────────┤
                  ├─-r─┤
                  └─-t─┘
```

Lists all BladeCenter-chassis managed objects in one of two formats:

Format	Information listed
report (default)	object ID and object name
terse	object ID

Options and operands

chassis_name
The name of the BladeCenter chassis.

net_address
A TCP/IP address.

userid
An IBM Director user ID that is authorized to access the managed object.

password
The password for the IBM Director user account that is authorized to access the managed object.

-r | -report
Enables a detailed (report) listing format. Refer to the command description for specific information about what is listed when applied to that command.

-t | -terse
Enables terse listing format. Refer to command description for specific information about what is listed when applied to that command.

Example

Discover BladeCenter chassis objects
The following command executes a discovery task for BladeCenter chassis objects:

```
dircli bladecenterchassis/discoverbcchassis
```

Chassis bundle

This topic describes Chassis bundle commands available through either IBM Director command-line interface: dircli or dircmd. Chassis bundle commands facilitate management and configuration of chassis managed objects.

General syntax (dircli)

```
►►──dircli──chassis──┬─/─┬──┤ Command-argument syntax ├──────────────────────►◄
                     └───┘
```

General syntax (dircmd)

```
►►──dircmd──-s──server──-u──user──-p──password──┤ options ├──────────────────►

►──chassis──┤ Command-argument syntax ├──────────────────────────────────────►◄
```

-s *server*
Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to “IBM Director dircmd usage” on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

```
├──┬─chassislist──┬─────┬──────────────────┬──────────────────────────────────────┤
   │              ├─-r──┤                  │
   │              └─-t──┘                  │
   ├─chassissubsystemlist──chassis_oid──┬─────┬┤
   │                                    ├─-r──┤│
   │                                    └─-t──┘│
   ├─chassissubsystemtypelist──────────────────┤
   ├─help──────────────────────────────────────┤
   └─list──────────────────────────────────────┘
```

chassislist

```
├──chassislist──┬─────┬────────────────────────────────────────┤
                ├─-r──┤
                └─-t──┘
```

Lists all chassis managed objects in one of two formats:

Format	Information listed
report (default)	object ID and object name
terse	object ID

chassissubsystemlist

```
├──chassissubsystemlist──chassis_oid──┬─────┬──────────────────┤
                                      ├─-r──┤
                                      └─-t──┘
```

List subsystems present on a chassis managed object in one of two formats:

Format	Information listed
report (default)	object ID and object name
terse	object ID

chassissubsystemtypelist

```
├──chassissubsystemtypelist────────────────────────────────────┤
```

Lists supported subsystem types of a chassis managed object.

help

```
├──help────────────────────────────────────────────────────────┤
```

Displays general help for the bundle usage.

list

```
├──list──────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

Options and operands

-r | -report
Enabled a detailed (report) listing format. Refer to the command description for specific information about what is listed when applied to that command.

-t | -terse
Enables terse listing format. Refer to command description for specific information about what is listed when applied to that command.

chassis_oid
The managed object ID of the chassis.

Example

Display help for the chassis bundle
In the following example, an IBM Director superuser connects to the management server with the host name `IDWorld`, using the user ID `InfoDeveloper` and the password `passw0rd`, and invokes the help function.

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd chassis help
```

Command-line interface bundle

This topic describes command-line interface bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. Command-line interface bundle commands enable users to list available commands and bundles.

General syntax (dircli)

```
►►──dircli──cli──┬─/─┬──┤ Command-argument syntax ├──────────────────────────►◄
                 └───┘
```

General syntax (dircmd)

```
►►──dircmd──-s──server──-u──user──-p──password──┤ options ├──────────────────►

►──cli──┤ Command-argument syntax ├──────────────────────────────────────────►◄
```

-s *server*
Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.

- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to “IBM Director dircmd usage” on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

```
├──┬─help──────┬──────────────────────────────────────────────────┤
   ├─list──────┤
   └─listbundle┘
```

help

```
├──help──────────────────────────────────────────────────────────┤
```

Displays general help for the bundle usage.

list

```
├──list──────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

listbundle

```
├──listbundle────────────────────────────────────────────────────┤
```

Lists the installed bundles and commands for IBM Director.

Options and operands

This command does not accept any options or operands.

Example

List bundles

The following command lists the installed command bundles:

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd cli listbundle
```

This command returns the following output:

```
BladeCenterChassis/addbcchassis
BladeCenterChassis/discoverbcchassis
BladeCenterChassis/help
BladeCenterChassis/list
BladeCenterChassis/listbcchassis
Chassis/chassislist
Chassis/chassissubsystemlist
Chassis/chassissubsystemtypelist
Chassis/help
Chassis/list
cli/help
cli/listbundle
event/applyeventactionplan
event/createeventactionplan
```

```
event/help
event/list
event/listeventactionplans
event/listeventactions
event/listevents
event/listeventtypes
event/listfilters
monitor/applythreshold
monitor/help
monitor/list
monitor/listthresholds
mpa/help
mpa/list
mpa/listobjectattributes
mpa/listobjectattributevalues
mpa/listobjectsbyattribute
mpa/mpcli
mpa/setcredentials
mpa/setsysinterconnectconnection
native/addsystem
native/help
native/list
native/listsystems
native/startdiscovery
procmon/applypmtask
procmon/createpmtask
procmon/help
procmon/list
procmon/listpmtasks
scheduler/canceljobactivation
scheduler/cancelrdmtaskactivation
scheduler/getjobactivationlog
scheduler/getjobstatus
scheduler/help
scheduler/list
scheduler/listjobactivations
scheduler/listjobactivationsbysystem
scheduler/listjobs
server/accessobjects
server/addtostaticgroup
server/createdynamicgroup
server/createstaticgroup
server/deletegroups
server/deleteobjects
server/discoverall
server/doconstraintdump
server/help
server/list
server/listdynamicgroupcriteria
server/listgroupattributes
server/listgroupmembers
server/listgroups
server/listgroupsbyattribute
server/listinventoryvalues
server/listnoninteractivetasks
server/listobjectattributes
server/listobjects
server/listobjectsbyattribute
server/listtaskactivationstatus
server/pingobjects
server/removefromstaticgroup
server/renameobject
server/runtask
snmp/addsystem
snmp/get
snmp/getbulk
snmp/getnext
```

```
snmp/help
snmp/inform
snmp/list
snmp/listsystems
snmp/set
snmp/startdiscovery
snmp/trap
snmp/walk
user/addgroupaccess
user/addtaskaccess
user/help
user/list
user/listgroups
user/listprivilegetokens
user/listtasks
user/listuserattributes
user/listusers
user/modifyuserattributes
user/removegroupaccess
user/removetaskaccess
```

Event-management bundle

This topic describes event-management bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. Event-management bundle commands facilitate management of events and event action plans.

General syntax (dircli)

```
►►──dircli──event──┬─/─┬──┤ Command-argument syntax ├──────────────────────────►◄
                   └───┘
```

General syntax (dircmd)

```
►►──dircmd──-s──server──-u──user──-p──password──┤ options ├─────────────────────►

►─event──┤ Command-argument syntax ├────────────────────────────────────────────►◄
```

-s *server*
: Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
: Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
: Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to "IBM Director dircmd usage" on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

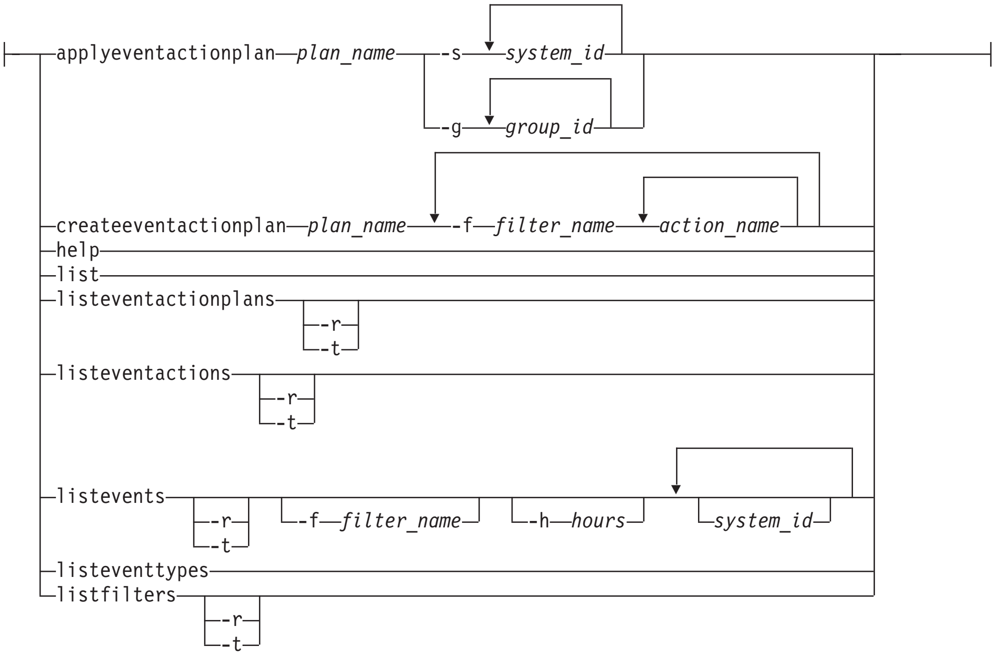

applyeventactionplan

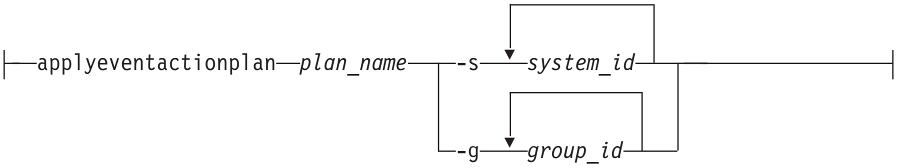

Applies an event action plan to a managed object or group.

createeventactionplan

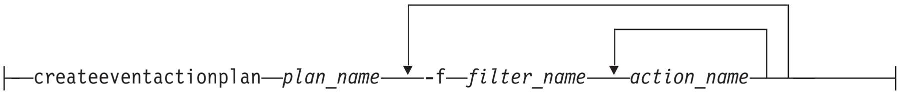

Creates an event action plan.

help

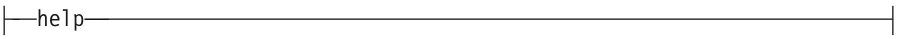

Displays general help for the bundle usage.

list

```
├──list───────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

listeventactionplans

```
├──listeventactionplans──┬────┬───────────────────────────────────┤
                         ├─-r─┤
                         └─-t─┘
```

Lists all event action plans in one of three formats:

Format	Information listed
report	event action plan name, key, and read-only status
terse	event action plan name
neither (default)	event action plan name and key

listeventactions

```
├──listeventactions──┬────┬───────────────────────────────────────┤
                     ├─-r─┤
                     └─-t─┘
```

Lists all event actions in one of three formats:

Format	Information listed
report	event action name, key, read-only status, run-able status, and logging properties
terse	event action name
neither (default)	event action name and key

listevents

```
                                                      ┌─────────────┐
                                                      ▼             │
├──listevents──┬────┬──┬───────────────┬──┬─────────┬──┬───────────┬┴──┤
               ├─-r─┤  └─-f─filter_name─┘  └─-h─hours─┘  └─system_id─┘
               └─-t─┘
```

Lists the contents of the event log in one of three formats:

Format	Information listed
report	event type, event date and time, event system, severity, category, sending system, and associated text description
terse	event and event system
neither (default)	event type, event date and time, event system, severity, and category

By default, all events in the past 24 hours are listed. You can limit the list by specifying a filter, setting a shorter time frame, or specifying the managed objects.

listeventtypes

```
├──listeventtypes──────────────────────────────────────────────────────────────┤
```

Lists the published event list.

listfilters

```
├──listfilters──┬──────┬───────────────────────────────────────────────────────┤
                ├─-r──┤
                └─-t──┘
```

Lists all event filters in one of three formats:

Format	Information listed
report	event filter name, key, and read-only status
terse	event filter name
neither (default)	event filter name and key

Options and operands

plan_name
The unique event action plan name.

system_id
The unique object ID for the managed system. The **listobjects** command can be used to list valid system IDs.

group_id
A group ID.

filter_name
The unique event filter name. The **listfilters** command can be used to list valid filter names.

action_name
The event action name. The **listeventactions** command can be used to list valid action names.

-r | -report
Enables a detailed (report) listing format. Refer to the command description for specific information about what is listed when applied to that command.

-t | -terse
Enables terse listing format. Refer to command description for specific information about what is listed when applied to that command.

hours
Specifies the number of hours.

Examples

List fatal events

In the following example, an IBM Director `superuser` connects to the management server with the host name `IDWorld`, using the user ID `InfoDeveloper` and the password `passw0rd`. When the user invokes the **listevents** function of the events-management bundle, the following command returns a list of all fatal events that occurred in the previous 8 hours.

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd event listevents -f
"Fatal Events" -h 8
```

List event types

In the following example, the user from the previous example invokes the **listeventtypes** function in combination with a **grep** command to list all IBM Director event types that are associated with security.

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd event listeventtypes |
grep -i "security"
```

List events for managed objects listed by the listgroupmembers command

The following example illustrates piping data from one command to another **dircmd** command. The server/listgroupmembers command is executed using the **dircli** command-line interface, then its output is piped to the **dircmd** command, which uses the **-r** option to indicate it will receive piped input. The entire command should be entered on one line.

```
dircli server/listgroupmembers -t 17D | dircmd -r -s IDworld -u
InfoDeveloper -p passw0rd event listevents
```

Level 2 managed-system bundle

This topic describes Level 2 managed-system bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. Managed-system bundle commands facilitate adding and listing Level-2 managed systems with IBM Director Agent.

General syntax (dircli)

```
►►──dircli──native──┬─/─┬─┤ Command-argument syntax ├────────────────────►◄
                    └───┘
```

General syntax (dircmd)

```
►►──dircmd──-s──server──-u──user──-p──password──┤ options ├──────────────►

►─native──┤ Command-argument syntax ├────────────────────────────────────►◄
```

-s *server*

Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*

Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*

Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to “IBM Director dircmd usage” on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

```
├──┬─addsystem──system_name──protocol──net_address─┬──────────────────────────────┤
   ├─help────────────────────────────────────────────┤
   ├─list────────────────────────────────────────────┤
   ├─listsystems──┬─────┬──────────────────────────────┤
   │              ├─-r─┤
   │              └─-t─┘
   └─startdiscovery──────────────────────────────────┘
```

addsystem

```
├──addsystem──system_name──protocol──net_address──────────────────────┤
```

Creates a Level-2 managed-system object on the management server. This command is equivalent to right-clicking the Group Contents pane of IBM Director Console and then clicking **New** → **Systems**.

listsystems

```
├──listsystems──┬─────┬───────────────────────────────────────────────┤
                ├─-r─┤
                └─-t─┘
```

Lists all level 2 managed systems in one of three formats:

Format	Information listed
report	system name, object ID (OID), unique ID (UID), MAC address, universal unique ID (UUID), IBM Director Agent version, state, whether access is denied, operating system information, IP address, and host name for each managed system
terse	object ID
neither (default)	system name and object ID

help

```
├──help───────────────────────────────────────────────────────────────┤
```

Displays general help for the bundle usage.

list

```
├──list───────────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

startdiscovery

```
├──startdiscovery─────────────────────────────────────────────────────┤
```

Discovers managed systems.

Options and operands

system_name
The name of the managed system.

protocol
The network protocol used by the managed object.

net_address
A TCP/IP address.

-r | -report
Enables a detailed (report) listing format. Refer to the command description for specific information about what is listed when applied to that command.

-t | -terse
Enables terse listing format. Refer to command description for specific information about what is listed when applied to that command.

Examples

Add a managed object

The following example invokes the **addsystem** function to add a managed-system object to the IBM Director environment. The new managed system is displayed in IBM Director Console as `TechWriter2` with TCP/IP as the network protocol and has an IP address of `160.0.0.27`.

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd native addsystem TechWriter2
TCPIP 160.0.0.27
```

List managed systems

The following three examples illustrate the different reporting formats for the **listsystems** command.

```
dircli native listsystems
20F  IDWORLD
230  EBERLEIN

dircli native listsystems -t
20F
230

dircli native listsystems -r
System Name .................IDWORLD
Object ID (oid) .............20F
Unique ID (uid) .............2E29E0074D56EB43
MAC Address .................{ '000255AC026C' }
Universal Unique ID (uuid)...407C020540B44A128E753890A12EC67A
Agent Version................5.10
State .......................Online
Access Denied ...............false
Operating System ............WINDOWS_NT
OS Major Version ............5
OS Minor Version ............0
IP Address ..................{ '9.42.174.12' }
Hostname ....................{ 'idworld.raleigh.ibm.com' }

System Name .................EBERLEIN
Object ID (oid) .............230
Unique ID (uid) .............45FBE8284C4591B9
MAC Address .................{ '00096BA7F407' }
Universal Unique ID (uuid)...E33688B7C6648D119F20864E64D4C382
Agent Version................-
State .......................Online
```

```
Access Denied ...............true
Operating System ...........WINDOWS_NT
OS Major Version ...........5
OS Minor Version ...........1
IP Address .................{ '9.42.174.13' }
Hostname ...................{ 'Eberlein.raleigh.ibm.com' }
```

Management Processor Assistant bundle

This topic describes Management Processor Assistant (MPA) bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. Management Processor Assistant bundle commands allow you to list information about managed objects or set credentials for communicating with a service processor.

General syntax (dircli)

```
►►──dircli──mpa──┬─/─┬──┤ Command-argument syntax ├───────────────────────────►◄
                 └───┘
```

General syntax (dircmd)

```
►►──dircmd──-s──server──-u──user──-p──password──┤ options ├───────────────────►

►──mpa──┤ Command-argument syntax ├───────────────────────────────────────────►◄
```

-s *server*
Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to "IBM Director dircmd usage" on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

```
├──┬─help────────────────────────────────────────────────────────────┬───┤
   ├─list────────────────────────────────────────────────────────────┤
   ├─listobjectattributes──┬─────┬───────────────────────────────────┤
   │                       ├─-r──┤                                   │
   │                       └─-t──┘                                   │
   │                              ┌──────────────┐                   │
   ├─listobjectattributevalues────▼─system_id────┴───────────────────┤
   │                                 ┌────────────────────────────┐  │
   ├─listobjectsbyattribute──┬─────┬─▼─attribute──=──value────────┴──┤
   │                         └─-t──┘                                 │
   │                                                     ┌────────┐  │
   └─setcredentials──-u──userid──-p──password──system_oid▼system_oid┴┘
```

help

```
├──help──────────────────────────────────────────────────────────────┤
```

Displays general help for the bundle usage.

list

```
├──list──────────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

listobjectattributes

```
├──listobjectattributes──┬─────┬─────────────────────────────────────┤
                         ├─-r──┤
                         └─-t──┘
```

Lists the managed-object attributes in one of three formats:

Format	Information listed
report	name, data type, and value range
terse	name
neither (default)	name and data type

The data listed can be used as input to the **listobjectsbyattribute** command.

listobjectattributevalues

```
                               ┌──────────────┐
├──listobjectattributevalues───▼─system_id────┴──────────────────────┤
```

Lists the current values of the following attributes of the specified managed objects:

- textID
- assetTag
- lastConnectionStatus
- promptAccess
- compEvents

- assetType

listobjectsbyattribute

```
                                       ┌─────────────────────────┐
                                       ▼                         │
├──listobjectsbyattribute──┬──────┬────attribute──=──value───────┴──────────────┤
                           └──-t──┘
```

Lists information about managed objects that meet the specified criteria. You can determine valid managed-object attributes and the range of possible values by using this command.

setcredentials

```
                                                                ┌──────────────┐
                                                                ▼              │
├──setcredentials──-u──userid──-p──password──system_oid──────system_oid────────┴───────┤
```

Specifies the user ID and password for communicating with a service processor.

Options and operands

-r | -report
Enables a detailed (report) listing format. Refer to the command description for specific information about what is listed when applied to that command.

-t | -terse
Enables terse listing format. Refer to command description for specific information about what is listed when applied to that command.

system_id
The unique object ID for the managed system. The **listobjects** command can be used to list valid system IDs.

attribute
The name of a managed-object attribute.

value
The value of a managed-object attribute. *value* is case-sensitive.

userid
An IBM Director user ID that is authorized to access the managed object.

password
The password for the IBM Director user account that is authorized to access the managed object.

system_oid
The unique ID of the managed object.

Examples

List MPA-related attributes for managed objects
In the following example, an IBM Director superuser connects to the management server with the host name `IDWorld`, using the user ID `InfoDeveloper` and the password `passw0rd`. The user invokes the **listobjectattributevalues** command to return the current values of MPA-related attributes for the managed system with the object ID `16B`.

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd mpa listobjectsbyattribute 16B
```

Set credentials for managed objects

In the following example, the user sets the user ID `madison` and the password `lucas` to access the service processors contained in the managed objects with the object IDs `1F0`, `1F1`, and `1F2`.

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd mpa setcredentials -u madison
-p lucas 1F0 1F1 1F2 1F3
```

Process-monitor bundle

This topic describes process-monitor bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. Process-monitor bundle commands are used to manage process-monitor tasks.

General syntax (dircli)

```
►►──dircli──procmon──┬─/─┬──┤ Command-argument syntax ├──────────────────►◄
                     └───┘
```

General syntax (dircmd)

```
►►──dircmd──-s──server──-u──user──-p──password──┤ options ├──────────────►

►──procmon──┤ Command-argument syntax ├──────────────────────────────────►◄
```

-s *server*
Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to "IBM Director dircmd usage" on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

```
                              ┌──────────────┐
                              ▼              │
├──┬─applypmtask─task_id───system_oid─┴───────────────────────────────────────────┬──────┤
   │                                   ┌─────────────────────────────────────────┐│
   │                                   ▼                                         ││
   ├─createpmtask─task_name───application_path─┬────┬──┬────┬──┬──────────────┬─┴┤
   │                                           └─+S─┘  └─+E─┘  └─+F─seconds───┘  │
   ├─help────────────────────────────────────────────────────────────────────────┤
   ├─list────────────────────────────────────────────────────────────────────────┤
   └─listpmtasks─┬─────┬─────────────────────────────────────────────────────────┘
                 ├─-r─┤
                 └─-t─┘
```

applypmtask

```
                                ┌──────────────┐
                                ▼              │
├──applypmtask─task_id───system_oid─┴──────────────────────────┤
```

Applies a process-monitor task to a managed system.

Use the **listpmtasks** function to determine valid task IDs. Use the **listobjects** function to determine valid system IDs.

createpmtask

```
                                   ┌─────────────────────────────────────────────┐
                                   ▼                                             │
├──createpmtask─task_name───application_path─┬────┬──┬────┬──┬──────────────┬─┴──┤
                                            └─+S─┘  └─+E─┘  └─+F─seconds───┘
```

Creates a process-monitor task for a program.

help

```
├──help──────────────────────────────────────────────────────────┤
```

Displays general help for the bundle usage.

list

```
├──list──────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

listpmtasks

```
├──listpmtasks─┬─────┬───────────────────────────────────────────┤
               ├─-r─┤
               └─-t─┘
```

Lists all the process-monitor tasks in one of three formats:

Format	Information listed
report	task name, object ID, process monitor programs, start monitor status, stop monitor status, fail monitor status, and fail timeout seconds
terse	task object ID
neither (default)	task name and object ID

Options and operands

task_id
The task ID.

system_oid
The unique ID of the managed object.

task_name
The task name.

application_path
The path and name of the application, for example, `c:\windows\notepad.exe`.

+S Generates an event when the program begins.

+E Generates an event when the program ends.

+F*seconds*
Generates an event when the program does not start correctly or fails after the specified number of seconds.

-r | -report
Enables a detailed (report) listing format. Refer to the command description for specific information about what is listed when applied to that command.

-t | -terse
Enables terse listing format. Refer to command description for specific information about what is listed when applied to that command.

Examples

Create a task
In the following example, an IBM Director superuser connects to the management server with the host name `IDWorld`, using the user ID `InfoDeveloper` and the password `passw0rd`. When the user invokes the **createpmtask** function, the following command creates a process-monitor task with the name `Notepad monitor` that generates an event if the program does not start correctly or fails after `5` seconds.

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd procmon createPMtask
"Notepad monitor" c:\winnt\notepad.exe+s+f5
```

List tasks
In the following example, the user from the previous example invokes the **listpmtasks** function to list all process-monitor tasks.

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd procmon listPMtasks
```

Resource-monitor bundle

This topic describes resource-monitor bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. Resource-monitor bundle commands apply or list resource-monitor threshold tasks.

General syntax (dircli)

```
►►──dircli──monitor──┬─/─┬──┤ Command-argument syntax ├──────────────────────►◄
                     └───┘
```

General syntax (dircmd)

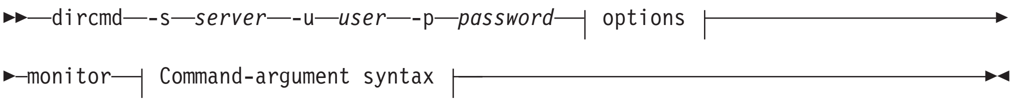

-s *server*
Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to "IBM Director dircmd usage" on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

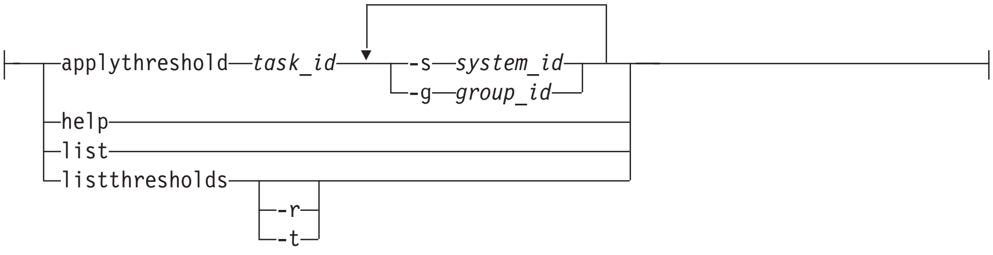

applythreshold

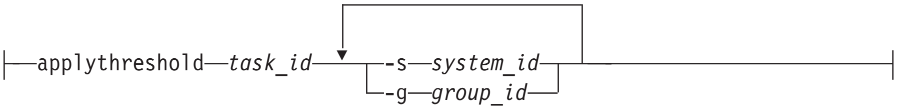

Applies a resource-monitor threshold task to a managed system or group.

help

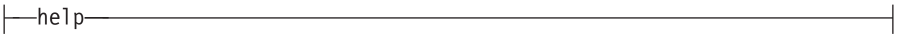

Displays general help for the bundle usage.

list

```
├──list──────────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

listthresholds

```
├──listthresholds──┬────┬────────────────────────────────────────────┤
                   ├─-r─┤
                   └─-t─┘
```

Lists all the resource-monitor threshold tasks in one of two formats:

Format	Information listed
report (default)	threshold name and task object ID
terse	task object ID

Options and operands

task_id
The task ID.

system_id
The unique object ID for the managed system. The **listobjects** command can be used to list valid system IDs.

group_id
A group ID.

-r | -report
Enables a detailed (report) listing format. Refer to the command description for specific information about what is listed when applied to that command.

-t | -terse
Enables terse listing format. Refer to command description for specific information about what is listed when applied to that command.

Examples

List threshold tasks

In the following example, an IBM Director superuser connects to the management server with the host name `IDWorld`, using the user ID `InfoDeveloper` and the password `passw0rd`. When the user invokes the **listthresholds** function, the following command lists all previously created threshold tasks.

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd monitor listthresholds -r
```

Apply threshold task

In the following example, the user from the previous example invokes the **applythreshold** function to apply the threshold task that is associated with OID `196` ("CPU utilization," in this case) to group `191` ("Systems with Windows 2000," in this case).

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd monitor applythreshold 196 -g 191
```

Scheduler bundle

This topic describes Scheduler bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. Scheduler bundle commands are used to retrieve information about scheduled jobs and to cancel activation of jobs.

General syntax (dircli)

```
►►—dircli—scheduler——┬─/─┬─┤ Command-argument syntax ├──────────────────────►◄
                     └───┘
```

General syntax (dircmd)

```
►►—dircmd——-s—server——-u—user——-p—password—┤ options ├──────────────────────►

►—scheduler—┤ Command-argument syntax ├──────────────────────────────────────►◄
```

-s *server*
: Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
: Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
: Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to "IBM Director dircmd usage" on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

```
├──┬─canceljobactivation—job_id—job_activation_id──────────┬──────────────┤
   ├─cancelrdmtaskactivation—job_id—job_activation_id—mo_id─┤
   ├─getjobactivationlog—job_id—job_activation_id─┬───────┬─┤
   │                                              └─mo_id─┘ │
   ├─getjobstatus—job_id────────────────────────────────────┤
   ├─help───────────────────────────────────────────────────┤
   ├─list───────────────────────────────────────────────────┤
   ├─listjobactivations—job_id─┬───────┬────────────────────┤
   │                           └─mo_id─┘                    │
   ├─listjobactivationsbysystem—mo_id───────────────────────┤
   └─listjobs───────────────────────────────────────────────┘
```

canceljobactivation

```
├──canceljobactivation──job_id──job_activation_id──────────────────────┤
```

Cancels activation of the specified job.

cancelrdmtaskactivation

```
├──cancelrdmtaskactivation──job_id──job_activation_id──mo_id───────────┤
```

Cancels activation of the specified Remote Deployment Manager (RDM) task.

getjobactivationlog

```
├──getjobactivationlog──job_id──job_activation_id──┬───────┬───────────┤
                                                   └─mo_id─┘
```

Lists the activation log for the specified job.

getjobstatus

```
├──getjobstatus──job_id────────────────────────────────────────────────┤
```

Returns the job status of the specified job.

help

```
├──help────────────────────────────────────────────────────────────────┤
```

Displays general help for the bundle usage.

list

```
├──list────────────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

listjobactivations

```
├──listjobactivations──job_id──┬───────┬───────────────────────────────┤
                               └─mo_id─┘
```

Lists all activations of all jobs with the specified job ID (and optionally, on the specified managed object).

listjobactivationsbysystem

```
├──listjobactivationsbysystem──mo_id───────────────────────────────────┤
```

Lists all job activations for the specified managed object.

listjobs

```
├──listjobs────────────────────────────────────────────────────────────┤
```

Lists all scheduled jobs.

Options and operands

job_id
The job ID.

job_activation_id
The job activation ID.

mo_id
The managed-object ID.

Examples

List scheduled jobs
The following example uses the **listjobs** command to list all the jobs defined on the system. In this example, only one job is defined.

```
dircli scheduler listjobs
```

The command results in the following output:

```
1e      sat_test_job
```

Display job status
The following example uses the **getjobstatus** command to display the status of the job returned in the previous example.

```
dircli scheduler getjobstatus 1e
```

The command results in the following output:

```
Job Status : 1e Active
```

Server-management bundle

This topic describes server-management bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. Server-management bundle commands provide the ability to list inventory values, and to create or list dynamic groups.

Note: Additional functions which were available through server-management bundle commands in IBM Director prior to version 5.10 are now provided through commands for the **dircli** command-line interface.

General syntax (dircli)

```
►►—dircli—server——┬─/─┬──┤ Command-argument syntax ├——————————————————————►◄
                  └───┘
```

General syntax (dircmd)

```
►►—dircmd——-s—server——-u—user——-p—password——┤ options ├————————————————————►

►—server——┤ Command-argument syntax ├——————————————————————————————————————►◄
```

-s *server*
Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
: Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
: Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to "IBM Director dircmd usage" on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

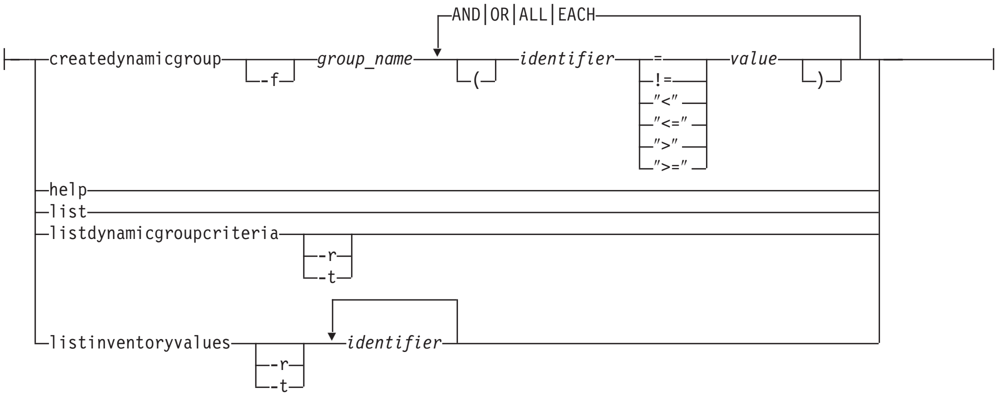

createdynamicgroup

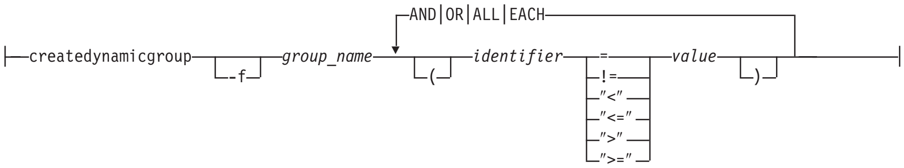

Creates a dynamic group. This is equivalent to creating a dynamic group in IBM Director Console.

help

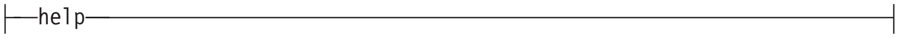

Displays general help for the bundle usage.

list

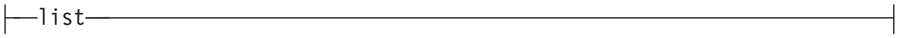

Lists the commands available in the bundle.

listdynamicgroupcriteria

```
├──listdynamicgroupcriteria──┬─────┬──────────────────────────────────┤
                             ├─-r─┤
                             └─-t─┘
```

Lists the criteria, based on database inventory, that are available for creating dynamic groups in one of three formats:

Format	Information listed
report	database, table, column, identifier, data type, multiple-rows-per-entity supported, operator-supported flag (=,!=), and operator-supported flag (>,>=,<,<=)
terse	identifier
neither (default)	identifier, database, table, and column

Notes: The database, table, and column name are language-translated strings.

You can use the returned value of this function as input to the createdynamicgroup function.

listinventoryvalues

```
                                  ┌────────────┐
                                  ▼            │
├──listinventoryvalues──┬─────┬───identifier─┴──────────────────────┤
                        ├─-r─┤
                        └─-t─┘
```

Lists the database inventory values for the specified identifiers in one of three formats:

Format	Information listed
report	database, table, column, identifier, data type, multiple-rows-per-entity supported, operator-supported flag (=,!=), operator-supported flag (>,>=,<,<=), and inventory values
terse	identifier and inventory values
neither (default)	identifier, database, table, column, and inventory values

Options and operands

-f Forces an equality relationship with a value that cannot be verified with the current database.

group_name
A name for the group being created.

identifier
The unique inventory identifier, in the form *DatabaseToken.TableToken.ColumnToken*.

value

An inventory value of the same type as the unique inventory identifier. It must be an existing inventory data, or, if the **-f** option is used, an unknown value can be used.

Successive *identifier-value* pairs in a series are linked using one of the following keywords:

- **AND** - all group criteria are met
- **OR** - any of the group criteria are met
- **ALL** - all group criteria are met for the same row
- **EACH** - any of the group criteria are met for the row

Use parentheses to nest group criteria for combining logical operations. The **ALL** and **EACH** keywords require that multiple rows per entity must be supported for the specified identifiers. When multiple rows per entity are supported, **AND** behaves like **ALL**.

-r | -report

Enables a detailed (report) listing format. Refer to the command description for specific information about what is listed when applied to that command.

-t | -terse

Enables terse listing format. Refer to command description for specific information about what is listed when applied to that command.

Examples

List inventory values

The following command lists inventory values for the identifier `PC_INV.TWG_INSTALLED_MEMORY.PHYSICAL_MEMORY_KB`:

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd server
    listinventoryvalues PC_INV.TWG_INSTALLED_MEMORY.PHYSICAL_MEMORY_KB
```

The command results in the following output (only one managed object with inventory):

```
PC_INV.TWG_INSTALLED_MEMORY.PHYSICAL_MEMORY_KB  "Inventory (PC)/Installed Memory/Physical Memory installed (KB)"
"508356"
```

List inventory values (report format)

The following command lists inventory values for the identifier `PC_INV.TWG_INSTALLED_MEMORY.PHYSICAL_MEMORY_KB`:

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd server
    listinventoryvalues -r PC_INV.TWG_INSTALLED_MEMORY.PHYSICAL_MEMORY_KB
```

The command results in the following output (only one managed object with inventory):

```
Database ...........................Inventory (PC)
Table ..............................Installed Memory
Column .............................Physical Memory installed (KB)
Identifer ..........................."PC_INV.TWG_INSTALLED_MEMORY.PHYSICAL_MEMORY_KB"
Data Type...........................Integer
Multiple Rows per Entity Supported..false
{=,!=} Operators Supported..........true
{<,<=,>,>=} Operators Supported.....true
Inventory Values:
"508356"
```

List inventory values (terse format)

The following command lists inventory values for the identifier `PC_INV.TWG_INSTALLED_MEMORY.PHYSICAL_MEMORY_KB`:

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd server
    listinventoryvalues -t PC_INV.TWG_INSTALLED_MEMORY.PHYSICAL_MEMORY_KB
```

The command results in the following output (two managed objects with inventory):

```
PC_INV.TWG_INSTALLED_MEMORY.PHYSICAL_MEMORY_KB
"508356" "8883945"
```

List dynamic group criteria

The following command lists identifiers which can be used to define dynamic groups or for listing inventory values:

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd server listdynamicgroupcriteria
```

The command results in the following output:

```
PC_INV.UMS_AOL.RetryInterval  "Inventory (PC)/Alert On LAN Settings/Retransmission Interval"
PC_INV.UMS_AOL.EventPollInterval  "Inventory (PC)/Alert On LAN Settings/Polling Interval"
PC_INV.UMS_AOL.HeartbeatInterval  "Inventory (PC)/Alert On LAN Settings/Heartbeat Interval"
PC_INV.UMS_AOL.Version  "Inventory (PC)/Alert On LAN Settings/Hardware Version"
PC_INV.UMS_AOL.AlertDestinationAddress  "Inventory (PC)/Alert On LAN Settings/Alert Destination"
PC_INV.UMS_AOL.EventAutoClearEnabled  "Inventory (PC)/Alert On LAN Settings/Auto Clear Events"
PC_INV.UMS_AOL.HeartbeatEnabled  "Inventory (PC)/Alert On LAN Settings/Heartbeat Enabled"
PC_INV.UMS_AOL.Enabled  "Inventory (PC)/Alert On LAN Settings/Hardware Enabled"
PC_INV.UMS_AOL.TimeoutInterval  "Inventory (PC)/Alert On LAN Settings/Watchdog Interval"
PC_INV.UMS_AOL.Enabled9  "Inventory (PC)/Alert On LAN Settings/Watchdog Enabled"
PC_INV.ASF_CFG.Name  "Inventory (PC)/Alert Standard Format Settings/Name"
PC_INV.ASF_CFG.DestinationType  "Inventory (PC)/Alert Standard Format Settings/Destination Type"
PC_INV.ASF_CFG.AlertDestinationAddress  "Inventory (PC)/Alert Standard Format Settings/Alert Destination Address"
PC_INV.ASF_CFG.MessageFormat  "Inventory (PC)/Alert Standard Format Settings/Message Format"
...
```

List dynamic group criteria (report format)

The following command lists identifiers which can be used to define dynamic groups or for listing inventory values:

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd server listdynamicgroupcriteria -r
```

The command results in the following output:

```
Database ...........................Inventory (PC)
Table ..............................Alert On LAN Settings
Column .............................Retransmission Interval
Identifer ..........................."PC_INV.UMS_AOL.RetryInterval"
Data Type...........................Integer
Multiple Rows per Entity Supported..false
{=,!=} Operators Supported..........true
{<,<=,>,>=} Operators Supported.....true

Database ...........................Inventory (PC)
Table ..............................Alert On LAN Settings
Column .............................Polling Interval
Identifer ..........................."PC_INV.UMS_AOL.EventPollInterval"
Data Type...........................Integer
Multiple Rows per Entity Supported..false
{=,!=} Operators Supported..........true
{<,<=,>,>=} Operators Supported.....true

Database ...........................Inventory (PC)
Table ..............................Alert On LAN Settings
Column .............................Heartbeat Interval
Identifer ..........................."PC_INV.UMS_AOL.HeartbeatInterval"
Data Type...........................Integer
Multiple Rows per Entity Supported..false
{=,!=} Operators Supported..........true
{<,<=,>,>=} Operators Supported.....true
...
```

List dynamic group criteria (terse format)

The following command lists identifiers which can be used to define dynamic groups or for listing inventory values:

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd server listdynamicgroupcriteria
```

The command results in the following output:

```
PC_INV.UMS_AOL.RetryInterval
PC_INV.UMS_AOL.EventPollInterval
PC_INV.UMS_AOL.HeartbeatInterval
PC_INV.UMS_AOL.Version
PC_INV.UMS_AOL.AlertDestinationAddress
PC_INV.UMS_AOL.EventAutoClearEnabled
PC_INV.UMS_AOL.HeartbeatEnabled
PC_INV.UMS_AOL.Enabled
PC_INV.UMS_AOL.TimeoutInterval
PC_INV.UMS_AOL.Enabled9
PC_INV.ASF_CFG.Name
PC_INV.ASF_CFG.DestinationType
PC_INV.ASF_CFG.AlertDestinationAddress
PC_INV.ASF_CFG.MessageFormat
...
```

Create a dynamic group

The following command creates a dynamic group named Under18GB which includes managed objects having disks smaller than 18 million kilobytes in size:

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd server createdynamicgroup
    -f Under18GB PC_INV.TWG_DISK.DISK_TOTAL_SIZE_KB "<=" 18000000
```

Create a dynamic group using OR and AND

The following command creates a dynamic group named MyGroup which includes managed objects which either have disks smaller than 18 million KB in size or have both a disk smaller than 36 million KB and less than 508356 KB installed memory:

```
dircli server/createdynamicgroup -f MyGroup4
    PC_INV.TWG_DISK.DISK_TOTAL_SIZE_KB "<" 18000000 OR
    (PC_INV.TWG_DISK.DISK_TOTAL_SIZE_KB "<" 36000000 AND
    PC_INV.TWG_INSTALLED_MEMORY.PHYSICAL_MEMORY_KB "<" 508356)
```

Create a dynamic group using ALL

The following command creates a dynamic group named Win2003-5 which includes managed objects for which the same operating system instance is named WIN2003 and is version 5:

```
dircli server/CreateDynamicGroup NC-Grp
    ((PC_INV.TWG_OPERATING_SYSTEM.OP_SYS_NAME = 'WIN2003') ALL
    (PC_INV.TWG_OPERATING_SYSTEM.OP_SYS_MAJ_VER = 5))
```

Create a dynamic group using EACH

The following command creates a dynamic group named Win2003andV5 which includes managed objects with one or more installed operating systems, one of which is named WIN2003 and one of which is version 5:

```
dircli server/CreateDynamicGroup Win2003andV5
    ((PC_INV.TWG_OPERATING_SYSTEM.OP_SYS_NAME = 'WIN2003') EACH
    (PC_INV.TWG_OPERATING_SYSTEM.OP_SYS_MAJ_VER = 5))
```

SNMP-device bundle

This topic describes SNMP-device bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. SNMP-device bundle commands facilitate management and configuration of managed SNMP devices.

General syntax (dircli)

```
►►──dircli──snmp──┬─/─┬──┤ Command-argument syntax ├──────────────────────►◄
                  └───┘
```

General syntax (dircmd)

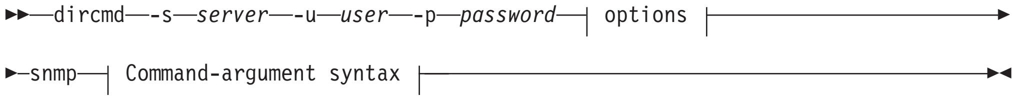

-s *server*
Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to “IBM Director dircmd usage” on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

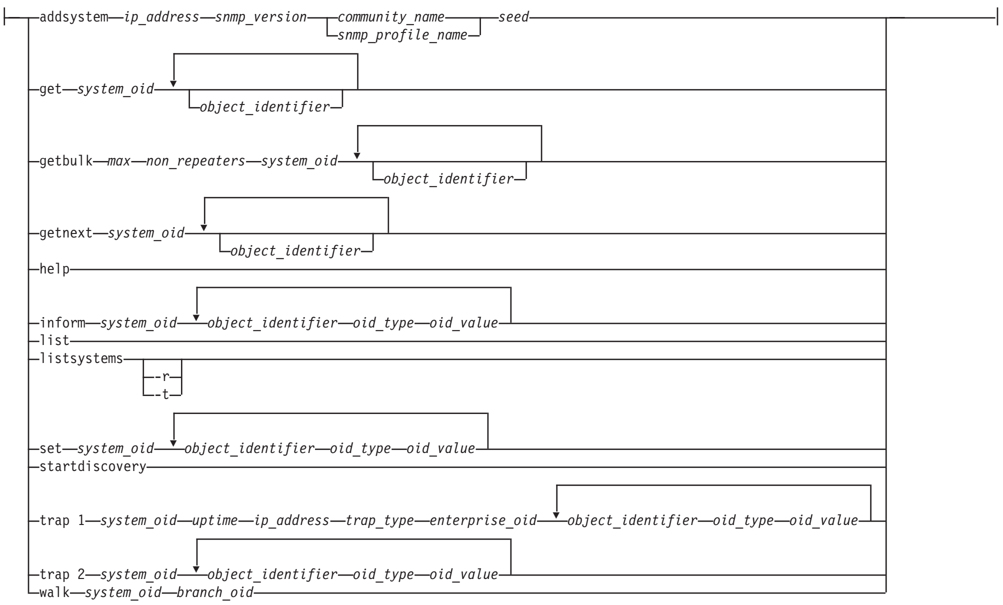
addsystem ip_address snmp_version community_name seed
snmp_profile_name
get system_oid
object_identifier
getbulk max non_repeaters system_oid
object_identifier
getnext system_oid
object_identifier
help
inform system_oid object_identifier oid_type oid_value
list
listsystems
-r
-t
set system_oid object_identifier oid_type oid_value
startdiscovery
trap 1 system_oid uptime ip_address trap_type enterprise_oid object_identifier oid_type oid_value
trap 2 system_oid object_identifier oid_type oid_value
walk system_oid branch_oid

addsystem

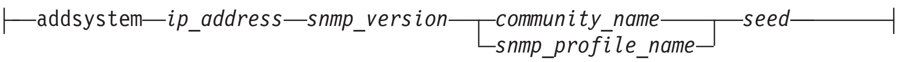

Creates an SNMP device on the management server. This is equivalent to right-clicking the Group Contents pane in IBM Director Console and then clicking **New → SNMP Devices**.

get

```
                   ┌──────────────────────┐
                   ▼                      │
├──get──system_oid───┬───────────────────┬─┴────────────────────────┤
                     └─object_identifier─┘
```

Performs an SNMP Get request against the SNMP device.

getbulk

```
                                                ┌──────────────────────┐
                                                ▼                      │
├──getbulk──max──non_repeaters──system_oid────────┬───────────────────┬─┴───┤
                                                  └─object_identifier─┘
```

Performs an SNMP Get Bulk request against the SNMP device.

getnext

```
                       ┌──────────────────────┐
                       ▼                      │
├──getnext──system_oid───┬───────────────────┬─┴────────────────────┤
                         └─object_identifier─┘
```

Performs an SNMP Get Next request against the SNMP device.

help

```
├──help──────────────────────────────────────────────────────────────┤
```

Displays general help for the bundle usage.

inform

```
                        ┌──────────────────────────────────────────┐
                        ▼                                          │
├──inform──system_oid─────object_identifier──oid_type──oid_value──┴───┤
```

Performs an SNMP Inform request against the SNMP device.

list

```
├──list──────────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

listsystems

```
├──listsystems──┬────┬───────────────────────────────────────────────┤
                ├─-r─┤
                └─-t─┘
```

Lists all SNMP devices in one of three formats:

Format	Information listed
report	system name, object ID, state, IP address, host name, MAC address, MIB2 system name, MIB2 system contact, MIB2 system location, MIB2 system object ID, and MIB2 system uptime
terse	object ID
neither (default)	system name and object ID

set

```
                      ┌──────────────────────────────────────────┐
                      ▼                                          │
├──set──system_oid────object_identifier──oid_type──oid_value──┴────────────┤
```

Performs an SNMP Set request against the SNMP device.

startdiscovery

```
├──startdiscovery──────────────────────────────────────────────────────────┤
```

Discovers all SNMP devices.

trap 1

```
├──trap 1──system_oid──uptime──ip_address──trap_type──enterprise_oid──────►

   ┌───────────────────────────────────────────┐
   ▼                                           │
►────object_identifier──oid_type──oid_value──┴─────────────────────────────┤
```

Sends an SNMPv1 trap to the SNMP device.

trap 2

```
                          ┌──────────────────────────────────────────┐
                          ▼                                          │
├──trap 2──system_oid────object_identifier──oid_type──oid_value──┴────────┤
```

Sends an SNMPv2 trap to the SNMP device.

walk

```
├──walk──system_oid──branch_oid────────────────────────────────────────────┤
```

Performs a walk on a branch of the MIB tree of the SNMP device.

Options and operands

ip_address
The IP address of the SNMP device or trap destination.

snmp_version
The version of SNMP to use. Valid values are 1, 2, and 3.

community_name
The SNMPv1 or SNMPv2 community name of the SNMP device.

snmp_profile_name
The SNMPv3 profile name of the SNMP device.

seed
A boolean switch indicating whether or not the SNMP device will be a seed for SNMP discovery. When true, the device will be a seed.

system_oid
The unique ID of the managed object.

object_identifier
An object identifier. Examples include `1.3.6.1.2.1.1.4.0` for sysContact and `1.3.6.1.2.1.1.1.0` for sysDescr.

max
The maximum number of get-next attempts to make when retrieving remaining objects.

non_repeaters
The number of objects that can be retrieved with a simple get-next operation.

oid_type
The type of the object identifier, which is one of the following values:

```
bits
counter
counter64
gauge
integer
ipaddress
nsapaddress
octets
oid
opaque
timeticks
unsigned32
```

oid_value
The value of the object identifier.

-r | -report
Enables a detailed (report) listing format. Refer to the command description for specific information about what is listed when applied to that command.

-t | -terse
Enables terse listing format. Refer to command description for specific information about what is listed when applied to that command.

uptime
The system uptime of the trap sender.

trap_type
The type of the trap being sent, which is one of the following values:

- `0` = coldStart
- `1` = warmStart
- `2` = linkDown
- `3` = linkUp
- `4` = authenticationFailure
- `5` = egpNeighborLoss
- `6` = enterprise-specific trap identified by an enterprise ID and a specific trap number chosen by the enterprise that defined the trap.

enterprise_oid
The enterprise object ID of the trap.

branch_oid
The object identifier of the branch. For example, you might use `1.3.6.1.2.1.1` to walk through all of the items in the system subtree.

Examples

Discover SNMP devices
In the following example, an IBM Director superuser connects to the management server with the host name `IDWorld`, using the user ID `InfoDeveloper` and the password `passw0rd`. When the user invokes the **startdiscovery** function, the following command discovers SNMP devices.

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd snmp startdiscovery
```

List discovered SNMP devices
In this example, the user invokes the **listsystems** function. This command produces a list of all discovered SNMP devices.

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd snmp listsystems
```

Display information about an SNMP device
In this example, the user invokes the **get** function to request an SNMP sysDescr (`1.3.6.1.2.1.1.1.0`) request from object `21B`.

```
dircmd -s IDWorld -u InfoDeveloper -p passw0rd snmp get 21B 1.3.6.1.2.1.1.1.0
```

The output from this command will be similar to the following sample output:

```
1.34.6.1.2.1.1.0 ,octets. = Hardware: x86 Family 15 model 1 Stepping 1 AT/AT
COMPATIBLE - Software: Windows 2000 Version 5.0 (Build 2195 Multiprocessor Free)
```

User-administration bundle

This topic describes user administration bundle commands available through either IBM Director command-line interface: **dircli** or **dircmd**. User-administration bundle commands facilitate reporting and modifying user access privileges.

General syntax (dircli)

```
►►─dircli─user──┬─/─┬──┤ Command-argument syntax ├────────────────────►◄
                └───┘
```

General syntax (dircmd)

```
►►─dircmd──-s─server──-u─user──-p─password─┤ options ├─────────────────►

►─user──┤ Command-argument syntax ├─────────────────────────────────────►◄
```

-s *server*
Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
Specifies the password for the superuser on the management server on which IBM Director Server is installed.

Note:

- The **dircli** keyword is required.
- The **dircmd** command-line interface is only available if IBM Director Server is installed on an IBM xSeries server.
- Refer to "IBM Director dircmd usage" on page 641 for complete information about **dircmd** options.

Command-argument syntax

The following syntax diagram illustrates all the commands available in the bundle. Details for each command follow in this section.

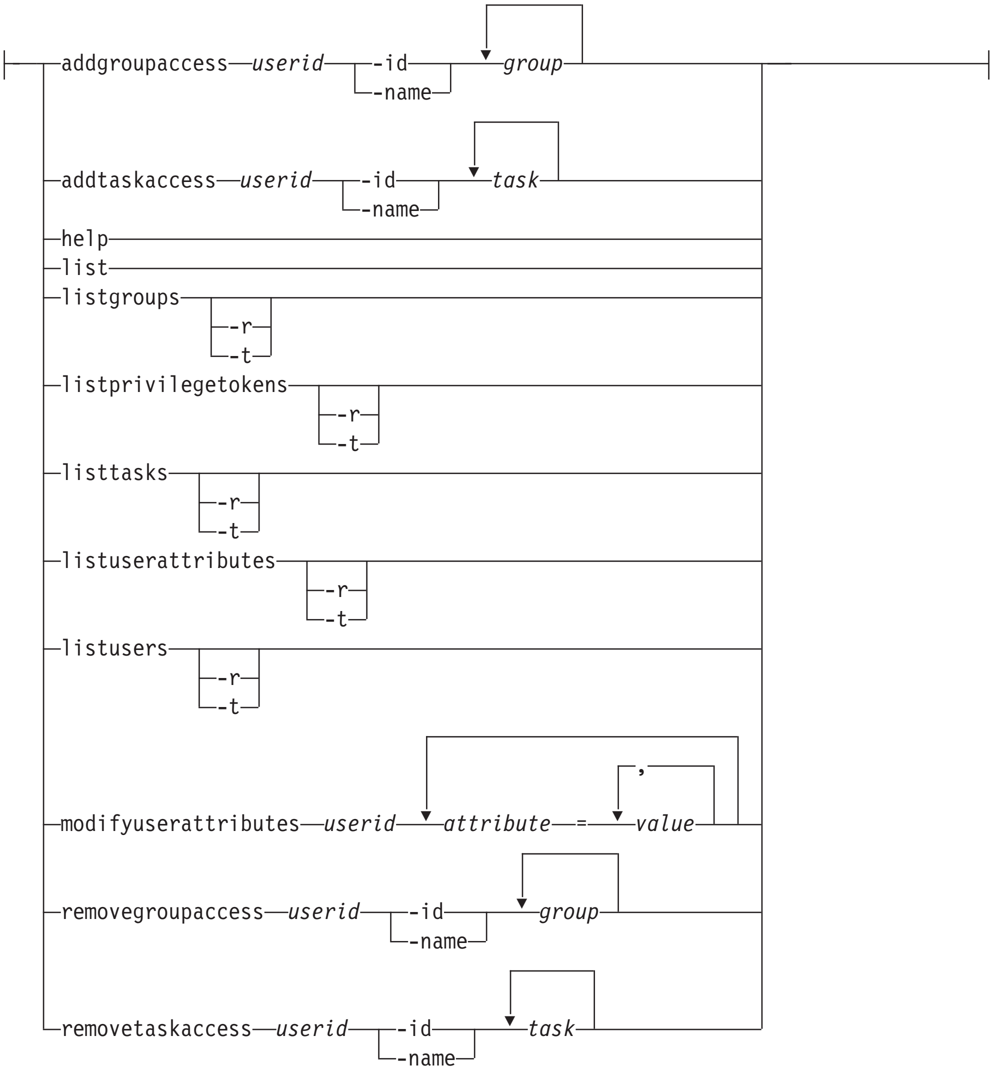

addgroupaccess

```
                                        ┌──────────┐
                                        ▼          │
├──addgroupaccess──userid──┬─-id────┬────group───┴──────────────────────┤
                           └─-name──┘
```

Adds access to the specified groups for the user.

addtaskaccess

```
                                      ┌───────┐
                                      ▼       │
├──addtaskaccess──userid──┬─-id────┬────task──┴──────────────────────────┤
                          └─-name──┘
```

Adds access to the specified tasks for the user.

help

```
├──help───────────────────────────────────────────────────────────────────┤
```

Displays general help for the bundle usage.

list

```
├──list───────────────────────────────────────────────────────────────────┤
```

Lists the commands available in the bundle.

listgroups

```
├──listgroups──┬──────┬───────────────────────────────────────────────────┤
               ├─-r───┤
               └─-t───┘
```

Lists groups in one of two formats:

Format	Information listed
report (default)	group name, object ID, and type
terse	group object ID

listprivilegetokens

```
├──listprivilegetokens──┬──────┬──────────────────────────────────────────┤
                        ├─-r───┤
                        └─-t───┘
```

Lists available privilege tokens in one of two formats:

Format	Information listed
report (default)	token ID and descriptive name
terse	token ID

listtasks

```
├──listtasks──┬────┬───────────────────────────────────────────────────┤
              ├─-r─┤
              └─-t─┘
```

Lists tasks in one of three formats:

Format	Information listed
report	task name, task object ID, and task ID
terse	task object ID
neither (default)	task name and task object ID

listuserattributes

```
├──listuserattributes──┬────┬──────────────────────────────────────────┤
                       ├─-r─┤
                       └─-t─┘
```

Lists the available user attributes in one of three formats:

Format	Information listed
report	attribute name, type, and expected value
terse	attribute name
neither (default)	attribute name and type

listusers

```
├──listusers──┬────┬───────────────────────────────────────────────────┤
              ├─-r─┤
              └─-t─┘
```

Lists users in one of three formats:

Format	Information listed
report	user name, superuser status, full name, description, e-mail, pager, locale, group access, task access, and privileges
terse	user name
neither (default)	user name, group access, task access, and privileges

modifyuserattributes

```
                                       ┌──────────────────────────┐
                                       │               ┌─,───┐    │
                                       ▼               ▼     │    │
├──modifyuserattributes──userid────attribute──=────value─┴────┴──────┤
```

Updates the specified user attributes.

removegroupaccess

```
                                             ┌──────┐
                                             ▼      │
├──removegroupaccess──userid──┬─-id───┬──────group──┴──────────────────────────┤
                              └─-name─┘
```

Denies group access to the user for the specified groups.

removetaskaccess

```
                                          ┌─────┐
                                          ▼     │
├──removetaskaccess──userid──┬─-id───┬─────task─┴──────────────────────────────┤
                             └─-name─┘
```

Denies task access to the user for the specified tasks.

Options and operands

userid
An IBM Director user ID that is authorized to access the managed object.

-id
Specifies that tasks or groups are identified by ID for the command.

-name
Specifies that tasks or groups are identified by name for the command.

group
The group name or ID.

task
The task name or ID.

-r | -report
Enables a detailed (report) listing format. Refer to the command description for specific information about what is listed when applied to that command.

-t | -terse
Enables terse listing format. Refer to command description for specific information about what is listed when applied to that command.

Examples

List IBM Director users

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd user listusers
```

The command results in the following output:

```
$$DEFAULT$$ group_access_limited=true task_access_limited=true group_access_readonly=true

SysAdmin group_access_limited=true task_access_limited=true group_access_readonly=true

IDWORLD\SysAdmin privset={TWGMSCS.ModifyCluster, engine.auditAdmin, engine.databaseadmin,
engine.discoverreq, engine.discoveryprops, engine.encryptAdmin, engine.modifyCatAccess,
engine.modifyLicenses, engine.modifyMOs, engine.powerdownMOs, engine.restartMOs,
engine.secureclients, engine.serverFileAccess, engine.serverprops, engine.shutdownMOs,
engine.useradmin, engine.wakeupMOs} group_access_limited=false task_access_limited=false
group_access_readonly=false
```

List IBM Director users (report format)

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd user listusers -r
```

The command results in the following output:

```
NAME..........IDWORLD\Administrator
SUPERUSER.....true
FULLNAME......
DESCRIPTION...Built-in account for administering the computer/domain
EMAIL.........
PAGER.........
LOCALE........enUS
GROUP ACCESS..ALL GROUPS
TASK ACCESS...ALL TASKS
PRIVILEGES....{ Allow modification of cluster settings, Allow auditing
administration, Allow database configuration, Allow discovery requests, Allow
access to discovery preferences, Allow encryption administration, Allow changes
to category access, Allow product license manipulation, Allow system create/
modify/delete operations, Allow power down of systems, Allow restart of systems,
Allow secure/unsecure actions on agents, Allow access to server file system,
Allow access to server preferences, Allow shutdown of systems, Allow user
account administration, Allow power on of systems }

NAME..........Administrator
SUPERUSER.....true
FULLNAME......
DESCRIPTION...Built-in account for administering the computer/domain
EMAIL.........
PAGER.........
LOCALE........enUS
GROUP ACCESS..(READONLY) {  }
TASK ACCESS...{  }

NAME..........$$DEFAULT$$
SUPERUSER.....false
FULLNAME......
DESCRIPTION...
EMAIL.........
PAGER.........
LOCALE........enUS
GROUP ACCESS..(READONLY) {  }
TASK ACCESS...{  }
```

List IBM Director users (terse format)

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd user listusers -t
```

The command results in the following output:

```
$$DEFAULT$$
Administrator
IDWORLD\Administrator
```

List user attributes

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd user listuserattributes
```

The command results in the following output:

```
STRING::email
STRING::pager
STRING ARRAY::privset
BOOLEAN::group_access_limited
BOOLEAN::group_access_readonly
BOOLEAN::task_access_limited
LONG ARRAY::taskaccesslist
LONG ARRAY::filteraccesslist
```

List groups in IBM Director

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd user listgroups
```

The command results in the following output:

```
1CC     Systems with Asset ID
1CD     Systems with SNMP Agent
1CF     Racks with Members
1D0     Systems with CIM
1D3     Scalable Systems and Members
1DB     Clusters and Cluster Members
1DD     Systems with ASF
```

```
1DE     Systems with ASF Secure Remote Management
1EB     All Managed Objects
204     Hardware Status Critical
205     Hardware Status Warning
206     Hardware Status Information
207     Platforms and Platform Members
208     Chassis and Chassis Members
20D     Systems with Windows 2000
210     Level 2: IBM Director Agents
22F     Systems with Windows XP
234     Systems with Linux
```

List privilege tokens

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd user listprivilegetokens
```

The command results in the following output:

```
TWGMSCS.ModifyCluster   Allow modification of cluster settings
engine.auditAdmin       Allow auditing administration
engine.databaseadmin    Allow database configuration
engine.discoverreq      Allow discovery requests
engine.discoveryprops   Allow access to discovery preferences
engine.encryptAdmin     Allow encryption administration
engine.modifyCatAccess  Allow changes to category access
engine.modifyLicenses   Allow product license manipulation
engine.modifyMOs        Allow system create/modify/delete operations
engine.powerdownMOs     Allow power down of systems
engine.restartMOs       Allow restart of systems
engine.secureclients    Allow secure/unsecure actions on agents
engine.serverFileAccess Allow access to server file system
engine.serverprops      Allow access to server preferences
engine.shutdownMOs      Allow shutdown of systems
engine.useradmin        Allow user account administration
engine.wakeupMOs        Allow power on of systems
```

List tasks

```
dircmd -s IDworld -u InfoDeveloper -p passw0rd user listtasks
```

The command results in the following output:

```
CA      SNMP Browser
D0      Hardware Status
D6      Scheduler
DA      Remote Session
DB
DD      Remote Control
E0      Process Management
108     Rack Manager
112     MPCLI
11F     Management Processor Assistant
153     Inventory
15A     File Transfer
162     Event Log
164     Event Action Plan Builder
165     Event Action Plans
169     System Accounts
17A     Network Configuration
183     Configure SNMP Agent
185     CIM Browser
187     Configure Alert Standard Format
18A     System Identification
190     BladeCenter Assistant
19A     Asset ID
19E     Active Console Viewer
1A3     Set Presence Check Interval
1A5     Microsoft Cluster Browser
1A8     Resource Monitors
1B3
```

```
1BA     Software Distribution
1BF     Software Health Check
1D5     External application launch
213     Capacity Manager
```

IBM Director dircmd usage

This topic describes the syntax and options specific to **dircmd** but not associated with any command bundles.

Note: Use of **dircli** instead of **dircmd** is strongly recommended for four reasons:

- **dircli** has enhanced security when compared with **dircmd**
- **dircli** is a high-performance client and runs approximately 10-100 times faster than **dircmd**, depending on configuration
- **dircli** is available on all platforms, while **dircmd** is limited to IBM xSeries servers
- support for the **dircmd** command-line interface is unlikely to continue beyond version 5.10 of IBM Director.

Although **dircmd**-bundle commands can be executed with **dircli**, the reverse is not true. Commands specific to **dircli** cannot be executed with **dircmd**.

Syntax

```
►►──dircmd──-s──server──-u──user──-p──password────────────────────────────►

►─┬──────────────────────────────────────────────┬────────────────────────►
  └─┬─────────────────┬─┬──────────────────────────┬─┘
    └─-k──data_link───┘ └─-o──data_link_params─────┘

►──┬─-h──────────────────────────────────────────────────┬───────────────►◄
   ├─-?──────────────────────────────────────────────────┤
   └─┬──────────────┬──bundle_id──┤ Command-argument syntax ├─┘
     ├─-f──filename─┤
     └─-r───────────┘
```

Options and operands

-s *server*
Specifies the DNS resolvable host name or TCP/IP address for establishing a connection to the management server on which IBM Director Server is installed.

-u *user*
Specifies the user name of a superuser on the management server on which IBM Director Server is installed.

-p *password*
Specifies the password for the superuser on the management server on which IBM Director Server is installed.

-k *data_link*
Overrides the default TCP/IP data-link connection class with the class specified by *data_link*. This advanced option should be used with caution. When the default option is overridden, data-link parameter specifications could also be needed to configure the link correctly. The default link class is

`com.tivoli.twg.libs.TWGTCPIPLink`. In addition the secure link class "com.tivoli.twg.libs.TWGSSLLink" is available for creating encrypted data connections to the server.

-o *data_link_params*
Overrides the default data link parameters with the parameters specified by *data_link_params*. This advanced option should be used with caution. The default TCP/IP data-link parameter is `2034`, which sets the socket port for the `com.tivoli.twg.libs.TWGTCPIPLink` network link.

-h Provides general but detailed information about the command.

Note: This help is derived at the server and not locally; therefore, a server connection must be present.

-? Provides general but detailed information about the command.

Note: This help is derived at the server and not locally; therefore, a server connection must be present.

-f *filename*
Directs the command to receive operand data from the file specified by *filename*. Typically the file contains a list of object IDs.

-r Directs the command to receive operand data from an input pipe, allowing output from another command to be used as input for the current command. For an example of this option, refer to "Examples" on page 610.

bundle_id
Specifies the command bundle being used. The following table lists valid values.

Bundle ID	Bundle
bladecenterchassis	"BladeCenter-chassis bundle" on page 600
cli	"Command-line interface bundle" on page 604
chassis	"Chassis bundle" on page 602
event	"Event-management bundle" on page 607
native	"Level 2 managed-system bundle" on page 611
mpa	"Management Processor Assistant bundle" on page 614
procmon	"Process-monitor bundle" on page 617
monitor	"Resource-monitor bundle" on page 619
scheduler	"Scheduler bundle" on page 622
server	"Server-management bundle" on page 624
snmp	"SNMP-device bundle" on page 629
user	"User-administration bundle" on page 634

Exit Codes

The **dircmd** command returns the following exit codes:

Exit code	Meaning	Value
OK	Successful completion	0
USAGE	Errors due to missing or improper arguments	1

Exit code	Meaning	Value
NOT_FOUND	Command or bundle not found	2
SECURITY_FAILURE	Security failure due to unauthorized client	3
COMMAND_EXCEPTION	Command implementation caused and exception	4
FAIL	General request of action failed	5

Note: Additional exit codes with unique positive values are defined by some bundles.

Deprecated dircmd commands

This topic lists commands formerly available through the **dircmd** command-line interface that are now deprecated. The equivalent functionality is now implemented through new commands in the **dircli** command-line interface.

Deprecated dircmd bundle/command	Equivalent command in dircli
bladecenterconfiguration/xmlfile	"mkcmprof" on page 575
server/accessobjects	"accessmo" on page 532
server/addtostaticgroup	"chgp" on page 541
server/createstaticgroup	"mkgp" on page 577
server/deletegroups	"rmgp" on page 587
server/deleteobjects	"rmmo" on page 589
server/discoverall	"discover" on page 546
server/listgroupattributes	"lsgp" on page 556
server/listgroupmembers	"lsmo" on page 561
server/listgroups	"lsgp" on page 556
server/listgroupsbyattribute	"lsgp" on page 556
server/listnoninteractivetasks	"lstask" on page 569
server/listobjectattributes	"lsmo" on page 561
server/listobjects	"lsmo" on page 561
server/listobjectsbyattribute	"lsmo" on page 561
server/listtaskactivationstatus	"lstask" on page 569
server/pingobjects	"pingmo" on page 583
server/removefromstaticgroup	"chgp" on page 541
server/renameobject	"chmo" on page 543
server/runtask	"runtask" on page 595

Miscellaneous IBM Director commands

This topic lists miscellaneous IBM Director commands.

Note: Not all commands are supported on all operating systems.

Configuration/management commands

Command	Description
“cimsubscribe” on page 646	This command is used for subscription management.
“cfgsecurity” on page 645	This command opens a security-configuration dialog on level-1 and level-2 managed systems running Linux.
“getfru” on page 653	This command retrieves field-replaceable unit (FRU) information from a server.
“Genevent” on page 651	This command sends a custom event from the managed system to IBM Director Server.
“twgend” on page 653	This command stops IBM Director Agent on i5/OS operating systems.
“twgipccf” on page 654	This command opens the Network driver configuration window for IBM Director Agent.
“twgreset” on page 655	This command returns IBM Director Server to its installation default values and clears all tables in the database.
“twgrestore” on page 655	This command restores the IBM Director persistent data.
“twgsave” on page 656	This command saves the IBM Director persistent data.
“twgstart” on page 657	This command starts IBM Director Server and IBM Director Agent on i5/OS and Linux operating systems.
“twgstat” on page 658	This command returns the active state of IBM Director Server on AIX, i5/OS and Linux operating systems.
“twgstop” on page 659	This command stops IBM Director Agent on i5/OS and Linux operating systems.

Database commands

Command	Description
“cfgdb”	This command starts the graphical database-configuration tool for IBM Director.
“cfgdbcmd” on page 645	This command configures the IBM Director database using a database-configuration response file.
“dbpasswd” on page 650	This command sets the password for the IBM Director database.
“uncfgdb” on page 659	This command starts the graphical database-removal tool for IBM Director.
“uncfgdbcmd” on page 660	This command removes the IBM Director database configuration using a database-removal response file.

cfgdb

This topic provides information about the **cfgdb** command. This command starts the graphical database-configuration tool for IBM Director.

```
►►—cfgdb——————————————————————————————————————————————————————————————►◄
```

Parameters

This command takes no parameters.

Examples

To configure the IBM Director database: Type the following command to launch the database configuration tool:

```
cfgdb
```

cfgdbcmd

This topic provides information about the **cfgdbcmd** command. This command configures the IBM Director database using a database-configuration response file.

```
►►──cfgdbcmd──-rspfile──response_file──┬─────────┬──────────────────────────────────────►◄
                                       └─-silent─┘
```

Parameters

–rspfile *response_file*
: Specifies that database configuration information is located in a response file. The *response_file* is the fully qualified name of the database-configuration response file.

–silent
: Specifies that the database configuration occurs silently. Results are logged to the `cfgdbcmd.output` file in the `director\log` directory. If the database configuration is successfully completed, a return code of `0` is generated. If the database configuration fails to be successfully completed, a non-zero return code is generated.

Example

Configure the IBM Director database
: The following command configures the IBM Director database using the data stored in the `MyDatabaseConfig.txt` file:

```
cfgdbcmd -rspfile c:\temp\MyDatabaseConfig.txt
```

cfgsecurity

This topic provides information about the **cfgsecurity** command. This command opens a security-configuration dialog on level-1 and level-2 managed systems running Linux.

Encryption between IBM Director Server and IBM Director Agent or IBM Director Core Services is configured on both IBM Director Server (using IBM Director Console) and on the managed systems.

For Windows managed systems, encryption is configured by modifying the installation via **Control Panel** → **Add or Remove Programs**.

For Linux managed systems, encryption is configured using the **cfgsecurity** command.

```
►►──cfgsecurity──────────────────────────────────────────────────────────►◄
```

Note: Restart IBM Director Agent or IBM Director Core Services after issuing the **cfgsecurity** command

Parameters

This command takes no parameters.

Example

Open the security-configuration dialog on Linux
The following command opens a security-configuration dialog on a system with IBM Director Agent installed in `/opt/IBM/director/`:

```
/opt/IBM/director/bin/cfgsecurity
```

cimsubscribe

This topic provides information about the **cimsubscribe** command. This command is used for subscription management.

The **cimsubscribe** command, available on Windows and Linux, is located in cimom/bin where IBM Director Agent is installed. Used during installation to set up the default local subscriptions, the **cimsubscribe** command can be run either in interactive mode or in silent mode, so it may be scripted. The **cimsubscribe** command can be used to connect to a local or remote CIMOM and has parameters for authentication.

Other clients such as cimcli or wbemtest may also be used for subscription management.

```
►►──cimsubscribe──┬─-b──batch_filename──────────────────────────────────────┬──────►◄
                  ├─-cf──-fn──filter_name──-q──query──┤ Options ├───────────┤
                  ├─-ch──-hn──handler_name──-d──destination──┤ Options ├────┤
                  ├─-cs──-fn──filter_name──-hn──handler_name──┤ Options ├───┤
                  ├─-df──-fn──filter_name──┤ Options ├──────────────────────┤
                  ├─-dh──-hn──handler_name──┤ Options ├─────────────────────┤
                  ├─-ds──-sn──subscription_name──┤ Options ├────────────────┤
                  ├─-h──────────────────────────────────────────────────────┤
                  ├─-i──┬────────────────────────────────────────┬──────────┤
                  │     └─-outputfile──┬───────────────────┬─────┘          │
                  │                    └─output_filename───┘                │
                  ├─-lf─────────────────────────────────────────────────────┤
                  ├─-lh─────────────────────────────────────────────────────┤
                  ├─-ls─────────────────────────────────────────────────────┤
                  └─-?──────────────────────────────────────────────────────┘
```

Options:

```
   ┌─-s──┬──────┬─────────────────────────────────────┐
   │     └─no───┘                                     │
├──┼──────────────────────────────────────────────────┼──┬────────────────┬──────►
   └─-s yes──┬─-u──user_name──-p──password──────────┬─┘  └─-n──namespace──┘
             └─-c──certificate_file──-k──key_file───┘
```

```
►─┬──────────────┬────────────────────────────────────────────────────────────┤
  └─-l—location──┘
```

Parameters

-b *batch_filename*
Specifies that the command will read and execute command parameters in the specified batch file. A **cimsubscribe** batch file is a text file with one or more lines, each of which specifies parameters for a **cimsubscribe** command but not the **cimsubscribe** keyword itself.

-cf | -createfilter
Creates a new CIM_IndicationFilter.

-fn | -filtername *filter_name*
Specifies the name of a specified CIM_IndicationFilter instance.

-q | -query *query*
Specifies a Windows Management Instrumentation Query Language (WQL) query used to filter CIM_Indication instances.

-ch | -createhandler
Creates a new CIM_IndicationHandlerCIMXML.

-hn | -handlername *handler_name*
Specifies the name of a specific CIM_IndicationHandlerCIMXML instance.

-d | -destination *destination*
Specifies a destination uniform resource locator (URL) used to identify an End Consumer location. For example, `http://localhost:5992/CIMListener/syslog`

Note: The port number specified by *destination* should be appropriate to the selected communications protocol.

-cs | -createsubscription
Creates a new CIM_IndicationSubscription after checking to ensure the specified handler and filter exist.

-df | -deletefilter
Deletes the specified CIM_IndicationFilter if the specified filter exists and is not being used by an existing CIM_IndicationSubscription.

-dh | -deletehandler
Deletes the specified CIM_IndicationHandlerCIMXML if the specified handler exists and is not being used by an existing CIM_IndicationSubscription.

-ds | -deletesubscription
Deletes the specified CIM_IndicationSubscription if the specified subscription exists.

-sn *subscription_name*
Specifies a subscription name for deletion.

-h | -?
Displays command-line help information about the **cimsubscribe** command.

-i Specifies that the command will be executed in interactive mode; the user is prompted for values which are not specified on the command line.

Note: When deleting handlers, filters, or subscriptions in interactive mode, all objects (handlers, filters, or subscriptions) are displayed in a numbered list. Because this list can be long, it may be necessary to increase the buffer size to allow you to scroll through the list to identify the number of the object to delete.

-outputfile [*output_filename*]
Specifies that in interactive mode, the **cimsubscribe** commands corresponding to the user's input are saved to a text file which can be used later for batch processing.

-lf | -listfilters
Lists the current set of CIM_IndicationFilter.

-lh | -listhandlers
Lists the current set of CIM_IndicationHandlerCIMXML.

-ls | -listsubscriptions
Lists the current set of CIM_IndicationSubscription.

-s | -secure
Specifies whether the HTTP or HTTPS protocol is used for communications. When followed by `yes`, specifies that the command will connect securely to port 5989 (HTTPS). When followed by `no` or omitted, specifies that the command will connect using port 5988 (HTTP).

Note: HTTPS communications may decrease performance because of increased processing. The port number specified by *destination* should be appropriate to the selected communications protocol.

-u | -username *user_name*
Specifies a user name for server authentication.

-p | -password *password*
Specifies a password for server authentication.

-c | -certificate *certificate_file*
Specifies the full system path to the client x509 certificate. For example, `c:\test\client.cert`. Only valid if using secure communications via HTTPS.

-k | -key *key_file*
Specifies the full system path to the client private key file. For example, `c:\test\client.key`. Only valid if using secure communications via HTTPS.

-n | -namespace *namespace*
Specifies a namespace. If no namespace is specified, the default value is `root/ibmsd`.

-l | -location *location*
Specifies the fully-qualified host name and port of the system on which to modify subscriptions. For example, `remotehost.raleigh.ibm.com:5988`. A user name and password must be specified for non-localhost destinations. A fully qualified host name must be specified when creating a remote filter, handler, or subscription. If no host name is specified, the default value is `localhost:5988`.

Note: There is no default WBEM port for indications. The ports 5992-5998 are not reserved, so any of these ports is a good candidate for a default port. Although the system administrator can change the listener port, it is not recommended.

Examples

The following examples illustrate common uses for the **cimsubscribe** command.

Connect to a remote host over HTTP

The following example illustrates creation of a handler named SNMP on listener 9.44.169.107, port 5988 (HTTP). The user name is heather; the password is m1n2b3. The destination for the handler is http://localhost:6988/CIMListener/SnmpConsumer.

```
cimsubscribe -ch -hn "SNMP" -d "http://localhost:6988/CIMListener/SnmpConsumer"
-l 9.44.169.107:5988 -u heather -p m1n2b3
```

Connect to a remote host over HTTPS

The following examples illustrate creation of a handler named SNMP on listener 9.44.169.107, port 5989 (HTTPS). The destination for the handler is http://localhost:6988/CIMListener/SnmpConsumer. The first example uses username/password authentication; the second specifies a client certificate and key.

```
cimsubscribe -ch -hn "SNMP" -d "http://localhost:6988/CIMListener/SnmpConsumer"
-l 9.44.169.107:5989 -s yes -u heather -p m1n2b3
```

```
cimsubscribe -ch -hn "SNMP" -d "http://localhost:6988/CIMListener/SnmpConsumer"
-l 9.44.169.107:5989 -s yes -c client.cert -k client.key
```

Create filter

The following example illustrates creation of a filter named Sev2 on the default listener.

```
cimsubscribe -cf -fn Sev2 -q "SELECT * FROM IBM_AlertIndication where Severity = 2"
```

Get CIM data on another namespace

The following example illustrates creation of a handler named SNMP for a namespace (root/cimv2) other than the default namespace.

```
cimsubscribe -ch -hn "SNMP" -d "http://localhost:6988/CIMListener/SnmpConsumer"
-n root/cimv2
```

Use batch processing

The following example illustrates batch processing with the **cimsubscribe** command. The defaultHandlers.dat batch file is processed as though each line in the file was entered as parameters for the **cimsubscribe** command.

```
cimsubscribe -b defaultHandlers.dat
```

The defaultHandlers.dat batch file contains the following lines specifying handler-creation instructions:

```
-ch -hn "SNMP" -d "http://localhost:6988/CIMListener/SnmpConsumer"
-ch -hn "TEC" -d "http://localhost:6988/CIMListener/TivoliConsumer"
-ch -hn "Log" -d "http://localhost:6988/CIMListener/LogConsumer"
-ch -hn "Health" -d "http://localhost:6988/CIMListener/HealthConsumer"
-ch -hn "SMS" -d "http://localhost:6988/CIMListener/SMSConsumer"
-ch -hn "PopUp" -d "http://localhost:6988/CIMListener/PopupConsumer"
```

Each line is executed as if the **cimsubscribe** command were typed with the specified parameters.

Use interactive processing

The following example is a transcript of an interactive processing session for creating a new filter. User responses are indicated in boldface type. After the filter is created, the interactive session prompts for a new action. Note the interactive command does not explicitly indicate that the new filter was created, but re-displays the query criteria for the filter.

```
C:\Program Files\IBM\Director\cimom\bin>cimsubscribe -i
What system would you like to connect to?
1. localhost
2. remote host
1
What port would you like to connect to?
1. 5988 (HTTP)
2. 5989 (HTTPS)
3. Another port
1
Enter the namespace.
root/ibmsd
Do you want a secure connection?
1. Yes
2. No
2
Connecting to localhost:5988...
Interactive mode

What would you like to do?
1. Create a filter
2. Create a handler
3. Create a subscription (bind an existing filter and handler)
4. Delete an existing filter
5. Delete an existing handler
6. Delete an existing subscription
7. Exit
1

What is the name of this filter?
new_filter

What is the WQL query for this filter?
SELECT * from IBMPSG_ProcessorPFEvent where PerceivedSeverity = 2

Name new_filter
Query SELECT * from IBMPSG_ProcessorPFEvent where PerceivedSeverity = 2

What would you like to do?
1. Create a filter
2. Create a handler
3. Create a subscription (bind an existing filter and handler)
4. Delete an existing filter
5. Delete an existing handler
6. Delete an existing subscription
7. Exit
7
```

dbpasswd

This topic provides information about the **dbpasswd** command. This command sets the password for the IBM Director database.

```
►►—dbpasswd—-user—user_id—-pwd—password—-confirmpwd—password————————————►◄
```

Parameters

user_id
The administrator user for the IBM Director database.

password
The new password for the IBM Director database. The password must be specified twice to confirm that the password was entered correctly.

Examples

To set the password: The following code sets the password for user admin to sh1bb0leth.

```
dbpasswd -user admin -pwd sh1bb0leth -confirmpwd sh1bb0leth
```

Genevent

This topic provides information about the **Genevent** command. This command sends a custom event from the managed system to IBM Director Server.

```
►►──Genevent── /type:"──event_type──"── /text:"──event_description──"──────────►

►──┬─────────────────────────────────────────┬──┬──────────────────┬──────────►◄
   │         ┌─@EventServer─────────────────┐│  │        ┌─UNKNOWN──┐│
   └─ /dest:─┴─protocol──::──server_address─┘┘  └─ /sev:─┼─FATAL────┤┘
                                                         ├─CRITICAL─┤
                                                         ├─MINOR────┤
                                                         ├─WARNING──┤
                                                         └─HARMLESS─┘
```

Note: The **Genevent** keyword and parameters are case-sensitive on Linux.

Parameters

/type:
Specifies the type of event to send to the management server.

event_type
A dot-delimited string specifying an event type. Quote *event_type* if it contains a space.

/text:
Specifies a text description for the event.

event_description
A quoted string specifying descriptive text for the event.

/dest:
Specifies the destination management server to which the event will be sent.

protocol
Specifies the protocol for the management *server_address*, one of `TCPIP`, `NETBIOS`, or `IPX`.

If more than network interface is enabled for the same protocol, the interfaces are identified by suffixing a number to the second through *n*th interface. For example, the second TCPIP network card would be identified by specifying a *protocol* of `TCPIP1`. Use the **twgipccf** command to open the Network driver configuration window, which lists enabled network drivers on the managed system.

server_address
The address of the management server to which the event will be sent, in the format specified by *protocol*.

@EventServer
Using **@EventServer** as the destination causes **Genevent** to send the event to any management server that has discovered the managed system. This is the default behavior if no destination is specified.

/sev:

Specifies the severity level for the event, one of: FATAL, CRITICAL, MINOR, WARNING, HARMLESS, or UNKNOWN. If **/sev:** is omitted, the severity level is UNKNOWN.

Examples

The following example queries the management server at IP address 9.44.206.162 to see if it is online:

```
Genevent /type:"Test.IPC" /text:"Is the primary management server online?"
/dest:TCPIP::9.44.206.162
```

The following example sends an event to the management server IDWORLD.raleigh.ibm.com that will add your IBM Director Agent as a managed node.

```
Genevent /type:Director.Topology.Online /text:"System Added"
/dest:TCPIP::IDWORLD.raleigh.ibm.com
```

The following example sends a test event to the management server to verify the communications status:

```
Genevent /type:"Test.Event" /text:"Checking communication..."
/dest:TCPIP2::ZURICH.IBM.COM
```

Return codes

Code	Meaning
0	The event was sent successfully to the given destination.
0xFFFF0000L	**Genevent** failed; unable to send packet on network.
0xFFFF0001L	**Genevent** failed due to security issues.
0xFFFF0002L	**Genevent** failed due to timeout.
0xFFFF0003L	**Genevent** failed due to service failure.
0xFFFF0004L	**Genevent** failed due to encryption failure.
0xFFFF0005L	**Genevent** failed due to invalid destination.
0xFFFF0006L	**Genevent** failed due to service node problem.
0xFFFF0007L	**Genevent** failed due to creation of packet problem.
0xFFFF0008L	**Genevent** failed due to failure to queue problem.
0xFFFF1000L	Generic SendBuffer failure.
0xFFFF10xxL	SendBuffer failures, error not in sender’s SN; specific error indicated by xx.
0xFFFF11xxL	SendBuffer failures, error in sender’s SN; specific error indicated by xx.

Additional information

Genevent is installed as part of IBM Director Agent. **Genevent** is *not* available on managed systems with only IBM Director Core Services installed.

The **Genevent** command has numerous applications. Some uses include:

- agent-initiated discovery of the managed system by using a batch file with **Genevent** that sends an event to the management server with the managed system’s name and IP address

- as an action in an event action plan, **Genevent** can be used to send a message to another management server when the plan is triggered by an event

Note: When cloning systems with IBM Director Agent installed, do not start IBM Director Agent on the system to be cloned until after it has been cloned. If started before cloning, IBM Director Agent will calculate a unique ID for the system, which will be copied to any clones, resulting in multiple systems with the same ″unique″ ID and problems with IBM Director Server.

getfru

This topic provides information about the **getfru** command. This command retrieves field-replaceable unit (FRU) information from a server.

Note: The **getfru** requires FTP access to the IBM Support FTP site (or to your specified FTP server) through your firewall. For the **getfru** command to run successfully, the managed system must have firewall access through a standard FTP port.

The **getfru** command is located in the `/CIMOM/bin` directory under the directory in which IBM Director is installed.

```
►►──getfru──┬──────────────────┬──┬─────────────────┬──────────────────────────►◄
            └─-s──ftp_server──┘  └─-d──directory──┘
```

Parameters

-s *ftp_server*
: Specifies the FTP address of the network server from which to retrieve FRU data files. If you do not specify an address, the command uses a default value of `ftp.software.ibm.com`.

-d *directory*
: Specifies the directory where the FRU data files are stored. If you do not specify a directory, the command uses a default value of `pc/pccbbs/bp_server`.

Examples

Get FRU information (Linux)

The following command retrieves FRU data from `ftp.shibbolethsystems.com/public/frudata`:

```
./opt/IBM/director/CIMOM/bin/getfru -s ftp.shibbolethsystems.com
-d /public/frudata
```

Get FRU information (Windows)

The following command retrieves FRU data from `ftp.shibbolethsystems.com/public/frudata`:

```
c:\Program Files\IBM\Director\cimom\bin\getfru -s ftp.shibbolethsystems.com
-d /public/frudata
```

twgend

This topic provides information about the **twgend** command. This command stops IBM Director Agent on i5/OS operating systems.

```
►►──twgend─────────────────────────────────────────────────────────────────────►◄
```

To stop IBM Director Agent, type one of the following commands and press Enter:

For i5/OS	`/qibm/userdata/director/bin/twgend`
For Linux (32-bit operating systems or AMD64)	`/opt/IBM/bin/director/twgstop`
For Linux (Intel Itanium, IBM iSeries, or IBM pSeries)	`/opt/ibm/bin/director/twgstop`
For NetWare	`unload twgipc`
For Windows	`net stop twgipc`

Note: Stopping IBM Director Agent on Linux, NetWare and Windows does not require use of the **twgend** command; the information is listed here for completeness.

Parameters

This command takes no parameters.

twgipccf

This topic provides information about the **twgipccf** command. This command opens the Network driver configuration window for IBM Director Agent.

```
►►—twgipccf——————————————————————————————————————————————————►◄
```

Parameters

This command takes no parameters.

Example

Open the Network driver configuration window

```
twgipccf
```

The Network driver configuration window opens.

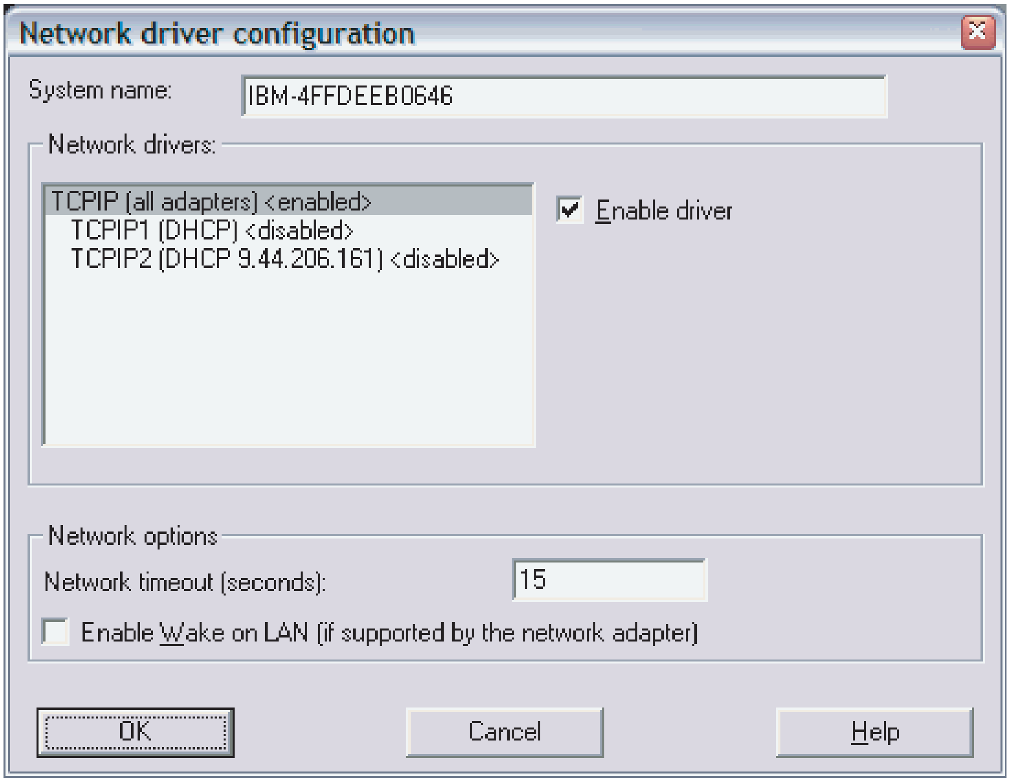

Network driver configuration

System name: IBM-4FFDEEB0646

Network drivers:

TCPIP (all adapters) <enabled>
  TCPIP1 (DHCP) <disabled>
  TCPIP2 (DHCP 9.44.206.161) <disabled>

☑ Enable driver

Network options

Network timeout (seconds): 15

□ Enable Wake on LAN (if supported by the network adapter)

twgreset

This topic provides information about the **twgreset** command. This command returns IBM Director Server to its installation default values and clears all tables in the database.

CAUTION:
twgreset changes the configuration of IBM Director Server and cannot be un-done except by manually re-configuring IBM Director Server.

```
►►──twgreset──┬────┬──────────────────────────────────────────────────────►◄
              └─-i─┘
```

Parameters

-i Specifies that **twgreset** will erase the system's unique identification files. This may be used after a restore to make sure that only the data from the saved directory will be in the IBM Director system.

Example

Reset the IBM Director Server configuration to the installation default values
The following command resets the IBM Director configuration and erases the system's unique identification files:

```
twgreset -i
```

twgrestore

This topic provides information about the **twgrestore** command. This command restores the IBM Director persistent data.

```
►►──twgrestore──directory──┬────┬──────────────────────────────────────────►◄
                           └─-t─┘
```

Parameters

directory
Specifies the directory from which the persistent data is restored. The data that you restore must be from the same version of IBM Director Server or IBM Director Agent that is installed.

-t Specifies that neither the system unique identifier or system name are restored.

Note: This command must be run locally. Before you run this command, stop all IBM Director processes that are running on the system.

Examples

Restore all IBM Director persistent data
The following command restores all IBM Director persistent data:

```
twgrestore /opt/IBM/director.save.1
```

Exclude the unique system identifier and name
The following command restores all IBM Director persistent data except the unique system identifier and name:

```
twgrestore restore /opt/IBM/director.save.1 -t
```

Return codes

The **twgrestore** command returns the following codes.

Code	Meaning
0	The persistent data was successfully restored.
1	An invalid parameter was issued.
2	An IBM Director service is still running.
3	The tar command failed.
15	An inaccessible directory was specified.

twgsave

This topic provides information about the **twgsave** command. This command saves the IBM Director persistent data.

The **twgsave** command saves the IBM Director persistent data to the *installation*.save.*n* directory. The variable *installation* is the directory where IBM Director Server or IBM Director Agent is installed. The variable *n* is an integer starting at 1 that is incremented each time the **twgsave** command is run. For example, if IBM Director is installed in the default directory, the directory that is used is `/opt/IBM/director.save.`*n*. The persistent data includes configuration and working IBM Director data. It does not include database information. This command must be run locally from the system on which you want to save data. Before you run this command, stop all IBM Director processes that are running on the system.

```
►►──twgsave──┬────┬──────────────────────────────────────────────►◄
             └─-s─┘
```

Parameters

-s Specifies that the software packages that are used by the Software Distribution task are not saved

Examples

Save the IBM Director persistent data

The following command saves all IBM Director persistent data:

```
twgsave
```

Exclude software packages

The following command saves all IBM Director persistent data except software packages:

```
twgsave -s
```

Return codes

The **twgsave** command returns the following codes.

Code	Meaning
0	The command was successful.
1	An invalid parameter was issued.
2	An IBM Director service is still running.
15	The directory was not found.

twgstart

This topic provides information about the **twgstart** command. This command starts IBM Director Server and IBM Director Agent on i5/OS and Linux operating systems.

```
►►──twgstart──────────────────────────────────────────────────────►◄
```

To start IBM Director Agent, type one of the following commands and press Enter:

For i5/OS	`/qibm/userdata/director/bin/twgstart`
For Linux (32-bit operating systems or AMD64)	`/opt/IBM/director/bin/twgstart`
For Linux (Intel Itanium, IBM iSeries, or IBM pSeries)	`/opt/ibm/director/bin/twgstart`
For NetWare	`load twgipc`
For Windows	`net start twgipc`

Note: Starting IBM Director Agent on NetWare and Windows does not require use of the **twgstart** command; the information is listed here for completeness.

Parameters

This command takes no parameters.

twgstat

This topic provides information about the **twgstat** command. This command returns the active state of IBM Director Server on AIX, i5/OS and Linux operating systems.

The **twgstat** command is not implemented or needed on Windows operating systems; on Windows, the active status of IBM Director Server is displayed in the system tray.

```
►►──twgstat──┬────┬──────────────────────────────────────────────────────►◄
             └─-r─┘
```

To return the active state of IBM Director Server, type one of the following commands and press **Enter**:

For AIX or Linux	`/opt/ibm/director/bin/twgstat`
For i5/OS	`/qibm/userdata/director/bin/twgstat`

Note: The **twgstat** command must be executed on the management server where the installation of IBM Director Server to be monitored is located.

Parameters

-r Specifies that the command will run recursively and check the state of IBM Director Server every five seconds. Whenever a status change occurs, the system time and the new status are displayed.

> **Note:** If needed, the polling interval of five seconds may be changed by editing the **twgstat** script. Use the following procedure with caution, as changing the script could make it inoperable.
>
> 1. Make a backup copy of the **twgstat** script before editing, in case of errors.
> 2. Open the script in a text editor such as vi and locate the following line near the end of the script:
>
>    ```
>    sleep 5
>    ```
> 3. Replace `5` with the number of seconds you wish the script to wait between status checks.
> 4. Save the modified script.
>
> Execute the modified script with the **-r** parameter to check the state of IBM Director Server using the new interval.

Return codes

Code	Meaning
0	**Active**: Process is fully active and ready for work.

Code	Meaning
1	**Starting**: Process is starting but is not yet ready for work.
2	**Ending**: Process was requested to end but has not yet ended.
3	**Inactive**: Process has ended or was never started.
4	**Error**: Process has ended abnormally.
7	**Bad parameters**: Incorrect parameters entered for the **twgstat** command.

twgstop

This topic provides information about the **twgstop** command. This command stops IBM Director Agent on i5/OS and Linux operating systems.

```
►►—twgstop——————————————————————————————————————————————————►◄
```

To stop IBM Director Agent, type one of the following commands and press **Enter**:

For i5/OS	`/qibm/userdata/director/bin/twgend`
For Linux (32-bit operating systems or AMD64)	`/opt/IBM/bin/director/twgstop`
For Linux (Intel Itanium, IBM iSeries, or IBM pSeries)	`/opt/ibm/bin/director/twgstop`
For NetWare	`unload twgipc`
For Windows	`net stop twgipc`

Note: Stopping IBM Director Agent on NetWare and Windows does not require use of the **twgstop** command; the information is listed here for completeness.

Parameters

This command takes no parameters.

uncfgdb

This topic provides information about the **uncfgdb** command. This command starts the graphical database-removal tool for IBM Director.

```
►►—uncfgdb——————————————————————————————————————————————————————————►◄
```

Parameters

This command takes no parameters.

Examples

To remove the IBM Director database: Type the following command to launch the database removal tool:

```
uncfgdb
```

uncfgdbcmd

This topic provides information about the **uncfgdbcmd** command. This command removes the IBM Director database configuration using a database-removal response file.

```
►►──uncfgdbcmd──-rspfile──response_file──-silent────────────────────────────────────────────►◄
```

Parameters

–rspfile *response_file*
Specifies that database removal information is located in a response file. Specify the fully qualified name for the database-removal response file.

–silent
Specifies that the database removal occurs silently. Results are logged to the `cfgdbcmd.output` file in the `director\log` directory. If the database configuration is successfully unconfigured, a return code of `0` is generated. If the database configuration fails to be successfully removed, a non-zero return code is generated.

Example

Remove the database configuration
The following command removes the IBM Director database using the information stored in the `MyDatabaseRemoval.txt` file:

```
uncfgdbcmd -rspfile c:\temp\MyDatabaseRemoval.txt
```

Management Processor Command-Line Interface (MPCLI) overview

This topic provides a brief overview of the IBM Management Processor Command-Line Interface management tool for IBM systems running Linux or Microsoft Windows.

MPCLI provides system management functions from an easy-to-use command-line interface that connects to a service processor. Using this command-line interface, you can access and set a wide range of information about the health, configuration, communication, and state of your system. These functions are immediately available after you install the command-line interface and make a connection to a service processor.

Instructions for installing and using MPCLI are available online at www.ibm.com/support/docview.wss?rs=0&uid=psg1MIGR-54214&loc=en_US.

Once installed, MPCLI may be started in one of two ways:

From the command line
Type ″MPCLI″ and press Enter.

From IBM Director Console
Right-click the managed object on which you want to use the task, then select the MPCLI task from that menu.

On Windows systems, MPCLI may also appear in the Start menu.

Appendix B. Configuration Manager XML file example

The Configuration Manager saves an XML configuration file that contains server or BladeCenter-unit configuration information.

If you installed IBM Director Server in the default location, the XML configuration file is located in the one of the following directories:

For Linux	opt/IBM/director
For i5/OS	/QIBM/UserData/Director
For Windows	*c*:\Program Files\IBM\Director

where *c* is the drive letter of the hard disk drive on which IBM Director Server is installed.

```
<configurationmanager xsi:noNamespaceSchemaLocation="config_manager.xsd"
   title="Configuration Manager"
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   date="2005-6-29"
   author="Configuration Manager" >
     <profileName>asdf</profileName>
     <serverchassis>
          <type>IBM:SVR:Chassis:RackMount</type>
          <detectapply>false</detectapply>
          <computersystem>
               <type>none</type>
               <detectapply>false</detectapply>
               <plugin name="spconfig" version="1.0">
                    <configBlock><sp>
                         <module>
                               <username>testUser</username>
                               <password>passw0rd</password>
                               <replaceuserid>false</replaceuserid>
                               <useraccess>administrator</useraccess>
                         </module>
                    </sp></configBlock>
               </plugin>
               <plugin name="spnetworkconfig" version="1.0">
                    <configBlock><sp>
                         <module>
                               <protocol>
                               <type>snmpv3</type>
                               <state>enabled</state>
                               <snmpProfileName>myContext</snmpProfileName>
                               <username>testUser</username>
                               <authenProtocol>
                                     <type>MD5</type>
                                     <password>passw0rd</password>
                               </authenProtocol>
                               <privacyProtocol>
                                     <type>DES</type>
                                     <password>encrypti0n</password>
                               </privacyProtocol>
                               <ipTrap>1.2.3.4</ipTrap>
                               </protocol>
                        </module>
                    </sp></configBlock>
               </plugin>
          </computersystem>
     </serverchassis>
</configurationmanager>
```

.

Appendix C. Database

This topic provides reference information about the IBM Director database.

Database-configuration response file

This topic provides information about the database-configuration response file.

The cfgdbcmd.rsp file is located in the director\data directory. Modify the cfgdbcmd.rsp file or create additional database-configuration response files as needed. The following is a commented sample file used as an example.

```
===============================================================
;Database Configuration Response File
; A semicolon in the first column indicates a comment statement
;===============================================================
;===============================================================
; Apache Derby
;===============================================================
;DbmsApplication = Apache Derby
;
;===============================================================
; DB2
;===============================================================
;DbmsApplication = DB2
;DbmsTcpIpListenerPort = 50000
;DbmsServerName = xxxxx
;DbmsDatabaseName = xxxxx
;DbmsUserId = xxxxx
;DbmsPassword = xxxxx
;===============================================================

;===============================================================
; SQL Server
;===============================================================
;DbmsApplication = Microsoft SQL Server
;DbmsTcpIpListenerPort = 1433
;DbmsServerName = xxxxx
;DbmsDatabaseName = xxxxx
;DbmsUserId = xxxxx
;DbmsPassword = xxxxx
;===============================================================

;===============================================================
; Oracle
;===============================================================
DbmsApplication = Oracle
DbmsTcpIpListenerPort = 1521
DbmsServerName = xxxxx
DbmsDatabaseName = xxxxx
DbmsUserId = xxxxx
DbmsPassword = xxxxx
DbmsAdminId = xxxxx
DbmsAdminPassword = xxxxx
DbmsTableName = IBM_DIRECTOR_DATA_TS
DbmsTableFile = IBM_DIRECTOR_DATA_DF
DbmsTableFileSize = 500
DbmsTempTableName = IBM_DIRECTOR_TEMP_TS
DbmsTempTableFile = IBM_DIRECTOR_TEMP_DF
DbmsTempTableFileSize = 50
;===============================================================
;===============================================================
; PostgreSQL
;===============================================================
;DbmsApplication = PostgreSQL
;DbmsTcpIpListenerPort = 5432
;DbmsServerName = xxxxx
;DbmsDatabaseName = xxxxx
;DbmsUserId = xxxxx
;DbmsPassword = xxxxx
;===============================================================

;===============================================================
; Database Disabled
;===============================================================
;DbmsApplication = noDatabase
```

Database-removal response file

This topic provides information about the database-removal response file.

The uncfgbcmd.rsp file is located in the director\data directory. Modify the uncfgbcmd.rsp file or create additional database-removal response files as needed. The following is a commented sample file used as an example.

```
;===============================================================
; Database Configuration Removal Response File
; A semicolon in the first column indicates a comment statement
;===============================================================
;
;===============================================================
; Un-Install Keywords. 0=keep 1=remove
;===============================================================
;DbmsRemoveTables = 1
;DbmsRemoveConfiguration = 1
```

Supported database applications

This topic provides information about the database applications that are supported for use with IBM Director. IBM Director Server uses an SQL database to store inventory data for the systems in the environment.

The following tables list the database applications supported by IBM Director Server. They also provide information about whether the database application can be installed locally (on the management server) or remotely.

Table 38. Database applications for management servers running AIX

Database application	Type of installation
Apache Derby (embedded in IBM Director Server)	Local only
IBM DB2 Universal Database™ 8.1with Fix Pack 9a	Local or remote
Oracle Server, versions 9.2 and 10g	Local or remote

Table 39. Database applications for management servers running i5/OS

Database application	Type of installation
IBM DB2 Universal Database for iSeries (part of i5/OS)	Local only

Table 40. Database applications for management servers running Linux (xSeries, System p5 and pSeries, and System z9 and zSeries)

Database application	Type of installation
Apache Derby (embedded in IBM Director Server)	Local only
IBM DB2 Universal Database 8.1with Fix Pack 9a	Local or remote
Microsoft SQL Server 2000 with Service Pack 3a	Remote only
Oracle Server, versions 9.2 and 10g	Local or remote
PostgreSQL, versions 7.2., 7.3, and 7.4	Local or remote

Table 41. Database applications for management servers running Windows

Database application	Type of installation
Apache Derby (embedded in IBM Director Server)	Local only
IBM DB2 Universal Database 8.1, Fix Pack 9a	Local or remote
Microsoft Data Engine (MSDE) 2000 (aka Microsoft SQL Server 2000 Desktop Engine) with Service Pack 3a (local use only)	Local only
Microsoft SQL Server 2000 with Service Pack 3a	Local or remote
Oracle Server, versions 9.2 and 10g	Local or remote

The Microsoft Jet database supported in previous releases of IBM Director continues to be supported only for IBM Director upgrades. MS Jet is *not* available as a database option for new installations.

Appendix D. Default subscriptions for CIM agents

This topic provides information about the default CIM subscriptions for IBM Director Core Services and IBM Director Agent.

The following script will run during installation of either IBM Director Core Services or IBM Director Agent on a managed system.

Note: Long command lines have been broken, with subsequent lines indented for clarity.

```
##########################################################################################
## Filters
##########################################################################################
cimsubscribe -cf -fn "Voltage Sensor Normals"   -q "SELECT * from IBMPSG_VoltageEvent where severity = 2"
cimsubscribe -cf -fn "Voltage Sensor Warnings"  -q "SELECT * from IBMPSG_VoltageEvent where severity = 4"
cimsubscribe -cf -fn "Voltage Sensor Criticals" -q "SELECT * from IBMPSG_VoltageEvent where severity = 6"

cimsubscribe -cf -fn "Temperature Sensor Normals"
                 -q "SELECT * from IBMPSG_TemperatureEvent where severity = 2"
cimsubscribe -cf -fn "Temperature Sensor Warnings"
                 -q "SELECT * from IBMPSG_TemperatureEvent where severity = 4"
cimsubscribe -cf -fn "Temperature Sensor Criticals"
                 -q "SELECT * from IBMPSG_TemperatureEvent where severity = 6"

cimsubscribe -cf -fn "Tachometer Normals"          -q "SELECT * from IBMPSG_FanEvent where severity = 2"
cimsubscribe -cf -fn "Tachometer Sensor Warnings"  -q "SELECT * from IBMPSG_FanEvent where severity = 4"
cimsubscribe -cf -fn "Tachometer Sensor Criticals" -q "SELECT * from IBMPSG_FanEvent where severity = 6"

cimsubscribe -cf -fn "Lease Normals"   -q "SELECT * from IBMPSG_LeaseExpirationEvent where severity = 2"
cimsubscribe -cf -fn "Lease Warnings"  -q "SELECT * from IBMPSG_LeaseExpirationEvent where severity = 4"
cimsubscribe -cf -fn "Lease Criticals" -q "SELECT * from IBMPSG_LeaseExpirationEvent where severity = 6"

cimsubscribe -cf -fn "Warranty Normals"
                 -q "SELECT * from IBMPSG_WarrantyExpirationEvent where severity = 2"
cimsubscribe -cf -fn "Warranty Warnings"
                 -q "SELECT * from IBMPSG_WarrantyExpirationEvent where severity = 4"
cimsubscribe -cf -fn "Warranty Criticals"
                 -q "SELECT * from IBMPSG_WarrantyExpirationEvent where severity = 6"

cimsubscribe -cf -fn "Processor PFA Normals"
                 -q "SELECT * from IBMPSG_ProcessorPFEvent where severity = 2"
cimsubscribe -cf -fn "Processor PFA Warnings"
                 -q "SELECT * from IBMPSG_ProcessorPFEvent where severity = 4"
cimsubscribe -cf -fn "Processor PFA Criticals"
                 -q "SELECT * from IBMPSG_ProcessorPFEvent where severity = 6"

cimsubscribe -cf -fn "Memory PFA Normals"   -q "SELECT * from IBMPSG_MemoryPFEvent where severity = 2"
cimsubscribe -cf -fn "Memory PFA Warnings"  -q "SELECT * from IBMPSG_MemoryPFEvent where severity = 4"
cimsubscribe -cf -fn "Memory PFA Criticals" -q "SELECT * from IBMPSG_MemoryPFEvent where severity = 6"

cimsubscribe -cf -fn "Power Supply Normals"
                 -q "SELECT * from IBMPSG_PowerSupplyEvent where severity = 2"
cimsubscribe -cf -fn "Power Supply Warnings"
                 -q "SELECT * from IBMPSG_PowerSupplyEvent where severity = 4"
cimsubscribe -cf -fn "Power Supply Criticals"
                 -q "SELECT * from IBMPSG_PowerSupplyEvent where severity = 6"

cimsubscribe -cf -fn "Service Processor Error Log Normals"
                 -q "SELECT * from IBMPSG_SP_ErrorLogEvent where severity = 2"
cimsubscribe -cf -fn "Service Processor Error Log Warnings"
                 -q "SELECT * from IBMPSG_SP_ErrorLogEvent where severity = 4"
cimsubscribe -cf -fn "Service Processor Error Log Criticals"
                 -q "SELECT * from IBMPSG_SP_ErrorLogEvent where severity = 6"

cimsubscribe -cf -fn "Service Processor PFA Normals"
                 -q "SELECT * from IBMPSG_SP_PFAEvent where severity = 2"
cimsubscribe -cf -fn "Service Processor PFA Warnings"
                 -q "SELECT * from IBMPSG_SP_PFAEvent where severity = 4"
cimsubscribe -cf -fn "Service Processor PFA Criticals"
                 -q "SELECT * from IBMPSG_SP_PFAEvent where severity = 6"

cimsubscribe -cf -fn "Service Processor Remote Login Normals"
                 -q "SELECT * from IBMPSG_SP_RemoteLoginEvent where severity = 2"
cimsubscribe -cf -fn "Service Processor Remote Login Warnings"
                 -q "SELECT * from IBMPSG_SP_RemoteLoginEvent where severity = 4"
cimsubscribe -cf -fn "Service Processor Remote Login Criticals"
                 -q "SELECT * from IBMPSG_SP_RemoteLoginEvent where severity = 6"

cimsubscribe -cf -fn "Service Processor DASD Backplane Normals"
```

```
                 -q "SELECT * from IBMPSG_SP_DASDBackplaneEvent where severity = 2"
cimsubscribe -cf -fn "Service Processor DASD Backplane Warnings"
                 -q "SELECT * from IBMPSG_SP_DASDBackplaneEvent where severity = 4"
cimsubscribe -cf -fn "Service Processor DASD Backplane Criticals"
                 -q "SELECT * from IBMPSG_SP_DASDBackplaneEvent where severity = 6"

cimsubscribe -cf -fn "Service Processor Generic Fan Normals"
                 -q "SELECT * from IBMPSG_SP_GenericFanEvent where severity = 2"
cimsubscribe -cf -fn "Service Processor Generic Fan Warnings"
                 -q "SELECT * from IBMPSG_SP_GenericFanEvent where severity = 4"
cimsubscribe -cf -fn "Service Processor Generic Fan Criticals"
                 -q "SELECT * from IBMPSG_SP_GenericFanEvent where severity = 6"

cimsubscribe -cf -fn "Service Processor Generic Volatge Normals"
                 -q "SELECT * from IBMPSG_SP_GenericVoltageEvent where severity = 2"
cimsubscribe -cf -fn "Service Processor Generic Volatge Warnings"
                 -q "SELECT * from IBMPSG_SP_GenericVoltageEvent where severity = 4"
cimsubscribe -cf -fn "Service Processor Generic Volatge Criticals"
                 -q "SELECT * from IBMPSG_SP_GenericVoltageEvent where severity = 6"

cimsubscribe -cf -fn "Storage Normals"   -q "SELECT * from IBMPSG_StorageEvent where severity = 2"
cimsubscribe -cf -fn "Storage Warnings"  -q "SELECT * from IBMPSG_StorageEvent where severity = 4"
cimsubscribe -cf -fn "Storage Criticals" -q "SELECT * from IBMPSG_StorageEvent where severity = 6"

cimsubscribe -cf -fn "SMART Drive Normals"   -q "SELECT * from IBMPSG_SMARTEvent where severity = 2"
cimsubscribe -cf -fn "SMART Drive Warnings"  -q "SELECT * from IBMPSG_SMARTEvent where severity = 4"
cimsubscribe -cf -fn "SMART Drive Criticals" -q "SELECT * from IBMPSG_SMARTEvent where severity = 6"

cimsubscribe -cf -fn "RAID Normals"   -q "SELECT * from IBMPSG_StorageRAIDEvent where severity = 2"
cimsubscribe -cf -fn "RAID Warnings"  -q "SELECT * from IBMPSG_StorageRAIDEvent where severity = 4"
cimsubscribe -cf -fn "RAID Criticals" -q "SELECT * from IBMPSG_StorageRAIDEvent where severity = 6"

cimsubscribe -cf -fn "RAID System Health Normals"
                 -q "SELECT * from IBMPSG_StorageRAIDHealthEvent where severity = 2"
cimsubscribe -cf -fn "RAID System Health Warnings"
                 -q "SELECT * from IBMPSG_StorageRAIDHealthEvent where severity = 4"
cimsubscribe -cf -fn "RAID System Health Criticals"
                 -q "SELECT * from IBMPSG_StorageRAIDHealthEvent where severity = 6"

cimsubscribe -cf -fn "Redundant NIC Normals"
                 -q "SELECT * from IBMPSG_RedundantNetworkAdapterEvent where severity = 2"
cimsubscribe -cf -fn "Redundant NIC Warnings"
                 -q "SELECT * from IBMPSG_RedundantNetworkAdapterEvent where severity = 4"
cimsubscribe -cf -fn "Redundant NIC Criticals"
                 -q "SELECT * from IBMPSG_RedundantNetworkAdapterEvent where severity = 6"

cimsubscribe -cf -fn "Redundant NIC Switchover Normals"
                 -q "SELECT * from IBMPSG_RedundantNetworkAdapterSwitchoverEvent where severity = 2"
cimsubscribe -cf -fn "Redundant NIC Switchover Warnings"
                 -q "SELECT * from IBMPSG_RedundantNetworkAdapterSwitchoverEvent where severity = 4"
cimsubscribe -cf -fn "Redundant NIC Switchover Criticals"
                 -q "SELECT * from IBMPSG_RedundantNetworkAdapterSwitchoverEvent where severity = 6"

cimsubscribe -cf -fn "Redundant NIC Switchback Normals"
                 -q "SELECT * from IBMPSG_RedundantNetworkAdapterSwitchbackEvent where severity = 2"
cimsubscribe -cf -fn "Redundant NIC Switchback Warnings"
                 -q "SELECT * from IBMPSG_RedundantNetworkAdapterSwitchbackEvent where severity = 4"
cimsubscribe -cf -fn "Redundant NIC Switchback Criticals"
                 -q "SELECT * from IBMPSG_RedundantNetworkAdapterSwitchbackEvent where severity = 6"

cimsubscribe -cf -fn "Network Adapter Normals"
                 -q "SELECT * from IBMPSG_NetworkAdapterFailedEvent where severity = 2"
cimsubscribe -cf -fn "Network Adapter Warnings"
                 -q "SELECT * from IBMPSG_NetworkAdapterFailedEvent where severity = 4"
cimsubscribe -cf -fn "Network Adapter Criticals"
                 -q "SELECT * from IBMPSG_NetworkAdapterFailedEvent where severity = 6"

cimsubscribe -cf -fn "Network Adapter Offline Normals"
                 -q "SELECT * from IBMPSG_NetworkAdapterOfflineEvent where severity = 2"
cimsubscribe -cf -fn "Network Adapter Offline Warnings"
                 -q "SELECT * from IBMPSG_NetworkAdapterOfflineEvent where severity = 4"
cimsubscribe -cf -fn "Network Adapter Offline Criticals"
                 -q "SELECT * from IBMPSG_NetworkAdapterOfflineEvent where severity = 6"

cimsubscribe -cf -fn "Network Adapter Online Normals"
                 -q "SELECT * from IBMPSG_NetworkAdapterOnlineEvent where severity = 2"
cimsubscribe -cf -fn "Network Adapter Online Warnings"
                 -q "SELECT * from IBMPSG_NetworkAdapterOnlineEvent where severity = 4"
cimsubscribe -cf -fn "Network Adapter Online Criticals"
                 -q "SELECT * from IBMPSG_NetworkAdapterOnlineEvent where severity = 6"

cimsubscribe -cf -fn "Chassis Normals"   -q "SELECT * from IBMPSG_ChassisEvent where severity = 2"
cimsubscribe -cf -fn "Chassis Warnings"  -q "SELECT * from IBMPSG_ChassisEvent where severity = 4"
cimsubscribe -cf -fn "Chassis Criticals" -q "SELECT * from IBMPSG_ChassisEvent where severity = 6"

cimsubscribe -cf -fn "LAN Leash Normals"   -q "SELECT * from IBMPSG_LANLeashEvent where severity = 2"
cimsubscribe -cf -fn "LAN Leash Warnings"  -q "SELECT * from IBMPSG_LANLeashEvent where severity = 4"
cimsubscribe -cf -fn "LAN Leash Criticals" -q "SELECT * from IBMPSG_LANLeashEvent where severity = 6"

##########################################################################################
```

```
## Handlers
##########################################################################################
cimsubscribe -ch -hn "SNMP"   -d "http://localhost:8888/CIMListener/snmp"
cimsubscribe -ch -hn "TEC"    -d "http://localhost:8888/CIMListener/tec"
cimsubscribe -ch -hn "Log"    -d "http://localhost:8888/CIMListener/log"
cimsubscribe -ch -hn "Health" -d "http://localhost:8888/CIMListener/health"
cimsubscribe -ch -hn "SMS"    -d "http://localhost:8888/CIMListener/sms"
cimsubscribe -ch -hn "PopUp"  -d "http://localhost:8888/CIMListener/popup"



##########################################################################################
## Subscriptions
##########################################################################################
##########################################################################################
## SNMP Subscriptions
##########################################################################################

cimsubscribe -cs -fn "Voltage Sensor Normals" -hn "SNMP"
cimsubscribe -cs -fn "Voltage Sensor Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Voltage Sensor Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Temperature Sensor Normals" -hn "SNMP"
cimsubscribe -cs -fn "Temperature Sensor Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Temperature Sensor Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Tachometer Normals" -hn "SNMP"
cimsubscribe -cs -fn "Tachometer Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Tachometer Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Lease Normals" -hn "SNMP"
cimsubscribe -cs -fn "Lease Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Lease Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Warranty Normals" -hn "SNMP"
cimsubscribe -cs -fn "Warranty Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Warranty Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Processor PFA Normals" -hn "SNMP"
cimsubscribe -cs -fn "Processor PFA Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Processor PFA Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Memory PFA Normals" -hn "SNMP"
cimsubscribe -cs -fn "Memory PFA Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Memory PFA Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Power Supply Normals" -hn "SNMP"
cimsubscribe -cs -fn "Power Supply Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Power Supply Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Service Processor Error Log Normals" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor Error Log Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor Error Log Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Service Processor PFA Normals" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor PFA Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor PFA Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Service Processor Remote Login Normals" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor Remote Login Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor Remote Login Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Service Processor DASD Backplane Normals" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor DASD Backplane Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor DASD Backplane Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Service Processor Generic Fan Normals" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor Generic Fan Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor Generic Fan Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Service Processor Generic Volatge Normals" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor Generic Volatge Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Service Processor Generic Volatge Criticals" -hn "SNMP"

cimsubscribe -cs -fn "SMART Drive Normals" -hn "SNMP"
cimsubscribe -cs -fn "SMART Drive Warnings" -hn "SNMP"
cimsubscribe -cs -fn "SMART Drive Criticals" -hn "SNMP"

cimsubscribe -cs -fn "RAID Normals" -hn "SNMP"
cimsubscribe -cs -fn "RAID Warnings" -hn "SNMP"
cimsubscribe -cs -fn "RAID Criticals" -hn "SNMP"

cimsubscribe -cs -fn "RAID System Health Normals" -hn "SNMP"
cimsubscribe -cs -fn "RAID System Health Warnings" -hn "SNMP"
cimsubscribe -cs -fn "RAID System Health Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Redundant NIC Normals" -hn "SNMP"
cimsubscribe -cs -fn "Redundant NIC Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Redundant NIC Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Redundant NIC Switchover Normals" -hn "SNMP"
cimsubscribe -cs -fn "Redundant NIC Switchover Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Redundant NIC Switchover Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Redundant NIC Switchback Normals" -hn "SNMP"
cimsubscribe -cs -fn "Redundant NIC Switchback Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Redundant NIC Switchback Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Network Adapter Normals" -hn "SNMP"
cimsubscribe -cs -fn "Network Adapter Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Network Adapter Criticals" -hn "SNMP"

cimsubscribe -cs -fn "Chassis Normals" -hn "SNMP"
cimsubscribe -cs -fn "Chassis Warnings" -hn "SNMP"
cimsubscribe -cs -fn "Chassis Criticals" -hn "SNMP"

cimsubscribe -cs -fn "LAN Leash Normals" -hn "SNMP"
cimsubscribe -cs -fn "LAN Leash Warnings" -hn "SNMP"
cimsubscribe -cs -fn "LAN Leash Criticals" -hn "SNMP"

##########################################################################################
## TEC Subscriptions
##########################################################################################

cimsubscribe -cs -fn "Voltage Sensor Normals" -hn "TEC"
cimsubscribe -cs -fn "Voltage Sensor Warnings" -hn "TEC"
cimsubscribe -cs -fn "Voltage Sensor Criticals" -hn "TEC"

cimsubscribe -cs -fn "Temperature Sensor Normals" -hn "TEC"
cimsubscribe -cs -fn "Temperature Sensor Warnings" -hn "TEC"
cimsubscribe -cs -fn "Temperature Sensor Criticals" -hn "TEC"

cimsubscribe -cs -fn "Tachometer Normals" -hn "TEC"
cimsubscribe -cs -fn "Tachometer Warnings" -hn "TEC"
cimsubscribe -cs -fn "Tachometer Criticals" -hn "TEC"

cimsubscribe -cs -fn "Lease Normals" -hn "TEC"
cimsubscribe -cs -fn "Lease Warnings" -hn "TEC"
cimsubscribe -cs -fn "Lease Criticals" -hn "TEC"

cimsubscribe -cs -fn "Warranty Normals" -hn "TEC"
cimsubscribe -cs -fn "Warranty Warnings" -hn "TEC"
cimsubscribe -cs -fn "Warranty Criticals" -hn "TEC"

cimsubscribe -cs -fn "Processor PFA Normals" -hn "TEC"
cimsubscribe -cs -fn "Processor PFA Warnings" -hn "TEC"
cimsubscribe -cs -fn "Processor PFA Criticals" -hn "TEC"

cimsubscribe -cs -fn "Memory PFA Normals" -hn "TEC"
cimsubscribe -cs -fn "Memory PFA Warnings" -hn "TEC"
cimsubscribe -cs -fn "Memory PFA Criticals" -hn "TEC"

cimsubscribe -cs -fn "Power Supply Normals" -hn "TEC"
cimsubscribe -cs -fn "Power Supply Warnings" -hn "TEC"
cimsubscribe -cs -fn "Power Supply Criticals" -hn "TEC"

cimsubscribe -cs -fn "Service Processor Error Log Normals" -hn "TEC"
cimsubscribe -cs -fn "Service Processor Error Log Warnings" -hn "TEC"
cimsubscribe -cs -fn "Service Processor Error Log Criticals" -hn "TEC"

cimsubscribe -cs -fn "Service Processor PFA Normals" -hn "TEC"
cimsubscribe -cs -fn "Service Processor PFA Warnings" -hn "TEC"
cimsubscribe -cs -fn "Service Processor PFA Criticals" -hn "TEC"

cimsubscribe -cs -fn "Service Processor Remote Login Normals" -hn "TEC"
cimsubscribe -cs -fn "Service Processor Remote Login Warnings" -hn "TEC"
cimsubscribe -cs -fn "Service Processor Remote Login Criticals" -hn "TEC"

cimsubscribe -cs -fn "Service Processor DASD Backplane Normals" -hn "TEC"
cimsubscribe -cs -fn "Service Processor DASD Backplane Warnings" -hn "TEC"
cimsubscribe -cs -fn "Service Processor DASD Backplane Criticals" -hn "TEC"

cimsubscribe -cs -fn "Service Processor Generic Fan Normals" -hn "TEC"
cimsubscribe -cs -fn "Service Processor Generic Fan Warnings" -hn "TEC"
cimsubscribe -cs -fn "Service Processor Generic Fan Criticals" -hn "TEC"

cimsubscribe -cs -fn "Service Processor Generic Volatge Normals" -hn "TEC"
cimsubscribe -cs -fn "Service Processor Generic Volatge Warnings" -hn "TEC"
cimsubscribe -cs -fn "Service Processor Generic Volatge Criticals" -hn "TEC"

cimsubscribe -cs -fn "SMART Drive Normals" -hn "TEC"
cimsubscribe -cs -fn "SMART Drive Warnings" -hn "TEC"
cimsubscribe -cs -fn "SMART Drive Criticals" -hn "TEC"

cimsubscribe -cs -fn "RAID Normals" -hn "TEC"
cimsubscribe -cs -fn "RAID Warnings" -hn "TEC"
cimsubscribe -cs -fn "RAID Criticals" -hn "TEC"

cimsubscribe -cs -fn "RAID System Health Normals" -hn "TEC"
cimsubscribe -cs -fn "RAID System Health Warnings" -hn "TEC"
cimsubscribe -cs -fn "RAID System Health Criticals" -hn "TEC"
cimsubscribe -cs -fn "Redundant NIC Normals" -hn "TEC"
cimsubscribe -cs -fn "Redundant NIC Warnings" -hn "TEC"
cimsubscribe -cs -fn "Redundant NIC Criticals" -hn "TEC"

cimsubscribe -cs -fn "Redundant NIC Switchover Normals" -hn "TEC"
cimsubscribe -cs -fn "Redundant NIC Switchover Warnings" -hn "TEC"
cimsubscribe -cs -fn "Redundant NIC Switchover Criticals" -hn "TEC"

cimsubscribe -cs -fn "Redundant NIC Switchback Normals" -hn "TEC"
cimsubscribe -cs -fn "Redundant NIC Switchback Warnings" -hn "TEC"
cimsubscribe -cs -fn "Redundant NIC Switchback Criticals" -hn "TEC"

cimsubscribe -cs -fn "Network Adapter Normals" -hn "TEC"
cimsubscribe -cs -fn "Network Adapter Warnings" -hn "TEC"
cimsubscribe -cs -fn "Network Adapter Criticals" -hn "TEC"

cimsubscribe -cs -fn "Chassis Normals" -hn "TEC"
cimsubscribe -cs -fn "Chassis Warnings" -hn "TEC"
cimsubscribe -cs -fn "Chassis Criticals" -hn "TEC"

cimsubscribe -cs -fn "LAN Leash Normals" -hn "TEC"
cimsubscribe -cs -fn "LAN Leash Warnings" -hn "TEC"
cimsubscribe -cs -fn "LAN Leash Criticals" -hn "TEC"

##########################################################################################
## Log Subscriptions
##########################################################################################

cimsubscribe -cs -fn "Voltage Sensor Warnings" -hn "Log"
cimsubscribe -cs -fn "Voltage Sensor Criticals" -hn "Log"

cimsubscribe -cs -fn "Temperature Sensor Warnings" -hn "Log"
cimsubscribe -cs -fn "Temperature Sensor Criticals" -hn "Log"

cimsubscribe -cs -fn "Tachometer Warnings" -hn "Log"
cimsubscribe -cs -fn "Tachometer Criticals" -hn "Log"

cimsubscribe -cs -fn "Lease Warnings" -hn "Log"
cimsubscribe -cs -fn "Lease Criticals" -hn "Log"

cimsubscribe -cs -fn "Warranty Warnings" -hn "Log"
cimsubscribe -cs -fn "Warranty Criticals" -hn "Log"

cimsubscribe -cs -fn "Processor PFA Warnings" -hn "Log"
cimsubscribe -cs -fn "Processor PFA Criticals" -hn "Log"

cimsubscribe -cs -fn "Memory PFA Warnings" -hn "Log"
cimsubscribe -cs -fn "Memory PFA Criticals" -hn "Log"

cimsubscribe -cs -fn "Power Supply Warnings" -hn "Log"
cimsubscribe -cs -fn "Power Supply Criticals" -hn "Log"

cimsubscribe -cs -fn "Service Processor Error Log Warnings" -hn "Log"
cimsubscribe -cs -fn "Service Processor Error Log Criticals" -hn "Log"

cimsubscribe -cs -fn "Service Processor PFA Warnings" -hn "Log"
cimsubscribe -cs -fn "Service Processor PFA Criticals" -hn "Log"

cimsubscribe -cs -fn "Service Processor Remote Login Warnings" -hn "Log"
cimsubscribe -cs -fn "Service Processor Remote Login Criticals" -hn "Log"

cimsubscribe -cs -fn "Service Processor DASD Backplane Warnings" -hn "Log"
cimsubscribe -cs -fn "Service Processor DASD Backplane Criticals" -hn "Log"

cimsubscribe -cs -fn "Service Processor Generic Fan Warnings" -hn "Log"
cimsubscribe -cs -fn "Service Processor Generic Fan Criticals" -hn "Log"

cimsubscribe -cs -fn "Service Processor Generic Volatge Warnings" -hn "Log"
cimsubscribe -cs -fn "Service Processor Generic Volatge Criticals" -hn "Log"

cimsubscribe -cs -fn "SMART Drive Warnings" -hn "Log"
cimsubscribe -cs -fn "SMART Drive Criticals" -hn "Log"

cimsubscribe -cs -fn "RAID Warnings" -hn "Log"
cimsubscribe -cs -fn "RAID Criticals" -hn "Log"

cimsubscribe -cs -fn "RAID System Health Warnings" -hn "Log"
cimsubscribe -cs -fn "RAID System Health Criticals" -hn "Log"

cimsubscribe -cs -fn "Redundant NIC Warnings" -hn "Log"
cimsubscribe -cs -fn "Redundant NIC Criticals" -hn "Log"

cimsubscribe -cs -fn "Redundant NIC Switchover Warnings" -hn "Log"
cimsubscribe -cs -fn "Redundant NIC Switchover Criticals" -hn "Log"

cimsubscribe -cs -fn "Redundant NIC Switchback Warnings" -hn "Log"
cimsubscribe -cs -fn "Redundant NIC Switchback Criticals" -hn "Log"

cimsubscribe -cs -fn "Network Adapter Warnings" -hn "Log"
cimsubscribe -cs -fn "Network Adapter Criticals" -hn "Log"
cimsubscribe -cs -fn "Chassis Warnings" -hn "Log"
cimsubscribe -cs -fn "Chassis Criticals" -hn "Log"

cimsubscribe -cs -fn "LAN Leash Warnings" -hn "Log"
cimsubscribe -cs -fn "LAN Leash Criticals" -hn "Log"

##########################################################################################
## Health Subscriptions
##########################################################################################

cimsubscribe -cs -fn "Voltage Sensor Normals" -hn "Health"
cimsubscribe -cs -fn "Voltage Sensor Warnings" -hn "Health"
cimsubscribe -cs -fn "Voltage Sensor Criticals" -hn "Health"

cimsubscribe -cs -fn "Temperature Sensor Normals" -hn "Health"
cimsubscribe -cs -fn "Temperature Sensor Warnings" -hn "Health"
cimsubscribe -cs -fn "Temperature Sensor Criticals" -hn "Health"

cimsubscribe -cs -fn "Tachometer Normals" -hn "Health"
cimsubscribe -cs -fn "Tachometer Warnings" -hn "Health"
cimsubscribe -cs -fn "Tachometer Criticals" -hn "Health"

cimsubscribe -cs -fn "Lease Normals" -hn "Health"
cimsubscribe -cs -fn "Lease Warnings" -hn "Health"
cimsubscribe -cs -fn "Lease Criticals" -hn "Health"

cimsubscribe -cs -fn "Warranty Normals" -hn "Health"
cimsubscribe -cs -fn "Warranty Warnings" -hn "Health"
cimsubscribe -cs -fn "Warranty Criticals" -hn "Health"

cimsubscribe -cs -fn "Processor PFA Normals" -hn "Health"
cimsubscribe -cs -fn "Processor PFA Warnings" -hn "Health"
cimsubscribe -cs -fn "Processor PFA Criticals" -hn "Health"

cimsubscribe -cs -fn "Memory PFA Normals" -hn "Health"
cimsubscribe -cs -fn "Memory PFA Warnings" -hn "Health"
cimsubscribe -cs -fn "Memory PFA Criticals" -hn "Health"

cimsubscribe -cs -fn "Power Supply Normals" -hn "Health"
cimsubscribe -cs -fn "Power Supply Warnings" -hn "Health"
cimsubscribe -cs -fn "Power Supply Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor Error Health Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor Error Health Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor Error Health Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor PFA Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor PFA Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor PFA Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor Remote Healthin Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor Remote Healthin Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor Remote Healthin Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor DASD Backplane Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor DASD Backplane Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor DASD Backplane Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor Generic Fan Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor Generic Fan Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor Generic Fan Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor Generic Volatge Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor Generic Volatge Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor Generic Volatge Criticals" -hn "Health"

cimsubscribe -cs -fn "Storage Normals" -hn "Health"
cimsubscribe -cs -fn "Storage Warnings" -hn "Health"
cimsubscribe -cs -fn "Storage Criticals" -hn "Health"

cimsubscribe -cs -fn "SMART Drive Normals" -hn "Health"
cimsubscribe -cs -fn "SMART Drive Warnings" -hn "Health"
cimsubscribe -cs -fn "SMART Drive Criticals" -hn "Health"

cimsubscribe -cs -fn "RAID Normals" -hn "Health"
cimsubscribe -cs -fn "RAID Warnings" -hn "Health"
cimsubscribe -cs -fn "RAID Criticals" -hn "Health"

cimsubscribe -cs -fn "RAID System Health Normals" -hn "Health"
cimsubscribe -cs -fn "RAID System Health Warnings" -hn "Health"
cimsubscribe -cs -fn "RAID System Health Criticals" -hn "Health"

cimsubscribe -cs -fn "Redundant NIC Normals" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Warnings" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Criticals" -hn "Health"

cimsubscribe -cs -fn "Redundant NIC Switchover Normals" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Switchover Warnings" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Switchover Criticals" -hn "Health"

cimsubscribe -cs -fn "Redundant NIC Switchback Normals" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Switchback Warnings" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Switchback Criticals" -hn "Health"

cimsubscribe -cs -fn "Network Adapter Normals" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Warnings" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Criticals" -hn "Health"

cimsubscribe -cs -fn "Network Adapter Offline Normals" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Offline Warnings" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Offline Criticals" -hn "Health"

cimsubscribe -cs -fn "Network Adapter Online Normals" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Online Warnings" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Online Criticals" -hn "Health"

cimsubscribe -cs -fn "Chassis Normals" -hn "Health"
cimsubscribe -cs -fn "Chassis Warnings" -hn "Health"
cimsubscribe -cs -fn "Chassis Criticals" -hn "Health"

cimsubscribe -cs -fn "LAN Leash Normals" -hn "Health"
cimsubscribe -cs -fn "LAN Leash Warnings" -hn "Health"
cimsubscribe -cs -fn "LAN Leash Criticals" -hn "Health"

##########################################################################################
## SMS Subscriptions
##########################################################################################

cimsubscribe -cs -fn "Voltage Sensor Normals" -hn "SMS"
cimsubscribe -cs -fn "Voltage Sensor Warnings" -hn "SMS"
cimsubscribe -cs -fn "Voltage Sensor Criticals" -hn "SMS"

cimsubscribe -cs -fn "Temperature Sensor Normals" -hn "SMS"
cimsubscribe -cs -fn "Temperature Sensor Warnings" -hn "SMS"
cimsubscribe -cs -fn "Temperature Sensor Criticals" -hn "SMS"

cimsubscribe -cs -fn "Tachometer Normals" -hn "SMS"
cimsubscribe -cs -fn "Tachometer Warnings" -hn "SMS"
cimsubscribe -cs -fn "Tachometer Criticals" -hn "SMS"

cimsubscribe -cs -fn "Lease Normals" -hn "SMS"
cimsubscribe -cs -fn "Lease Warnings" -hn "SMS"
cimsubscribe -cs -fn "Lease Criticals" -hn "SMS"

cimsubscribe -cs -fn "Warranty Normals" -hn "SMS"
cimsubscribe -cs -fn "Warranty Warnings" -hn "SMS"
cimsubscribe -cs -fn "Warranty Criticals" -hn "SMS"

cimsubscribe -cs -fn "Processor PFA Normals" -hn "SMS"
cimsubscribe -cs -fn "Processor PFA Warnings" -hn "SMS"
cimsubscribe -cs -fn "Processor PFA Criticals" -hn "SMS"

cimsubscribe -cs -fn "Memory PFA Normals" -hn "SMS"
cimsubscribe -cs -fn "Memory PFA Warnings" -hn "SMS"
cimsubscribe -cs -fn "Memory PFA Criticals" -hn "SMS"

cimsubscribe -cs -fn "Power Supply Normals" -hn "SMS"
cimsubscribe -cs -fn "Power Supply Warnings" -hn "SMS"
cimsubscribe -cs -fn "Power Supply Criticals" -hn "SMS"

cimsubscribe -cs -fn "Service Processor Error SMS Normals" -hn "SMS"
cimsubscribe -cs -fn "Service Processor Error SMS Warnings" -hn "SMS"
cimsubscribe -cs -fn "Service Processor Error SMS Criticals" -hn "SMS"

cimsubscribe -cs -fn "Service Processor PFA Normals" -hn "SMS"
cimsubscribe -cs -fn "Service Processor PFA Warnings" -hn "SMS"
cimsubscribe -cs -fn "Service Processor PFA Criticals" -hn "SMS"

cimsubscribe -cs -fn "Service Processor Remote SMSin Normals" -hn "SMS"
cimsubscribe -cs -fn "Service Processor Remote SMSin Warnings" -hn "SMS"
cimsubscribe -cs -fn "Service Processor Remote SMSin Criticals" -hn "SMS"

cimsubscribe -cs -fn "Service Processor DASD Backplane Normals" -hn "SMS"
cimsubscribe -cs -fn "Service Processor DASD Backplane Warnings" -hn "SMS"
cimsubscribe -cs -fn "Service Processor DASD Backplane Criticals" -hn "SMS"

cimsubscribe -cs -fn "Service Processor Generic Fan Normals" -hn "SMS"
cimsubscribe -cs -fn "Service Processor Generic Fan Warnings" -hn "SMS"
cimsubscribe -cs -fn "Service Processor Generic Fan Criticals" -hn "SMS"

cimsubscribe -cs -fn "Service Processor Generic Volatge Normals" -hn "SMS"
cimsubscribe -cs -fn "Service Processor Generic Volatge Warnings" -hn "SMS"
cimsubscribe -cs -fn "Service Processor Generic Volatge Criticals" -hn "SMS"

cimsubscribe -cs -fn "SMART Drive Normals" -hn "SMS"
cimsubscribe -cs -fn "SMART Drive Warnings" -hn "SMS"
cimsubscribe -cs -fn "SMART Drive Criticals" -hn "SMS"

cimsubscribe -cs -fn "RAID Normals" -hn "SMS"
cimsubscribe -cs -fn "RAID Warnings" -hn "SMS"
cimsubscribe -cs -fn "RAID Criticals" -hn "SMS"
cimsubscribe -cs -fn "RAID System SMS Normals" -hn "SMS"
cimsubscribe -cs -fn "RAID System SMS Warnings" -hn "SMS"
cimsubscribe -cs -fn "RAID System SMS Criticals" -hn "SMS"

cimsubscribe -cs -fn "Redundant NIC Normals" -hn "SMS"
cimsubscribe -cs -fn "Redundant NIC Warnings" -hn "SMS"
cimsubscribe -cs -fn "Redundant NIC Criticals" -hn "SMS"

cimsubscribe -cs -fn "Redundant NIC Switchover Normals" -hn "SMS"
cimsubscribe -cs -fn "Redundant NIC Switchover Warnings" -hn "SMS"
cimsubscribe -cs -fn "Redundant NIC Switchover Criticals" -hn "SMS"

cimsubscribe -cs -fn "Redundant NIC Switchback Normals" -hn "SMS"
cimsubscribe -cs -fn "Redundant NIC Switchback Warnings" -hn "SMS"
cimsubscribe -cs -fn "Redundant NIC Switchback Criticals" -hn "SMS"

cimsubscribe -cs -fn "Network Adapter Normals" -hn "SMS"
cimsubscribe -cs -fn "Network Adapter Warnings" -hn "SMS"
cimsubscribe -cs -fn "Network Adapter Criticals" -hn "SMS"

cimsubscribe -cs -fn "Chassis Normals" -hn "SMS"
cimsubscribe -cs -fn "Chassis Warnings" -hn "SMS"
cimsubscribe -cs -fn "Chassis Criticals" -hn "SMS"

cimsubscribe -cs -fn "LAN Leash Normals" -hn "SMS"
cimsubscribe -cs -fn "LAN Leash Warnings" -hn "SMS"
cimsubscribe -cs -fn "LAN Leash Criticals" -hn "SMS"

##########################################################################################
## Health Subscriptions
##########################################################################################

cimsubscribe -cs -fn "Voltage Sensor Normals" -hn "Health"
cimsubscribe -cs -fn "Voltage Sensor Warnings" -hn "Health"
cimsubscribe -cs -fn "Voltage Sensor Criticals" -hn "Health"

cimsubscribe -cs -fn "Temperature Sensor Normals" -hn "Health"
cimsubscribe -cs -fn "Temperature Sensor Warnings" -hn "Health"
cimsubscribe -cs -fn "Temperature Sensor Criticals" -hn "Health"

cimsubscribe -cs -fn "Tachometer Normals" -hn "Health"
cimsubscribe -cs -fn "Tachometer Warnings" -hn "Health"
cimsubscribe -cs -fn "Tachometer Criticals" -hn "Health"

cimsubscribe -cs -fn "Lease Normals" -hn "Health"
cimsubscribe -cs -fn "Lease Warnings" -hn "Health"
cimsubscribe -cs -fn "Lease Criticals" -hn "Health"

cimsubscribe -cs -fn "Warranty Normals" -hn "Health"
cimsubscribe -cs -fn "Warranty Warnings" -hn "Health"
cimsubscribe -cs -fn "Warranty Criticals" -hn "Health"

cimsubscribe -cs -fn "Processor PFA Normals" -hn "Health"
cimsubscribe -cs -fn "Processor PFA Warnings" -hn "Health"
cimsubscribe -cs -fn "Processor PFA Criticals" -hn "Health"

cimsubscribe -cs -fn "Memory PFA Normals" -hn "Health"
cimsubscribe -cs -fn "Memory PFA Warnings" -hn "Health"
cimsubscribe -cs -fn "Memory PFA Criticals" -hn "Health"

cimsubscribe -cs -fn "Power Supply Normals" -hn "Health"
cimsubscribe -cs -fn "Power Supply Warnings" -hn "Health"
cimsubscribe -cs -fn "Power Supply Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor Error Health Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor Error Health Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor Error Health Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor PFA Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor PFA Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor PFA Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor Remote Healthin Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor Remote Healthin Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor Remote Healthin Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor DASD Backplane Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor DASD Backplane Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor DASD Backplane Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor Generic Fan Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor Generic Fan Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor Generic Fan Criticals" -hn "Health"

cimsubscribe -cs -fn "Service Processor Generic Volatge Normals" -hn "Health"
cimsubscribe -cs -fn "Service Processor Generic Volatge Warnings" -hn "Health"
cimsubscribe -cs -fn "Service Processor Generic Volatge Criticals" -hn "Health"

cimsubscribe -cs -fn "Storage Normals" -hn "Health"
cimsubscribe -cs -fn "Storage Warnings" -hn "Health"
cimsubscribe -cs -fn "Storage Criticals" -hn "Health"
cimsubscribe -cs -fn "SMART Drive Normals" -hn "Health"
cimsubscribe -cs -fn "SMART Drive Warnings" -hn "Health"
cimsubscribe -cs -fn "SMART Drive Criticals" -hn "Health"

cimsubscribe -cs -fn "RAID Normals" -hn "Health"
cimsubscribe -cs -fn "RAID Warnings" -hn "Health"
cimsubscribe -cs -fn "RAID Criticals" -hn "Health"

cimsubscribe -cs -fn "RAID System Health Normals" -hn "Health"
cimsubscribe -cs -fn "RAID System Health Warnings" -hn "Health"
cimsubscribe -cs -fn "RAID System Health Criticals" -hn "Health"

cimsubscribe -cs -fn "Redundant NIC Normals" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Warnings" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Criticals" -hn "Health"

cimsubscribe -cs -fn "Redundant NIC Switchover Normals" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Switchover Warnings" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Switchover Criticals" -hn "Health"

cimsubscribe -cs -fn "Redundant NIC Switchback Normals" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Switchback Warnings" -hn "Health"
cimsubscribe -cs -fn "Redundant NIC Switchback Criticals" -hn "Health"

cimsubscribe -cs -fn "Network Adapter Normals" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Warnings" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Criticals" -hn "Health"

cimsubscribe -cs -fn "Network Adapter Offline Normals" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Offline Warnings" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Offline Criticals" -hn "Health"

cimsubscribe -cs -fn "Network Adapter Online Normals" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Online Warnings" -hn "Health"
cimsubscribe -cs -fn "Network Adapter Online Criticals" -hn "Health"

cimsubscribe -cs -fn "Chassis Normals" -hn "Health"
cimsubscribe -cs -fn "Chassis Warnings" -hn "Health"
cimsubscribe -cs -fn "Chassis Criticals" -hn "Health"

cimsubscribe -cs -fn "LAN Leash Normals" -hn "Health"
cimsubscribe -cs -fn "LAN Leash Warnings" -hn "Health"
cimsubscribe -cs -fn "LAN Leash Criticals" -hn "Health"

##########################################################################################
## PopUp Subscriptions
##########################################################################################
#############################   NONE   ###################################################
```

Appendix E. Discovery

This topic provides reference information about IBM Director discovery.

Discovery operations

This topic provides information about the types of discovery operations that IBM Director supports.

IBM Director supports four types of discovery concerning managed systems and SNMP devices:

Broadcast discovery

Broadcast discovery sends out a general broadcast packet to the local subnet.

Broadcast discovery also can send out a broadcast packet to remote subnets. If you specify the IP address and subnet mask for a system (a discovery seed address), IBM Director sends a broadcast packet to that specific subnet and discovers all managed systems on that subnet that do not filter out broadcast packets.

Multicast discovery

Multicast discovery operates by sending a packet to the multicast address. By default, IBM Director uses the following multicast addresses:

224.0.1.118

Used by IBM Director Server to discover Level-2 managed systems (IBM Director Agent).

239.255.255.253

Used by Service Location Protocol (SLP) to discover Level-1 managed systems (IBM Director Core Services), service processors, BladeCenter chassis, and SMI-S storage devices.

Managed objects monitor the applicable address and respond to the multicast from IBM Director Server. Multicasts are defined with maximum time to live (TTL), which is the number of times a packet is passed between subnets. After the TTL expires, the packet is discarded.

Multicasts are useful for networks that filter broadcasts but do not filter multicasts. Multicast discovery is available only for TCP/IP systems.

Unicast discovery

Unicast discovery sends a directed request to a specific address or range of addresses. This method is useful in networks where both broadcasts and multicasts are filtered. Unicast discovery is available only for TCP/IP systems.

Broadcast relay agents

Broadcast relay enables the server to discover TCP/IP and IPX systems when the systems are not directly reachable by broadcast packets because of network configuration. This situation can occur in networks where the management server and managed systems are in separate subnets and the network between them does not allow broadcast packets to pass from one subnet to the other.

This option generates less network traffic than unicast discovery and avoids many of the problems associated with filtered broadcasts. In broadcast relay, IBM Director Server sends a special discovery request message to a particular managed system, instructing the managed system to perform a discovery on the local subnet using a general broadcast. When managed systems on that subnet receive the discovery request, they reply to the instance of IBM Director Server that made the original request.

The management server performs all types of discovery simultaneously.

Discovering service processors

This topic provides information about discovering service processors.

In IBM Director Console, when you click **Discover** → **Physical Platform**, IBM Director can discover certain service processor models and create physical platform managed objects (PPMOs) for them. This is done using service location protocol (SLP).

For other service processor models, either IBM Director Core Services or IBM Director Agent must be installed. See the following table for details.

Table 42. Discovery options for service processors

Service processor	Discover Physical Platforms	Discover Level 0: Agentless Systems	Discover Level 1: IBM Director Core Services Systems	Discover Level 2: IBM Director Agents
Remote Supervisor Adapter	Yes	Yes	Yes	Yes
Remote Supervisor Adapter II	Yes	Yes	Yes	Yes
IPMI baseboard management controller	No	No	Yes	Yes
ASM PCI Adapter	No	No	Yes	Yes
ASM processor [1]	No	No	No	No
Integrated system management processor (ISMP)	Yes [2]	No	No	No

1. You can discover ASM processors using only IBM Director Agent, versions 4.10, 4.10.2, 4.11, 4.12, 4.20, 4.20.2, 4.21, and 4.22.
2. An ISMP service processor is discovered when an RS-485 gateway is discovered.

Appendix F. FRU data files

IBM Director obtains field-replaceable unit (FRU) data files for use with some tasks.

IBM Director obtains information about the field-replaceable unit (FRU) components that are installed in a managed system from the IBM Support FTP site (ftp://ftp.software.ibm.compc/pccbbs/bp_server). The FRU information is contained in a FRU data file that is:

- Specific to the managed system server model type
- Available only for xSeries servers that currently are supported by IBM

IBM Director makes one attempt to copy the FRU data file:

For managed systems running Linux	The copy occurs during the IBM Director Agent installation on the managed system.
For managed systems running Windows	The copy occurs the first time you restart the managed system after IBM Director Agent is installed.

For the copy to succeed, the managed system must be connected to the network and have firewall access through a standard FTP port. By default, IBM Director attempts to reach the IBM Support FTP site on FTP port 21. After IBM Director successfully copies the FRU data file to the managed system, the FRU data file is processed and the FRU information is stored in the CIM server. Then, IBM Director deletes the FRU data file from the managed system.

Copying FRU data files from the IBM Support FTP site

You can copy FRU data files to your network from the IBM Support FTP site. Use this procedure if a managed system cannot access the IBM Support FTP site.

To copy a FRU data file to your network, complete the following steps:

1. Access the IBM Support FTP site (ftp.software.ibm.com) using the FTP protocol. This FTP site uses an anonymous login.
2. Change to the pc/pccbbs/bp_server directory.
3. Use the `get` command to copy a FRU data file from the IBM Support FTP site to your network. To retrieve a FRU data file, you must provide the applicable FRU data file name. These file names use the following syntax:

   ```
   machine_type_numberums.txt
   ```

   where *machine_type_number* is the machine type number for the managed system. For example, if a server has a machine type number of 1234, the filename is 1234ums.txt. Use the Inventory task to determine the four-digit machine type number of your managed system.

   Note: You can retrieve only one FRU data file at a time. You cannot use the `mget` command.
4. Copy the FRU data file to a server and directory on your network. This server is your internal FTP site repository for your FRU data files. Your FTP site must use an anonymous login.

5. Select the applicable version of the getfru command for the operating system running on your managed system:

Option	Description
For Linux	/opt/IBM/director/CIMOM/bin
For Windows	*c*:\Program Files\IBM\Director\cimom\bin

   where *c* is the drive letter of the hard disk on which IBM Director Agent is installed and IBM Director Agent is installed in the default location.

6. Write a script using the applicable getfru command syntax to retrieve FRU data files from your FTP site.

Option	Description
For Linux	`./getfru -s` *`ftp_server_name`* `-d` *`directory_of_fru_files`*
For Windows	`getfru -s` *`ftp_server_name`* `-d` *`directory_of_fru_files`*

   where:

   - *ftp_server_name* is the FTP address of the network server to which you copied the FRU data files. If you do not specify an address, the command uses a default of ftp.software.ibm.com.
   - *directory_of_fru_files* is the directory on your network server that stores the FRU data files. If you do not specify a directory, the command uses a default of pc/pccbbs/bp_server.

7. Use the Process Management task to run the script to access the FRU data files located on your network. See Viewing and working with processes, services, and device-services information for more information.

Appendix G. IBM Director Agent features

The IBM Director Agent features vary depending on the operating system on which IBM Director Agent is installed.

Management Processor Assistant Agent for NetWare

IBM Director Agent, version 5.10, for NetWare provides the Management Processor Assistant (MPA) Agent feature.

Management Processor Assistant (MPA) Agent works with xSeries and eserver servers that contain one of the following service processors or adapters:

- Advanced System Management processor (ASM processor)
- Advanced System Management PCI adapter (ASM PCI adapter)
- Integrated system management processor (ISMP)
- Intelligent Platform Management Interface (IPMI) baseboard management controller
- Remote Supervisor Adapter
- Remote Supervisor Adapter II

You must install MPA Agent in order to use the MPA task to configure, monitor, and manage the service processors. MPA Agent handles in-band communication between service processors and IBM Director Server. For managed systems running NetWare, if supported by the service processor, MPA Agent handles in-band alert notification.

IBM Director Remote Control Agent

You can use IBM Director Remote Control Agent to perform remote desktop functions on Level-2 managed systems.

From IBM Director Console, you can control the mouse and keyboard of a Level-2 managed system on which IBM Director Remote Control Agent has been installed. This feature is supported only on Windows 32-bit and 64-bit operating systems.

Appendix H. IBM Director extensions

In addition to the standard IBM Director installation, you can add extensions to IBM Director. *Extensions* are tools or plug-ins that extend the functionality of IBM Director. Extensions can be free or available for purchase.

Icon	Extension	What it does
	BladeCenter Management	Configure IBM eServer BladeCenter chassis, management modules, and network devices; start a command-line interface to the management module, start a Web application to view BladeCenter product settings and information, or start vendor software to manage your switches or other network devices.
	Capacity Manager	A resource-management planning tool that you can use to monitor critical resources, such as processor usage, hard disk capacity, memory usage, and network traffic; identify current or potential bottlenecks for an individual server or a group of servers; generate performance-analysis reports and recommends ways to improve performance and prevent diminished performance or downtime; and forecast performance trends.
	Electronic Service Agent™	Monitor your xSeries servers and IBM eServer BladeCenter products for hardware errors. Hardware errors that meet certain criteria are reported to IBM. Electronic Service Agent also administers hardware and software inventory collections, and reports inventory changes to IBM. All information sent to IBM is stored in a secure IBM database and used for improved problem determination.
	Hardware Management Console	View the servers and hardware resources that are managed by the HMC for IBM eServer i5 and eServer p5 models, perform power control, and launch management tools.
	Remote Deployment Manager	Remotely perform configuration, deployment, and retirement operations on both IBM and non-IBM systems. You can use RDM to update system firmware, modify configuration settings, install operating systems and applications, back up and recover primary partitions, and securely erase data from disks.
	ServeRAID Manager	Configure, monitor, and maintain ServeRAID adapters or controllers that are installed locally or remotely on servers. You can view information that is related to controllers, arrays, logical drives, hot-spare drives, and hard disk drives. Also, you can view configuration settings and events and locate defunct hard disk drives.
	Software Distribution (Premium Edition)	Import applications and data, build a software package, and distribute the package to IBM Director managed systems.
	System Availability	Analyze the availability of a managed system or group. You can view statistics about managed-system uptime and downtime through reports and graphical representations.

Icon	Extension	What it does
	Virtual Machine Manager	Manage both physical and virtual machines from a single console. With IBM Virtual Machine Manager (VMM), you can manage virtual components from supported virtualization applications in an IBM Director environment. This includes applications from Microsoft and VMware.
None	Web-based Access	View managed system information, change alert standard format (ASF) alerts, change system settings and configurations, and more. When you install Web-based Access on a managed system, you can access IBM Director Agent and view real-time asset and health information about the managed system from a Web browser. This feature is supported only on Windows 32-bit operating systems.
	z/VM Center	Provision Linux systems on virtual hardware that is based on real System z9 and zSeries hardware and the z/VM hypervisor.

BladeCenter Management

Use the BladeCenter Management task to configure @server BladeCenter chassis, management modules, and network devices; start a command-line interface to the management module, start a Web application to view BladeCenter product settings and information, or start vendor software to manage your switches or other network devices.

BladeCenter Management has the following subtasks:

BladeCenter Configuration Manager
Create or update a BladeCenter chassis configuration profiles.

Network Device Manager
Starts vendor software to manage your switches or other network devices. Depending on the device, a Telnet window, Web interface, or other software interface is started.

Note: This subtask was previously called Switch Management launch pad.

Launch Web Browser
View BladeCenter chassis information via the management module Web page.

Note: You cannot access this task in the Tasks pane. Instead, right-click the managed object on which you want to use the task and select the task from that menu.

Launch Command Line Interface
Starts the Management Processor Command Line Interface (MPCLI) against the selected BladeCenter management module.

Note: You cannot access this task in the Tasks pane. Instead, right-click the managed object on which you want to use the task and select the task from that menu. For documentation of the MPCLI, see www.ibm.com/support/docview.wss?rs=0&uid=psg1MIGR-54214&loc=en_US.

Icon	
Supported IBM Director objects	• (For BladeCenter Configuration Manager) BladeCenter Chassis and Physical Platform managed objects. Also, Ethernet, Fibre, and Infiniband switch modules which are represented in the Groups pane as both Remote managed objects and SNMP device managed objects. • (For Network Device Manager) Ethernet, Fibre, and Infiniband switch modules which are represented in the Groups pane as both Remote managed objects and SNMP device managed objects. • (For Launch Web Browser) BladeCenter Chassis managed objects • (For Launch Command Line Interface) BladeCenter Chassis and Physical Platform managed objects
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Extension to the IBM Director product. The standard IBM Director product installation offers this extension as an optionally installed feature.
Required hardware or hardware limitations	BladeCenter chassis, management modules, network devices, and blade servers.
Required software	(For BladeCenter Configuration Manager) If you use the option to deploy operating systems to your blade servers, Remote Deployment Manager is required.
Required protocols	Telnet
Required device drivers	None
Mass Configuration support	No, but the BladeCenter Configuration Manager subtask provides a feature where you can apply profiles created in the BladeCenter Configuration Manager against one or more managed objects at a time.
Scheduler support	Yes, you can schedule to apply profiles created in the BladeCenter Configuration Manager against one or more managed objects.
Files associated with this task	(For BladeCenter Configuration Manager) An XML file for importing and exporting BladeCenter chassis configuration profile information.
Events associated with this task	• Configuration Manager • Most MPA events For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Capacity Manager

The Capacity Manager task is a resource-management planning tool that you can use to monitor critical resources, such as processor usage, hard disk capacity, memory usage, and network traffic; identify current or potential bottlenecks for an individual server or a group of servers; generate performance-analysis reports and recommends ways to improve performance and prevent diminished performance or downtime; and forecast performance trends.

Capacity Manager has three components:

Monitor Activator
Displays the status of resource and performance-analysis monitors on managed systems; you can specify which monitors are active. When new hardware is detected, corresponding monitors are activated automatically.

Report Generator
Provides Report Definitions, which you can customize for generating reports.

Report Viewer
Provides four views of your generated report data and graphs of monitor performance.

Icon	
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Extension to the standard IBM Director product, which must be purchased separately.
Required hardware or hardware limitations	Designed specifically for use on xSeries and Netfinity servers.
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	Any generated reports that you selected to save to a file. Unless you specify another directory in the report definition, these files are located in the following applicable directory: • (For Linux) /opt/ibm/*director*/*reports* where: • (For Windows) c:\Program Files\IBM\*director**reports* where: • *director* is the default directory in which IBM Director Console is installed. • *reports* is the default directory where the reports are stored and exported by the user.
Events associated with this task	Capacity Manager For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Electronic Service Agent

Use the Electronic Service Agent (IBM Director Extension Edition) to monitor your xSeries servers and @server BladeCenter products for hardware errors. Hardware errors that meet certain criteria are reported to IBM. Electronic Service Agent also administers hardware and software inventory collections, and reports inventory changes to IBM. All information sent to IBM is stored in a secure IBM database and used for improved problem determination.

Electronic Service Agent includes the following features:

- Places service calls to IBM automatically if the server is under a service agreement or warranty.
- Collects and sends scheduled system inventory and diagnostic inventory to an IBM database. This inventory information is available to IBM support representatives when they are solving your problem.
- Provides problem-definable threshold levels for error reporting.
- Communicates with IBM using a secure Internet connection using encryption and authentication, or by using a dial-up connection.

For more information, see the following Web site and documentation:

- Electronic Service Agent Web site at www.ibm.com/support/electronic/. Click **xSeries** for the applicable information.
- *IBM Electronic Service Agent on xSeries - Director Extension User's Guide*. This document is available from the Electronic Service Agent Web page. See the following table for availability.
- *IBM Electronic Services: Support for Business in an On Demand World*. This is an IBM Redbooks document that is available at www.ibm.com/redbooks/.

Icon	
Supported IBM Director objects	Level-2 managed systems **Note:** You can use the Electronic Service Agent (Standalone Edition) to monitor Level-0 and Level-1 managed systems. The Standalone Edition is not integrated into IBM Director Console. You can download this edition of the product from the Electronic Service Agent Web page at www.ibm.com/support/electronic/. Click **xSeries** for the applicable information.
Supported operating systems	• Linux • Windows For detailed operating-system support information, see the Electronic Service Agent documentation.
Availability	Extension to the IBM Director product. You can download the extension from the IBM Electronic Service Agent Web page.
Required hardware or hardware limitations	Designed specifically for use on xSeries servers and IBM eServer BladeCenter products. For a detailed list of supported servers and products, see the Electronic Service Agent documentation.
Required software	• (For Windows) AT&T Global Network Services • (For Linux) WvDial
Required protocols	HTTPS for connecting to the Internet.
Required device drivers	No

Mass Configuration support	No
Scheduler support	Yes
Files associated with this task	None
Events associated with this task	Yes, for a detailed list of events see the Electronic Service Agent documentation.

Hardware Management Console

Use the Hardware Management Console (HMC) task to view the servers and hardware resources that are managed by the HMC for IBM eServer i5 and eServer p5 models, perform power control, and launch management tools.

Icon	
Supported IBM Director objects	Level-1 and Level-2 managed systems
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Extension to the IBM Director product. It is available on CDs that come with the applicable iSeries and pSeries products. Also, you can download the extension from the IBM Support Web site at www.ibm.com/servers/eserver/xseries/systems_management/ibm_director/. **Note:** (iSeries only) IBM Director Server already includes the HMC extension. The iSeries CD provides the IBM Director Console HMC extension code.
Required hardware or hardware limitations	IBM eServer i5 and eServer p5 models
Required software	CIM
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	• CIM > System > IP Changed • CIM > System > Life Cycle • CIM > System > Power State • CIM > System > Service Event For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

IBM Remote Deployment Manager

Use the IBM Remote Deployment Manager (RDM) extension to remotely perform configuration, deployment, and retirement operations on both IBM and non-IBM systems. You can use RDM to update system firmware, modify configuration settings, install operating systems and applications, back up and recover primary partitions, and securely erase data from disks.

For more information, including the product documentation, see the IBM Support Web site at www.ibm.com/servers/eserver/xseries/systems_management/ibm_director/

Icon	
Supported IBM Director objects	Level-0, Level-1, and Level-2 managed systems are supported as distribution endpoints
Supported operating systems	• Linux • Windows For detailed operating-system support information, see the Remote Deployment Manager documentation.
Availability	Extension to the standard IBM Director product, which must be purchased separately. See the Remote Deployment Manager Web page for purchasing information.
Required hardware or hardware limitations	Yes, for detailed hardware support information see the *IBM Remote Deployment Manager Compatibility Guide*.
Required software	None
Required protocols	Yes, for detailed information see the Remote Deployment Manager documentation.
Required device drivers	None
Mass Configuration support	No
Scheduler support	Yes
Files associated with this task	Yes, for detailed information see the *IBM Remote Deployment Manager User's Guide*
Events associated with this task	None

ServeRAID Manager

Use the ServeRAID Manager task to configure, monitor, and maintain ServeRAID adapters or controllers that are installed locally or remotely on servers. You can view information that is related to controllers, arrays, logical drives, hot-spare drives, and hard disk drives. Also, you can view configuration settings and events (which are called *notifications* in the ServeRAID Manager task) and locate defunct hard disk drives.

Note: The ServeRAID Manager task for IBM Director is not the same program as the ServeRAID Manager (Standalone Edition) that is provided with the ServeRAID hardware option. It is recommended that you not install both versions on the same system.

Icon	
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Extension to the IBM Director product. You can download the extension from the IBM Support Web site at www.ibm.com/servers/eserver/xseries/systems_management/ibm_director/.
Required hardware or hardware limitations	Designed specifically for use on xSeries and Netfinity servers. The following adapters or controllers must be installed locally or remotely on these servers: • ServeRAID adapters • Integrated SCSI controllers with RAID capabilities • Serial ATA controllers with integrated RAID • Ultra320 SCSI controllers with integrated RAID
Required software	None
Required protocols	None
Required device drivers	Applicable ServeRAID device drivers that support ServeRAID hardware.
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	The ServeRAID Manager task for IBM Director generates the following events: • The ServeRAID events that are contained under the CIM > System event type. • (SNMP events under iso) The iBMServeRAID events that are contained under the ibmSystemMIB event type. • Storage > ServeRAID Controller **Note:** The ServeRAID Manager (Standalone Edition) generates the events that are contained under the SNMP > Hardware > Storage > RAID event type. For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Software Distribution (Premium Edition)

Use the Software Distribution (Premium Edition) task to import applications and data, build a software package, and distribute the package to IBM Director managed systems.

In addition to the function provided by IBM Director Software Distribution (Standard Edition), Software Distribution (Premium Edition), has the following additional capabilities:

- Import non-IBM or IBM software and build software packages that use the following wizards:
  - InstallShield Package wizard (Windows)
  - Microsoft Windows Installer wizard (Windows)
  - RPM Package wizard (AIX and Linux)
  - AIX InstallP wizard (AIX)
- Import non-IBM or IBM software and build a software package by using the Custom Package Editor
- Import a software package created in IBM Director by using the Previously Exported Package wizard
- Export a software package for use on another management server or as a backup
- Restore i5/OS libraries, objects and installed programs

Icon	
Supported IBM Director objects	Level-0, Level-1, and Level-2 managed systems as distribution endpoints. **Note:** Only packages in the Solution Install format can be distributed to Level-0 and Level-1 managed systems.
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Software Distribution (Premium Edition) must be purchased separately and installed on the management server.
Required hardware or hardware limitations	None
Required software	None
Required protocols	Default protocol is TCP. If you disable TCP-session support on a managed system, Software Distribution uses UDP.
Required device drivers	None
Mass Configuration support	No
Scheduler support	Yes
Files associated with this task	Depending on the operating system and the wizard that you use, the following files are associated with this task: • (For Director Update Assistant) The software-update file and an XML file that describes the software-update file • (For Solution Install) JAR or ZIP files or Solution Install data using the Solution Install directory structure. This does not include archive format. Solution Install packages contain two XML files that describe the package. • (For Linux and AIX) RPM file • (For AIX InstallP) BFF file • (For Microsoft Windows Installer) MSI file • (For InstallShield) setup.exe and the ISS response file • (For exported packages) SPB file • (For i5/OS) Library, licensed programs, or objects
Events associated with this task	None

System Availability

Use the System Availability task to analyze the availability of a managed system or group. You can view statistics about managed-system uptime and downtime through reports and graphical representations.

System Availability can identify problematic managed systems that have had too many unplanned outages over a specified period of time, a managed system that has availability data that is too old, or a managed system that fails to report data to IBM Director Server. When a system-availability report is generated, managed systems that meet the criteria that you specify as being problematic are flagged as such. You can run the System Availability task on a managed system or group immediately or schedule a System Availability task using the Scheduler task.

Icon	🖥
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Extension to the IBM Director product. You can download the extension from the IBM Support Web site at www.ibm.com/servers/eserver/xseries/systems_management/ibm_director/.
Required hardware or hardware limitations	Designed specifically for use on xSeries and Netfinity servers.
Required software	To use the function that identifies a managed system as problematic, IBM Director (version 4.1 or later) System Availability Agent must be installed on the managed system.
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	Yes
Files associated with this task	• (Windows only) The System Availability task uses information from the system log file; a damaged, missing, or full system log file affects this task. If you clear the system log file, all system-availability information is lost. • (Linux only) The System Availability task uses information from the /var/log/messages file. • IBM Director Server stores system-availability reports in the IBM\Director\Reports\System Availability directory on the management server. You can change the location where IBM Director Server stores system-availability reports in the Settings window.
Events associated with this task	System Availability For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Virtual Machine Manager

Use the Virtual Machine Manager (VMM) to manage both physical and virtual machines from a single console. With IBM Virtual Machine Manager (VMM), you can manage virtual components from supported virtualization applications in an IBM Director environment. This includes applications from Microsoft and VMware.

For more information, see the *IBM Virtual Machine Manager Installation and User's Guide*. This document is available from the Virtual Machine Manager Web page. See the following table for availability.

Icon	⧉
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	• Linux • Windows For detailed operating-system support information, see the Virtual Machine Manager documentation.
Availability	Extension to the IBM Director product. You can download the extension from the IBM Support Web site at www.ibm.com/servers/eserver/xseries/systems_management/ibm_director/.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	Yes
Files associated with this task	None
Events associated with this task	Yes, for a detailed list of events see the Virtual Machine Manager documentation.

Web-based Access

Use Web-based Access to view managed system information, change alert standard format (ASF) alerts, change system settings and configurations, and more. When you install Web-based Access on a managed system, you can access IBM Director Agent and view real-time asset and health information about the managed system from a Web browser. This feature is supported only on Windows 32-bit, Windows XP 64-bit, and Windows 2003 64-bit operating systems.

Web-based Access is useful in the following situations:

- You do not want to install IBM Director Console.
- You plan to manage only a few servers, desktop computers, or other devices.
- You want to remotely access managed systems when using a Web browser.

- You want to view the most up-to-date information about the assets, and operating-system state of a managed system.

Icon	None
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	Windows 32-bit, Windows XP 64-bit, and Windows 2003 64-bit operating systems. For detailed operating-system support information, see the *Web-based Access Installation and User's Guide*, which is available on the Web at:www.ibm.com/servers/eserver/xseries/ systems_management/ibm_director/.
Availability	Extension to the IBM Director product. You can download the extension from the IBM Director Web site at www.ibm.com/servers/eserver/xseries/ systems_management/ibm_director/.
Required hardware or hardware limitations	None
Required software	Yes, for detailed Web-browser, JVM, and Java Foundation Class/Swing library (JFC/Swing) support information see the *Web-based Access Installation and User's Guide*, which is available on the Web at:www.ibm.com/servers/eserver/xseries/ systems_management/ibm_director/
Required protocols	HTTP and HTTPS for connecting to the Internet.
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	None

z/VM Center

Use z/VM Center to provision Linux systems on virtual hardware that is based on real IBM System z9 and @server zSeries hardware and the z/VM hypervisor.

z/VM Center provides the following tasks:

Virtual Server Deployment

With the Virtual Server Deployment task, you can define configurations of z/VM virtual servers and save them as virtual server templates.

Use the Virtual Server Deployment task to manage individual z/VM virtual servers and operating system instances and to set up templates and Linux guest systems.

A Linux guest system is a combination of a Linux instance and the z/VM virtual server on which the Linux instance is installed.

Server Complexes

With the Server Complexes task you can manage configurations of Linux guest systems. A server complex is a configuration profile for Linux guest systems and includes both Linux and z/VM aspects. A server complex can define network settings, Linux configuration scripts, disk access, and VM Resource Manager (VMRM) performance goals.

You can automatically configure a Linux guest system by assigning it to a server complex. You can also create a new Linux guest system within a server complex. When creating a new Linux instance in a server complex, you automatically create a z/VM virtual server with a Linux instance that is configured according to the server complex. For creating a Linux guest system, you require a virtual server template and an operating system template that have been created by the Virtual Server Deployment task.

You can make changes to a server complex and then apply the configuration changes to all Linux instances in the server complex.

Use Server Complexes to manage numerous Linux instances with similar configurations.

Icon	**(z/VM Center)** (Virtual Server Deployment) (Server Complexes)
Supported IBM Director objects	Level-0, Level-1, and Level-2 managed systems The z/VM Manageability Access Point must be a Level-2 managed system.
Supported operating systems	The following Linux distributions: • SUSE LINUX Enterprise Server 9 for IBM System z9, zSeries and S/390 • Red Hat Enterprise Linux AS, version 4.0, for IBM System z9, zSeries and S/390 running as a guest operating system on either of: • z/VM 5.2 • z/VM 5.1, with the PTFs for APAR VM63804. For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Availability	Extension to the IBM Director product. It is available on CDs that come with the applicable System z9 and zSeries and iSeries products. Also, you can download the extension from the IBM Support Web site at www.ibm.com/servers/eserver/xseries/systems_management/ibm_director/ To be able to use z/VM Center, you must purchase the IBM Director Extensions, V5.10 feature of IBM Virtualization Engine and Infrastructure Services for Linux on System z9 and zSeries, V2.1. For more details see the z/VM Center setup information in the IBM Director information center at http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/vsd0_t_prepare.html.
Required hardware or hardware limitations	IBM System z9 or eServer zSeries
Required software	• The IBM Enterprise Storage Server® FlashCopy2 is recommended to enhance performance • A z/VM Directory Maintenance Product, for example: DirMaint FL510 with the following APARs: – PTF for APAR VM 63700 for IBM DirMaint FL510 to support 1 TB FCP-attached SCSI disks – PTF for APAR VM 63733 for IBM DirMaint – PTF for APAR VM VM63639 for IBM DirMaint • CPINT – The CPINT RPM is shipped with SUSE LINUX Enterprise Server 9. – For Red Hat Enterprise Linux AS, you can download the RPM from linuxvm.org/Patches/. You need version 2.5.3 or later. **Note:** Be aware that installing the RPM on Red Hat Linux AS might affect any support contract you may have for the distribution. For the latest information see the IBM Director 5.10 information center at http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	z/VM Center includes a file, zVMPersonalization*version*.rpm, that you must install in a Linux to create a master Linux. A master Linux is the starting point for cloning Linux instances. *version* represents a file version specification.
Events generated by this task	None

Virtual Server Deployment

Virtual Server Deployment is a subtask of z/VM Center. Virtual Server Deployment builds on z/VM to create and decommission virtual Linux servers (Linux guest systems).

This topic provides a brief introduction to the main objects you work with in the Virtual Server Deployment task. It also gives an overview of the task flow for deploying a z/VM virtual server that has an operating system installed. For more details see http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/vsd0_c_concepts.html.

Virtual server templates

A virtual server template defines the characteristics of a z/VM guest virtual machine (virtual hardware on z/VM). The template includes, for example, the processing power and memory size of the guest virtual machine. You use a virtual server template for creating a z/VM virtual server.

z/VM virtual servers

To z/VM, a z/VM virtual server is a z/VM guest virtual machine. To set up an operating system in a z/VM virtual server, you use an operating system template.

Operating system templates

You create an operating system template from an existing instance of an operating system. An operating system template includes:

- Disks that need to be unique for each operating system instances that is based on the template
- Disks that are read-only and are to be shared by all operating system instances that are based on the template
- Information on network interfaces configured for use by operating system instances that are based on the template

When you use an operating system template to set up an operating system instance in a z/VM virtual server, a wizard guides you through specifying a small amount of data that needs to be unique for each operating system instance, for example, IP addresses.

Decommissioning a z/VM virtual server

When a z/VM virtual servers is no longer needed, you can delete it to free all resources that are used by it for redeployment.

Server Complexes task

The Server Complexes task is a subtask of the z/VM Center extension to IBM Director 5.10. By creating and using server complexes, you can control the configuration of virtual Linux guests in a z/VM, in an automatic fashion.

This topic provides a brief introduction to the main objects you work with in the Server Complexes task and the functions you can perform using these objects.

Linux guest systems

A Linux guest system is an IBM Director managed object. This object represents a Linux system running as a guest operating system in a z/VM virtual server. From the point of view of IBM Director, a Linux guest system is simply a Linux system running an IBM Director agent.

Server complexes

A server complex is a virtual complex that is used to configure the properties of Linux guest systems in a z/VM. This includes both the z/VM side and the Linux side configuration. Within the context of a managed z/VM, you can create as many server complexes as required.

Tier

A tier is a subsection within a server complex. You can divide a server complex into multiple tiers, each with their own set of properties. You can use tiers to group Linux guest systems according to their functionality.

Server complex properties

Server complexes have four configuration domains: VMRM, minidisks, scripts, and network. You can set properties for some or all of these domains and apply these properties to Linux guest systems within the server complex. You can also set the properties of a single tier within a server complex.

Provisioning resources

Provisioning resources are the virtual server templates, operating system templates, and disk pools that can be used in the cloning process. These resources are maintained in the VSD task. Cloning is the process of creating a Linux guest system and configuring it according to server complex properties.

Appendix I. IBM Director tasks

The standard IBM Director installation provides tasks that you can use alone or in combination to work with managed objects in your systems-management environment. Managed objects include, but are not limited to, Level-0, Level-1, Level-2 managed systems, SNMP devices, BladeCenter management modules, platforms, and network devices.

You also can add extensions to IBM Director to extend the functionality of IBM Director.

Note: Tasks that are available from the menu and are not displayed in the Tasks pane do not have an icon.

Icon	Task	What it does
	Active Console Viewer	View who is currently logged in to the IBM Director Server. Start the Active Console Viewer to display a list of all active consoles attached to the server, which includes the user ID logged in, the host name of the system that the user logged in from, and the date and time the user logged in. The window is automatically updated as users log into and log out of the server.
	Asset ID	View lease, warranty, user, and system information, including hard disk drive serial numbers, system serial numbers, and system board serial numbers. You also can use Asset ID to create personalized data fields to add custom information.
	CIM Browser	View detailed information in the CIM layer.
	Configure Alert Standard Format	Configure Alert Standard Format (ASF) on managed systems that contain ASF-capable network interface cards (NICs).
	Configure SNMP Agent	Configure SNMP devices and agents for communication in your IBM Director environment.
	Discovery	Identify and establish connections with all managed objects in the IBM Director environment. The management server sends out a discovery request and waits for responses from all managed objects. The managed objects listen for this request and respond to the management server that sent the request.
	Event Action Plans	Specify actions that are performed in response to events that are generated by a managed object.
	Event Log	View details about all events or subsets of events that have been received and logged by IBM Director Server. You can view all events or view events for a managed object or by filter criteria.
	External Application Launch	Add applications to the IBM Director Console Tasks pane. You can add third-party applications, scripts, and command-line commands to the Tasks pane so you can start them from IBM Director Console.

Icon	Task	What it does
	File Transfer	Send files from one location to another and to synchronize files, directories, or drives. The File Transfer task is a secure alternative to FTP.
	Hardware Status	View details about all events or subsets of events that have been received and logged by IBM Director Server. You can view all events or view events for a managed object or by filter criteria.
	Inventory	Collect data about the hardware and software that is currently installed on the managed objects in your network.
None	Mass Configuration	Runs a single task on a group of managed objects.
	Message Browser	View messages sent to IBM Director Console as a result of an event action plan. The Message Browser is displayed automatically whenever a message is sent to the management console. You can choose to be notified in this manner when an event occurs by configuring an event action plan with the Send an Event Message to a Console User event action.
	Microsoft Cluster Browser	View the structure, nodes, and resources that are associated with a Microsoft Cluster Server (MSCS) cluster. You can determine the status of a cluster resource and view the associated properties of the cluster resources. The Microsoft Cluster Browser does not display the status of a cluster as a whole but displays the individual cluster resource statuses.
	Network Configuration	View and edit settings for Ethernet adapters, IP addresses, Domain Name Systems (DNS) configurations, Windows Internet Naming Service (WINS) configurations, Windows domains and workgroups, and modems of a managed system.
	Process Management	Manage individual processes on managed systems. Specifically, you can start, stop, and monitor processes and set up process monitors to generate events whenever applications change state. You can issue commands on managed systems also.
	Rack Manager	Group your equipment in rack suites. Using Rack Manager, you can create virtual racks by associating equipment, such as managed systems and devices, networking devices, power devices, and monitors, with a rack in order to visually represent an existing rack in your environment. If the inventory-collection function of IBM Director does not recognize a managed system or device in Rack Manager, you can associate it with a predefined component of a similar size.
	Remote Control	Manage a remote system by displaying the screen image of the managed system on a management console. You can cut, copy, and paste text on both the managed system and the management console.
	Remote Session	Run command-line programs on a remote managed system. Remote Session creates less network traffic and uses fewer system resources than the Remote Control task and therefore is useful in low-bandwidth situations.

Icon	Task	What it does
	Resource Monitors	View statistics about critical system resources, such as processor, disk, and memory usage. With resource monitors, you also can set thresholds to detect potential problems with managed systems or devices. When a threshold is met or exceeded, an event is generated. You create event action plans to respond to resource-monitor events. You can apply resource monitors to individual managed systems and devices and to groups.
	Scheduler	Run a single noninteractive task or set of noninteractive tasks at a later time. You can specify an exact date and time you want the task to be started, or you can schedule a task to repeat automatically at a specified interval. Scheduled tasks are referred to as *jobs*.
	Server Configuration Manager	Create or update server configuration profiles. Configuration includes the service processors in IBM eServer xSeries servers.
	SNMP Browser	View and configure the attributes of SNMP devices, such as hubs, routers, or other SNMP-compliant management devices. This task is useful for SNMP-based management, troubleshooting, or monitoring the performance of SNMP devices.
	Software Distribution (Standard Edition)	Import applications and data, build a software package, and distribute the package to IBM Director managed systems. With IBM Director Software Distribution (Standard Edition), you can import only software that is distributed by IBM and build a software package that uses only the IBM Update Assistant wizard.
	System Accounts	View and change user and group security profiles on managed systems.

Active Console Viewer

Use the Active Console Viewer to view who is currently logged in to the IBM Director Server. Start the Active Console Viewer to display a list of all active consoles attached to the server, which includes the user ID logged in, the host name of the system that the user logged in from, and the date and time the user logged in. The window is automatically updated as users log into and log out of the server.

Icon	
Supported IBM Director objects	Not applicable
Supported operating systems	Not applicable
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None

Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	None

Asset ID

Use the Asset ID task to view lease, warranty, user, and system information, including hard disk drive serial numbers, system serial numbers, and system board serial numbers. You also can use Asset ID to create personalized data fields to add custom information.

Icon	✓
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation. A comparable service is also available in the Web-based Access feature.
Required hardware or hardware limitations	None, although if a system is EEPROM-enabled the data is written to the EEPROM as well as to a file. If the system is not EEPROM-enabled the data is written to a file only. Systems that have EEPROM include, but are not limited to, NetVista and ThinkPad computers.
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	Yes
Scheduler support	No
Files associated with this task	• (For Windows) c:\Program Files\IBM\Director\data\asset.dat • (For Linux) /opt/ibm/director/cimom/data/asset.dat • (For i5/OS) /Qibm/UserData/Director/data/asset.dat

Events associated with this task	• CIM > System > Lease Expiration • CIM > System > Warranty Expiration • (SNMP events under iso) ibmSystemMIB > ibmSystemLeaseExpiration • (SNMP events under iso) ibmSystemMIB > ibmSystemWarrantyExpiration For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

CIM Browser

Use the CIM Browser task to browse detailed information in the CIM layer.

Icon	
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	An installed CIMOM on a managed system. The IBM Director CIM Agent detects the CIMOM and uses it to provide data through the CIM Browser.
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	CIM

Configure Alert Standard Format

Use the Configure Alert Standard Format task to configure Alert Standard Format (ASF) on managed systems that contain ASF-capable network interface cards (NICs).

After ASF is configured, IBM Director can perform some or all of the following actions on the managed system:

- Power on
- Power off
- Restart system

- Receive PET events

The actions that IBM Director can perform depend on the level of ASF supported on the NIC in the managed system. In IBM Director Console, ASF-capable systems are represented in the Systems with ASF and the Systems with ASF Secure Remote Management groups.

Before a managed system is recognized by IBM Director Server as ASF-capable, the Inventory task must be run on the managed system. If the managed system supports ASF 1.0, IBM Director Server adds it to the Systems with ASF group. If the managed system supports ASF 2.0, IBM Director Server adds it to both the Systems with ASF group and the Systems with ASF Secure Remote Management group.

Icon	
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	Managed systems that contain ASF-capable network interface cards (NICs) installed with the applicable device drivers.
Required software	None
Required protocols	None
Required device drivers	The applicable device drivers for ASF-capable NICs.
Mass Configuration support	Yes
Scheduler support	No
Files associated with this task	None
Events associated with this task	PET For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Configure SNMP Agent

Use the Configure SNMP Agent task to configure SNMP agents for communication in your IBM Director environment.

Icon	
Supported IBM Director objects	Level-2 managed systems

Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	SNMP must be installed on the managed system or device.
Required device drivers	None
Mass Configuration support	Yes
Scheduler support	No
Files associated with this task	None
Events associated with this task	SNMP For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Discovery

Use Discovery to identify and establish connections with all managed objects in the IBM Director environment. Managed objects include, but are not limited to, Level-0, Level-1, Level-2 managed systems, SNMP devices, BladeCenter management modules, platforms, switches, and z/VM systems. The management server sends out a discovery request and waits for responses from all managed objects. The managed objects listen for this request and respond to the management server that sent the request.

Note: Discovery is a background process. When you start this task, no window or progress indicator is displayed. Because it is a background process, you can use IBM Director Console while the discovery continues. As managed objects are discovered in the IBM Director environment, they are displayed in IBM Director Console.

Icon	🔦
Supported IBM Director objects	All managed objects
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None

Required device drivers	None
Mass Configuration support	No
Scheduler support	Yes
Files associated with this task	None
Events associated with this task	None

Event Action Plans

Use Event Action Plans to specify actions that are performed in response to events that are generated by a managed object. Managed objects include, but are not limited to, Level-0, Level-1, Level-2 managed systems, SNMP devices, BladeCenter management modules, platforms, and switches.

An event action plan is composed of two types of components:
- One or more event filters, which specify event types and any related parameters
- One or more event actions, which occur in response to filtered events

You can apply an event action plan to an individual managed object, several managed objects, or a group of managed objects.

By creating event action plans and applying them to specific managed objects, you can be notified by e-mail or pager, for example, when a specified threshold is reached or a specified event occurs. Or you can configure an event action plan to start a program on a managed object and change a managed-object variable when a specific event occurs. You can use any event, including process-monitor and resource-monitor events, to build an event action plan.

When you install IBM Director, a single event action plan is already defined, in addition to any that you created using the Event Action Plan wizard. The Log All Events event action plan has the following characteristics:
- It uses the event filter named All Events, a simple event filter that processes all events from all managed objects.
- It performs the action Add to the Event Log, a standard event action that adds an entry to the IBM Director Server event log.

Icon	⚙
Supported IBM Director objects	All managed objects
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None

Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	All events in an IBM Director environment are available for use in this task. For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Event Log

Use Event Log to view details about all events or subsets of events that have been received and logged by IBM Director Server. You can view all events or view events for a managed object or by filter criteria.

Icon	⚠
Supported IBM Director objects	All managed objects
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	This task displays all events generated by managed objects, software, and other IBM Director tasks. For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

External Application Launch

Use the External Application Launch task to add applications to the IBM Director Console Tasks pane. You can add third-party applications, scripts, and command-line commands to the Tasks pane so you can start them from IBM Director Console. The application, script, or command-line commands must be installed on the management console (the system running IBM Director Console).

Icon	<>
Supported IBM Director objects	Varies by the application that you add to the IBM Director Console Tasks pane. You can add applications that run on a managed object, for example, the Storage Manager application supports the DS4000 Series device managed object. Other applications are not associated with managed objects or IBM Director, but you can include them in IBM Director Console for convenience. These applications run on the management console only, for example, Telnet or operating system commands.
Supported operating systems	• Linux • Windows For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director product.
Required hardware or hardware limitations	Varies by the application that you add to the IBM Director Console Tasks pane.
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	Command task files that have the CMDExt file extension. These files are located in the following directory on the management server: • (For Linux) /opt/ibm/director/classes/extensions • (For Windows) c:\Program Files\IBM\Director\classes\extensions where these are the default directory paths in which IBM Director Server is installed. For information about creating command task files, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Events associated with this task	None

File Transfer

Use the File Transfer task to send files from one location to another and to synchronize files, directories, or drives. The File Transfer task is a secure alternative to FTP.

You can transfer individual files and directories between the following systems:
- The management console and the management server
- The management console and a managed system
- The management server and a managed system

File transfer between two managed systems is not supported directly. However, you can transfer a file from one managed system to a management console or management server and then transfer that file to a different managed system.

Icon	📁
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	None

Hardware Status

Use the Hardware Status task to view status information about managed-system hardware or to specify that IBM Director ignore all or certain hardware events for a managed system.

IBM Director Console provides an overview of managed-system hardware status. The Groups pane contains three hardware-status groups identified by distinct icons:
- Hardware Status Critical
- Hardware Status Information
- Hardware Status Warning

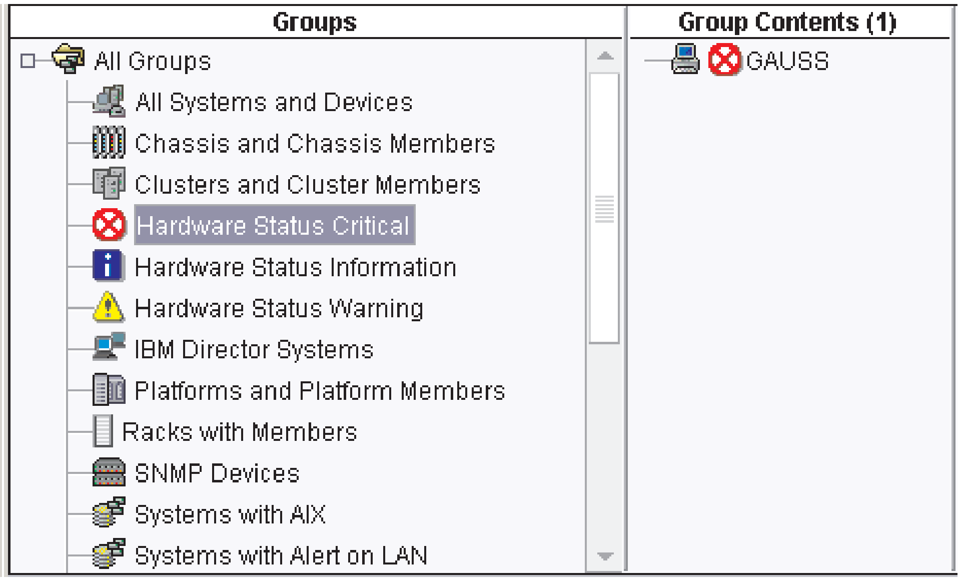

Figure 12. IBM Director Console displaying hardware-status groups

When a managed system generates a hardware event, IBM Director adds the system to the applicable hardware-status group. If a system generates multiple events with different severity levels, IBM Director adds the system to the group that reflects the most severe event. For example, if a managed system generates both a warning and a critical event, IBM Director adds the managed system to the Hardware Status Critical group. When you click a hardware-status group, the managed systems in that group are displayed in the Group Contents pane. The applicable hardware-severity icon is displayed next to the managed system.

Hardware status information also is presented in the lower-right corner of IBM Director Console, below the ticker tape.

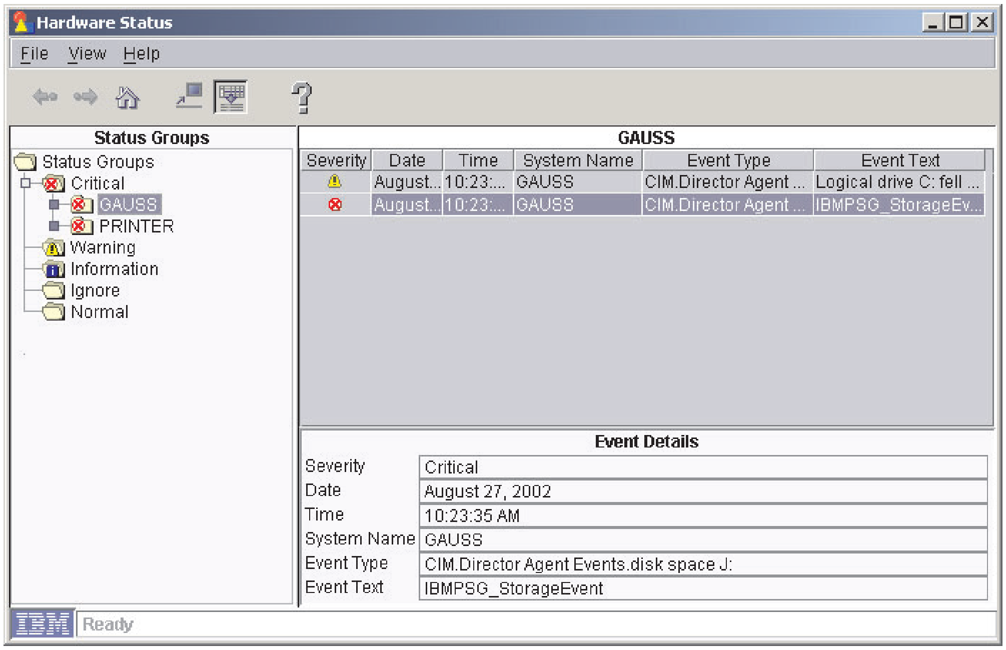

Severity	Date	Time	System Name	Event Type	Event Text
⚠	August...	10:23:...	GAUSS	CIM.Director Agent ...	Logical drive C: fell ...
⊗	August...	10:23:...	GAUSS	CIM.Director Agent ...	IBMPSG_StorageEv...

Severity	Critical
Date	August 27, 2002
Time	10:23:35 AM
System Name	GAUSS
Event Type	CIM.Director Agent Events.disk space J:
Event Text	IBMPSG_StorageEvent

Figure 13. IBM Director Console with hardware-status icons located in the bottom-right portion

The icons for the hardware-status groups are displayed, along with the number of managed systems that are in the hardware-status groups. If a hardware-status group does not contain any managed systems, its icon is unavailable.

Icon	
Supported IBM Director objects	Level-1 and Level-2 managed systems
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	This task is supported (although the support might be limited) whenever out-of-band events generated by a service processor or in-band events generated by CIM are supported on a server.
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	• CIM > System • MPA

Inventory

Use the Inventory task to collect data about the hardware and software that is currently installed on the managed objects in your network.

IBM Director can collect inventory data when a managed object is discovered initially and at regular intervals, or, rather than collecting inventory upon initial discovery, you can schedule an inventory collection at a more convenient time using the Scheduler feature. The default interval for refreshing the database is every 7 days. You can change the refresh interval and other inventory-collection parameters using the Inventory Collection Preferences page in the IBM Director Console Server Preferences window. You also can collect inventory data on a managed system or group immediately or schedule an inventory collection using the Scheduler task.

You can query the inventory database to display details about properties of a managed system, such as remaining disk space. You can use a standard query that is provided or create your own custom query.

You can use the inventory-software dictionary to track the software that is installed on your managed systems. You do not specify drives or directories that you want the Inventory task to search during the software-inventory collection process; the software-dictionary file contains predefined software profiles that recognize most standard software packages after they are installed. When you install software applications on servers, computers, or devices, the inventory-query browser displays the new software after the next inventory collection. If you have installed software that does not correspond to a predefined software profile that is included with IBM Director, you can edit the software-dictionary file to update your software inventory. Software with no predefined software profile includes software that is developed internally in your organization or a new version of software that is released after this version of IBM Director.

Icon	🔍
Supported IBM Director objects	All managed objects including all levels of managed systems.
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	Yes
Files associated with this task	None

Events associated with this task	Director > Inventory For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Mass Configuration

Use the Mass Configuration task to run a single task on a group of managed objects.

The following tasks are supported by Mass Configuration:
- Asset ID
- Configure Alert Standard Format
- Configure SNMP Agent
- Network Configuration

Icon	None
Supported IBM Director objects	All managed objects including all levels of managed systems.
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	Not applicable
Scheduler support	No
Files associated with this task	None
Events associated with this task	Mass Configuration For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Message Browser

Use the Message Browser task to view messages sent to IBM Director Console as a result of an event action plan. The Message Browser is displayed automatically whenever a message is sent to the management console. You can choose to be notified in this manner when an event occurs by configuring an event action plan with the Send an Event Message to a Console User event action.

Icon	✉

Supported IBM Director objects	Level-1 and Level-2 managed systems
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	None, but it does display messages sent as a result of an event-action plan.

Microsoft Cluster Browser

Use the Microsoft Cluster Browser task to view the structure, nodes, and resources that are associated with a Microsoft Cluster Server (MSCS) cluster. You can determine the status of a cluster resource and view the associated properties of the cluster resources. The Microsoft Cluster Browser does not display the status of a cluster as a whole but displays the individual cluster resource statuses.

Icon	
Supported IBM Director objects	Cluster managed objects
Supported operating systems	Windows only. For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director product.
Required hardware or hardware limitations	IBM Cluster hardware
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No

Files associated with this task	None
Events associated with this task	None

Network Configuration

Use the Network Configuration task to view and edit settings for Ethernet adapters, IP addresses, Domain Name Systems (DNS) configurations, Windows Internet Naming Service (WINS) configurations, Windows domains and workgroups, and modems of a managed system.

Icon	
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	Yes
Scheduler support	No
Files associated with this task	None
Events associated with this task	None

Process Management

Use the Process Management task to manage individual processes on managed systems. Specifically, you can start, stop, and monitor processes and set up process monitors to generate events whenever applications change state. You can issue commands on managed systems also.

Icon	
Supported IBM Director objects	Level-2 managed systems and devices
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	• No required hardware. • Cannot use this task on platforms or BladeCenter chassis. • (SNMP devices only) Can view processes on SNMP devices, but cannot affect the processes. • (SNMP printers only) Not supported.
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	Yes
Files associated with this task	None
Events associated with this task	Director > Director Agent > Process Monitors For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Rack Manager

Use the Rack Manager task to group your equipment in rack suites. Using Rack Manager, you can create virtual racks by associating equipment, such as managed systems and devices, networking devices, power devices, and monitors, with a rack in order to visually represent an existing rack in your environment. If the inventory-collection function of IBM Director does not recognize a managed system or device in Rack Manager, you can associate it with a predefined component of a similar size.

You can use Rack Manager to view hardware-status alerts that occur on managed systems or devices in a rack. If a rack component has a hardware-status alert, the rack component is outlined in red, blue, or yellow, depending on the severity level.

Icon	▥
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	Designed specifically for use on xSeries and Netfinity servers.
Required software	None
Required protocols	None
Required device drivers	None

Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	None

Remote Control

Use the Remote Control task to manage a remote system by displaying the screen image of the managed system on a management console. You can cut, copy, and paste text on both the managed system and the management console.

Remote Control has three control modes:

Active Remote-control mode. A management console controls the managed system, and the user of the managed system loses all use of the keyboard and mouse. Only one management console can be in control of a managed system in the active state; all other attached management consoles can monitor the managed-system display only.

Monitor
View-only mode. A management console that is attached to the managed system display the screen image and cursor movements of the managed system.

Suspend
View-only mode without image refresh. A management console that is attached to the managed system displays only the screen image of the managed system. The screen image that is displayed on the management console does not change when the screen image changes on the managed system.

Icon	🖥
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	Windows only. All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	Cannot be used on SNMP devices.
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No

Scheduler support	No
Files associated with this task	None
Events associated with this task	None

Remote Session

Use the Remote Session task to run command-line programs on a remote managed system. Remote Session creates less network traffic and uses fewer system resources than the Remote Control task and therefore is useful in low-bandwidth situations.

Note: You can have multiple remote sessions active at the same time, but you can have only one remote session through a management server to a single managed system.

Icon	🖥
Supported IBM Director objects	SNMP devices and Level-0, Level-1, and Level-2 managed systems
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	• (SNMP devices) The SSH or Telnet protocol is required. • (Managed systems running Linux or UNIX) The SSH protocol is required. **Note:** If the SSH server on the managed system does not respond, the Remote Session task attempts to use the Telnet protocol to connect to the managed system. • (Managed systems running i5/OS) The Telnet protocol is required.
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	None

Resource Monitors

Use the Resource Monitors task to view statistics about critical and noncritical system resources, such as processor, disk, and memory usage. With resource monitors, you also can set thresholds to detect potential problems with managed systems or devices. When a threshold is met or exceeded, an event is generated. You can create event action plans to respond to resource-monitor events. You can apply resource monitors to individual managed systems and devices and to groups. Also, you can record and view historical statistics and view current statistics.

Icon	📊
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	Director > Director Agent > Resource Monitors events which are customized events that you create. For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Scheduler

Use Scheduler to run a single noninteractive task or set of noninteractive tasks at a later time. You can specify an exact date and time you want the task to be started, or you can schedule a task to repeat automatically at a specified interval. Scheduled tasks are referred to as *jobs*.

Note: Only noninteractive tasks, which are tasks that do not require any user input or interaction, can be scheduled. The following tasks are noninteractive:

- BladeCenter Configuration Manager (the noninteractive subtask of BladeCenter Management)
- Capacity Manager
- Discovery
- Inventory

- Power Management
- Process Tasks, Process Monitors, and Remove Process Monitors (the noninteractive subtasks of Process Management)
- Software Distribution
- System Availability
- System Identification

IBM Director does not allow you to save changes to an existing job; you must always save changes to an existing job as a new job.

Icon	
Supported IBM Director objects	Any managed object and Level-0, Level-1, and Level-2 managed systems that are supported by noninteractive tasks.
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	Not applicable
Scheduler support	No
Files associated with this task	None
Events associated with this task	Director > Scheduler For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

Server Configuration Manager

Use the Server Configuration Manager task to create or update server configuration profiles. Configuration includes the service processors in eServer xSeries servers.

Note: The function provided by this task was previously provided as a part of the Management Processor Assistant task.

Icon	
Supported IBM Director objects	Physical Platform managed objects and Level-1 and Level-2 managed systems.

Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.p
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	Designed specifically for use on xSeries and Netfinity servers.
Required software	None
Required protocols	None
Required device drivers	For full Level-1 and Level-2 managed system support, you must install the applicable service processor device driver for the operating system that is running on the managed system.
Mass Configuration support	No, but the Server Configuration Manager task provides a feature where you can apply profiles against one or more managed objects at a time.
Scheduler support	Yes, you can schedule to apply profiles against one or more managed objects.
Files associated with this task	An XML file for importing and exporting server configuration profile information.
Events associated with this task	MPA For detailed events information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.

SNMP Browser

Use the SNMP Browser task to view and configure the attributes of SNMP devices, such as hubs, routers, or other SNMP-compliant management devices. This task is useful for SNMP-based management, troubleshooting, or monitoring the performance of SNMP devices.

Icon	
Supported IBM Director objects	SNMP devices only
Supported operating systems	All operating systems supported by IBM Director. For detailed operating-system support information, see the IBM Director information center at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	SNMP
Required device drivers	None
Mass Configuration support	No
Scheduler support	No

Files associated with this task	None
Events associated with this task	None

Software Distribution (Standard Edition)

Use the Software Distribution (Standard Edition) task to import applications and data, build a software package, and distribute the package to IBM Director managed systems. With IBM Director Software Distribution (Standard Edition), you can import only software that is distributed by IBM (such as UpdateXpress packages) and build a software package using only the IBM Update Assistant wizard.

For additional features and support, see the Software Distribution (Premium Edition) extension.

Icon	
Supported IBM Director objects	Level-0, Level-1, and Level-2 managed systems as distribution endpoints. **Note:** Only packages in the Solution Install format can be distributed to Level-0 and Level-1 managed systems.
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	Default protocol is TCP. If you disable TCP-session support on a managed system, Software Distribution (Standard Edition) uses UDP.
Required device drivers	None
Mass Configuration support	No
Scheduler support	Yes
Files associated with this task	The following files are associated with this task: • (For IBM Update Assistant) The software-update file and an XML file that describes the software-update file • (For Solution Install) JAR or ZIP files or Solution Install data using the Solution Install directory structure. This does not include archive format. Solution Install packages contain two XML files that describe the package.
Events associated with this task	None

System Accounts

Use the System Accounts task to view and change user and group security profiles on managed systems.

Icon	
Supported IBM Director objects	Level-2 managed systems
Supported operating systems	For detailed operating-system support information, see the IBM Director information center on the Web at publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo/fqm0_main.html.
Availability	Part of the standard IBM Director installation.
Required hardware or hardware limitations	None
Required software	None
Required protocols	None
Required device drivers	None
Mass Configuration support	No
Scheduler support	No
Files associated with this task	None
Events associated with this task	None

Appendix J. Icons

This topic provides information about IBM Director group and task icons.

Group icons

This topic provides information about the icons for the IBM Director groups.

Icon	Name	Description
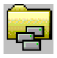	All Groups	This group contains all IBM Director groups. This group is always displayed.
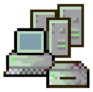	All Systems and Devices	This group contains all managed objects in the IBM Director environment. This group is always displayed.
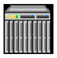	BladeCenter Chassis	This group contains all the BladeCenter chassis. This group is created after IBM Director discovers a BladeCenter chassis.
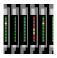	Chassis and Chassis Members	This group contains all chassis and the components that they contain. This group is always displayed.
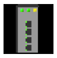	Chassis Ethernet Switch Members	This group contains Ethernet switch modules that are installed in a BladeCenter unit. This group is created after IBM Director discovers an Ethernet switch module in a BladeCenter chassis.
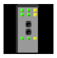	Chassis Fibre Channel Switch Members	This group contains Fibre Channel switch modules that are installed in a BladeCenter unit. This group is created after IBM Director discovers a Fibre Channel switch module in a BladeCenter chassis.
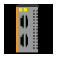	Chassis Infiniband Switch Members	This group contains Infiniband switch modules that are installed in a BladeCenter unit. This group is created after IBM Director discovers an Infiniband switch module in a BladeCenter chassis.
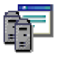	Clusters	This group contains cluster managed objects. This group is created when IBM Director discovers native managed objects that are part of a cluster and then creates a cluster managed object.
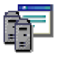	Clusters and Cluster Members	This group contains native managed objects (NMOs) that are part of a cluster. This group is always displayed. When IBM Director discovers native managed objects that are part of a cluster, it adds them to this group and creates a cluster managed object.
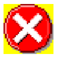	Hardware Status Critical	This group contains systems that have generated a critical-level event. This group is always displayed.
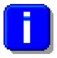	Hardware Status Information	This group contains systems that have generated an information-level event. This group is always displayed.
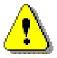	Hardware Status Warning	This group contains systems that have generated a warning-level event. This group is always displayed.
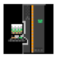	HMCs	This group contains Hardware Management Consoles (HMCs). This group is created when an HMC is discovered.

Icon	Name	Description
	HMCs and HMC Members	This group contains HMCs. It also contains the associated iSeries servers, pSeries servers, and logical partitions (LPARs). This group is created when an HMC is discovered.
	IBM Director Systems	This group contains all systems on which IBM Director Agent is installed. These are Level-2 managed systems. This group is displayed after IBM Director discovers a native managed object.
	Physical Platforms	This group contains the following types of managed objects: • Blade servers that are located in chassis that IBM Director discovers through Service Location Protocol (SLP) • Remote Supervisor Adapters and Remote Supervisor Adapters II that IBM Director discovers through SLP • Systems that Remote Deployment Manager (RDM) discovers as a result of a scan operation • Service processors that are installed in Level-2 managed systems • Service processors that are located in servers on an ASM interconnect network
	Platforms	This group contains all the managed objects that are represented by physical platform managed objects (PPMOs); it also contains LPARs.
	Platform and Platform Members	This group contains all members of the Physical Platforms and Platform groups. In addition, it also contains all Level-2 managed systems associated with those objects.
	Racks	This groups contains racks. This group is created after a user creates a rack-managed object.
	Racks with Members	This group contains racks and the managed objects that they contain. This group is always displayed.
	RMON Devices	This group contains SNMP devices with an RMON mib.
	Scalable Systems and Members	This group contains all members of Scalable Systems and Nodes groups.
	SMI-S Storage Devices	This group contains SMI-S Storage Devices. This group is created after IBM Director discovers SMI-S Storage Devices.
	SNMP Devices	This group contains SMP devices. This group is created after IBM Director discovers a device with an SNMP agent installed or embedded.
	Storage Devices	This group contains Storage Devices. This group is created after IBM Director discovers SMI-S Storage Devices.
	Systems with Asset ID	This group contains native managed objects that support Asset ID.

Icon	Name	Description
	Systems with ASF	This group contains native managed objects that support Alert Standard Format.
	Systems with ASF Secure Remote Management	This group contains native managed objects that support ASF 2.0.
	Systems with SNMP Agent.	This group contains native managed objects that have an SNMP agent installed or embedded.
	Systems with AIX	This group contains systems that are running AIX. This group is created after IBM Director discovers a native managed object running AIX. This group can be deleted.
	Systems with i5/OS	This group contains systems that are running i5/OS. This group is created after IBM Director discovers a native managed object running i5/OS. This group can be deleted.
	Systems with Linux	This group contains systems that are running Linux. This group is created after IBM Director discovers a native managed object running Linux. This group can be deleted.
	Systems with NetWare	This group contains systems that are running Novell NetWare. This group is created after IBM Director discovers a native managed object running NetWare. This group can be deleted.
	Systems with Windows 2000	This group contains systems that are running Windows 2000. This group is created after IBM Director discovers a native managed object running Windows 2000. This group can be deleted.
	Systems with Windows Server 2003	This group contains systems that are running Windows Server 2003. This group is created after IBM Director discovers a native managed object running Windows Server 2003. This group can be deleted.
	Systems with Windows XP	This group contains systems that are running Windows XP. This group is created after IBM Director discovers a native managed object running Windows XP. This group can be deleted.

Task icons

The task icons are displayed in the IBM Director Console Tasks pane.

Note: For each of these icons, clicking the button launches the specified task, and clicking the menu arrow opens a menu with additional options related to the specified task.

Icon	Task
	Active Console Viewer
	Asset ID

Icon	Task
	BladeCenter Management
	Capacity Manager
	CIM Browser
	Configure Alert Standard Format
	Configure SNMP Agent
	Discover Systems
	Electronic Service Agent
	Event Action Plans
	Event Log
	External Application Launch
	File Transfer
	Hardware Management Console
	Hardware Status
	Inventory
	Message Browser
	Microsoft Cluster Browser
	Network Configuration
	Process Management
	Rack Manager
	Remote Control
	Remote Deployment Manager
	Remote Session

Icon	Task
	Resource Monitors
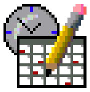	Scheduler
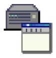	Server Configuration Manager
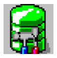	ServeRAID Manager
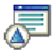	SNMP Browser
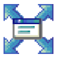	Software Distribution
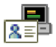	System Accounts
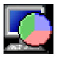	System Availability
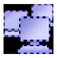	Virtual Machine Manager
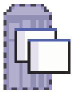	z/VM Center

Appendix K. Managed-object attributes

This topic lists the managed-object attributes for IBM Director managed objects.

Managed-object attributes may be displayed in the details view in IBM Director Console, and are used by some IBM Director command-line interface (CLI) commands.

The IBM Director CLI provides an interface to view and set managed object attributes using the **lsmo** and **chmo** commands. You must use the attribute-key and attribute-value-key keywords to specify attributes and their values. The attribute-name and attribute-value strings are translated display values and cannot be used as inputs to the **lsmo** or **chmo** commands. Values for attributes which are not writable cannot be changed using the **chmo** command.

Attribute key	Name	Writable	Values	Managed-object types
ClusterMO.members	Cluster Members	No	Integer: Managed-object identifier	• Clusters • Windows Clusters
CMO.chassisType	Chassis Type	No	Integer: One of the following values: `12` = Chassis `268` = Enterprise chassis `526` = Telco chassis	• BladeCenter Chassis • Chassis
CMO.subsystemList	Chassis Members	No	Integer: Managed-object identifier	• BladeCenter Chassis • Chassis
Complex.expectedChassis	Expected Scalable Nodes	No	Integer: The expected number of scalable nodes in this scalable system. This is the scalable system chassis count from the partition descriptor in nonvolatile random-access memory (NVRAM) on the service processor of a supported server.	• Scalable Systems
Complex.ID	Scalable System ID	No	String: Universal unique identifier (UUID) 16-byte hex value	• Scalable Systems
Complex.nodeOIDs	Scalable Nodes	No	String: Name(s) of the scalable nodes included in the scalable system	• Scalable Systems
LogicalPlatform.members	Host Membership	No	Integer: Managed-object identifier	• Logical Platforms • Scalable Partitions

Attribute key	Name	Writable	Values	Managed-object types
MH_ATTR_CONNECT_STATUS	Connection status	No	Byte: One of the following values: 0 = not attempted 1 = succeeded 2 = already connected 3 = failed 4 = failed no default stored 5 = failed missing login param 6 = failed mgmt hw not supported 7 = failed mgmt hw not stored 8 = failed addclient 9 = failed unknown host 10 = failed invalid gateway 11 = failed invalid target 12 = failed invalid gateway method 13 = failed no gateway comm 14 = failed no gateway factory 15 = failed name not found 16 = failed sp exception 17 = failed unknown method 18 = failed bad userid 19 = failed bad password 20 = failed disabled 21 = failed bad destination 22 = failed missing details 23 = failed unresolved default 24 = failed missing gateway moid 25 = failed invalid gateway family 26 = failed invalid gateway hardware 27 = failed no interconnect factory 28 = failed no server to agent sn 29 = failed no hw status 30 = failed discovery 31 = failed no factory 32 = failed no more channels 33 = failed no channel for moid 34 = failed no sn for channel 35 = failed invalid token 36 = failed unsupported sp hw 37 = failed no index for moid 38 = failed unknown hw family 39 = failed mo locked 40 = failed client service error 41 = failed ro gateway 50 = failed no details 60 = failed method deprecated 61 = failed exclusive lock 62 = success ip 63 = success rs485 64 = success ipc	• Physical Platforms
MH_ATTR_IP_ADDRESS	Management Module IP Address	No	String: Dot-delimited IP-address	• BladeCenter Chassis
MH_ATTR_PROMPT_FOR_ACCESS	Prompt for Access	Yes	String: One of the following values: false true	• • Physical Platforms
MH_ATTR_SP_TYPE	Service Processor Type	Physical platforms: Yes; Others: No	String: Service Processor type	• • Physical Platforms
MH_ATTR_TEXT_ID	Management Processor Text ID	Physical platforms: Yes; Others: No	String: Text ID	• • Physical Platforms

Attribute key	Name	Writable	Values	Managed-object types

MO.accessdenied	Access Denied	No	String: One of the following values:
			`false`
			`true`
				• BladeCenter Chassis
				• Chassis
				• Clusters
				• Level 0: Agentless Systems
				• Level 1: Core Services
				• Level 2: IBM Director Agents
				• Logical Platforms
				• Physical Platforms
				• Platforms
				• Racks
				• Remote I/O Enclosures
				• RMON Devices
				• Scalable Partitions
				• Scalable Systems
				• SMI-S Storage Devices
				• SNMP Devices
				• SNMP Printers
				• Storage Devices
				• Windows Clusters

MO.activeSlotNumber	Management Module Active Slot Number	No	Integer: One of the following values:
			`1`
			`2`
				• BladeCenter Chassis

MO.AgentDate	Agent Build Date	No	String of the form *dow mon dd hh:mm:ss zzz yyyy*: *dow* - Three-letter abbreviation for the day of the week (Sun, Mon, Tue, Wed, Thu, Fri, Sat) *mon* - Three-letter abbreviation for the month (Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec) *dd* - Two-digit representation of the day of the month (01 through 31) *hh* - Two-digit representation of the hour of the day (00 through 23) *mm* - Two-digit representation of the minute within the hour (00 through 59) *ss* - Two-digit representation of the second within the minute (00 through 59) *zzz* - Standard three-letter abbreviation for the time zone. The time zone abbreviation may indicate daylight savings time. If no time zone information is available, then *zzz* is empty—that is, it consists of no characters. *yyyy* - Four-digit representation of the year.	• Level 1: Core Services • Level 2: IBM Director Agents
MO.agenttimezoneoff	Agent Time Zone Offset (GMT + minutes)	Yes	Integer: GMT(+/-) minutes	• Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents

MO.AgentType	Agent Type	No	String: One of the following values:
			`Agentless`
			`Director_Server`
				• Level 0: Agentless Systems
				• Level 1: Core Services
				• Level 2: IBM Director Agents

MO.ArchType	Agent Architecture	No	String: One of the following values:
			`IA32`
			`IA64`
			`PowerPC`
			`PowerPC64`
			`RISC`
			`S390`
			`x86_64`
				• Level 1: Core Services
				• Level 2: IBM Director Agents

Attribute key	Name	Writable	Values	Managed-object types
MO.ComputerName	Computer Name	No	String: Name of the managed object	• BladeCenter Chassis • Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents
MO.ElementName	Nickname	No	String: User defined nickname for the device (for example, ″my storage device″)	• SMI-S Storage Devices
MO.EnclosureFRU	Chassis FRU Number	No	String: FRU	• BladeCenter Chassis
MO.EnclosureMachineTypeModel	Chassis Machine Type and Model	No	String: Machine type and model	• BladeCenter Chassis
MO.EnclosureSerialNumber	Chassis Serial Number	No	String: Serial number	• BladeCenter Chassis
MO.EnclosureUUID	Chassis UUID	No	String: Universal unique identifier (UUID) 16-byte hex value	• BladeCenter Chassis

MO.encryptionenabled	Encryption Enabled	No	String: One of the following values:
			`false`
			`true`
				• BladeCenter Chassis
				• Chassis
				• Clusters
				• Level 0: Agentless Systems
				• Level 1: Core Services
				• Level 2: IBM Director Agents
				• Logical Platforms
				• Physical Platforms
				• Platforms
				• Racks
				• Remote I/O Enclosures
				• RMON Devices
				• Scalable Partitions
				• Scalable Systems
				• SMI-S Storage Devices
				• SNMP Devices
				• SNMP Printers
				• Storage Devices
				• Windows Clusters

MO.Endpoint	Endpoint Unique Name	No	String: A unique identifier for the hardware device (for example, ″acme148 IBM 1S17001RS23A00AE DS400″)	• SMI-S Storage Devices

MO.hasLicense	Granted License	Yes	String: One of the following values:
			`false`
			`true`
				• Level 2: IBM Director Agents

MO.InteropSchemaNamespace	Interop Namespace	No	String: The interoperability CIM namespace used to obtain metadata base the SMI-S provider (for example, ″/interop″)	• SMI-S Storage Devices
MO.IPaddrs	TCP/IP Addresses	Physical platforms: Yes; Others: No	String: Comma-delimited list of dot-delimited IP-addresses, for example, `'192.168.2.3', '9.49.131.29'`	• BladeCenter Chassis • Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents • Logical Platforms • Physical Platforms • Platforms • RMON Devices • Scalable Partitions • SNMP Devices • SNMP Printers
MO.IPhosts	TCP/IP Hosts	Physical platforms: Yes; Others: No	String: Comma-delimited list of host names, for example, `'gmolson.pok.ibm.com'`, `'gmolson.rchland.ibm.com'`, `'sig-9-49-131-29.mts.ibm.com'`	• BladeCenter Chassis • Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents • Logical Platforms • Physical Platforms • Platforms • RMON Devices • Scalable Partitions • SNMP Devices • SNMP Printers

Attribute key	Name	Writable	Values	Managed-object types
MO.IPPort	SMI-S Provider IP Port Numbers	No	String: The port number through which the SMI-S CIM Server accepts connections (for example, 5988)	• SMI-S Storage Devices
MO.IPXaddrs	IPX Addresses		String: Comma-delimited list of IPX addresses in the following format: `network.node` *network* - Eight-character hexadecimal representation of the network number *node* - Twelve-character hexadecimal representation of the node number	• Level 2: IBM Director Agents
MO.iSCSISupport	Storage Device Type	No	String: The protocol type through which this storage device can be accessed through the SAN (for example, ″iSCSI, Fibre Channel″)	• SMI-S Storage Devices
MO.LparID	LPAR ID		String: Logical partition ID	• Level 0: Agentless Systems • Level 1: Core Services • Logical Platforms • Physical Platforms • Platforms
MO.MACAddress	MAC Address		String: 6-byte hex value	• Level 2: IBM Director Agents
MO.MACAddrList	MAC Addresses	Physical platforms: Yes; Others: No	String: Comma-delimited list of MAC-address 6-byte or 8-byte hex values. Platforms and partitions use the 8-byte form; SNMP and RMON devices use the 6-byte form.	• Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents • Logical Platforms • Physical Platforms • Platforms • RMON Devices • SNMP Devices • SNMP Printers • Scalable Partitions
MO.MachineTypeModel	Machine Type and Model	Physical platforms: Yes; Others: No	String	• Level 0: Agentless Systems • Level 1: Core Services • Logical Platforms • Physical Platforms • Platforms • Scalable Partitions
MO.ManufacturerId	Manufac-turer	No	String: The manufacturer of the hardware device (for example, ″IBM″)	• SMI-S Storage Devices

MO.MOFID	System Factory ID	No	String: One of the following values:
			`BladeCenter Chassis`
			`Chassis`
			`Clusters`
			`Level 0: Agentless Systems`
			`Level 1: Core Services`
			`Level 2: IBM Director Agents`
			`Logical Platforms`
			`Physical Platforms`
			`Platforms`
			`Racks`
			`Remote I/O Enclosures`
			`RMON Devices`
			`Scalable Partitions`
			`Scalable Systems`
			`SMI-S Storage Devices`
			`SNMP Devices`
			`SNMP Printers`
			`Storage Devices`
			`Windows Clusters`
				• BladeCenter Chassis
				• Chassis
				• Clusters
				• Level 0: Agentless Systems
				• Level 1: Core Services
				• Level 2: IBM Director Agents
				• Logical Platforms
				• Physical Platforms
				• Platforms
				• Racks
				• Remote I/O Enclosures
				• RMON Devices
				• Scalable Partitions
				• Scalable Systems
				• SMI-S Storage Devices
				• SNMP Devices
				• SNMP Printers
				• Storage Devices
				• Windows Clusters

Attribute key	Name	Writable	Values	Managed-object types
MO.name	System Name	Yes	String: Name	• BladeCenter Chassis • Chassis • Clusters • Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents • Logical Platforms • Physical Platforms • Platforms • Racks • Remote I/O Enclosures • RMON Devices • Scalable Partitions • Scalable Systems • SMI-S Storage Devices • SNMP Devices • SNMP Printers • Storage Devices • Windows Clusters
MO.OpSys	Operating System	No	String: One of the following values: AIX DOS FREEBSD HPUX IRIX LINUX NETWARE OPENBSD OS2 OS400 SCOUNIX SOLARIS SUNOS UNIXWARE UNIX_SYS_V VINES WINDOWS WINDOWS_NT [1] WINDOWS_9X	• Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents • RMON Devices • SNMP Devices
MO.OpSysArchType	OS Architecture	No	String: One of the following values: IA32 IA64 PowerPC PowerPC64 RISC S390 x86_64	• Level 0: Agentless Systems • Level 1: Core Services
MO.OpSysMajVer	OS Major Version	No	Integer	• Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents • RMON Devices • SNMP Devices
MO.OpSysMinVer	OS Minor Version	No	Integer	• Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents • RMON Devices • SNMP Devices

Attribute key	Name	Writable	Values	Managed-object types
MO.ping	System Presence Check Setting (minutes)	Yes	Non-negative integer: Number of minutes between system presence checks	• BladeCenter Chassis • Chassis • Clusters • Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents • Logical Platforms • Physical Platforms • Platforms • Racks • Remote I/O Enclosures • RMON Devices • Scalable Partitions • Scalable Systems • SMI-S Storage Devices • SNMP Devices • SNMP Printers • Storage Devices • Windows Clusters
MO.Protocols	Protocols Supported		String: Comma-delimited list of integers representing supported protocols	• Level 0: Agentless Systems • Level 1: Core Services
MO.RegisteredProfilesSupported	Registered Profiles Supported	No	String: The complete list of SMI-S profiles available (for example, ″SNIA: Array 1.0.2, SNIA:Server 1.0.2″)	• SMI-S Storage Devices

MO.secunsecsupport	Secure / Unsecure Supported	No	String: One of the following values:
			`false`
			`true`
			Note: `true` is the only valid value for IBM Director systems. For other managed objects, `false` is the only valid value.
				• BladeCenter Chassis
				• Chassis
				• Clusters
				• Level 0: Agentless Systems
				• Level 1: Core Services
				• Level 2: IBM Director Agents
				• Logical Platforms
				• Physical Platforms
				• Platforms
				• Racks
				• Remote I/O Enclosures
				• RMON Devices
				• Scalable Partitions
				• Scalable Systems
				• SMI-S Storage Devices
				• SNMP Devices
				• SNMP Printers
				• Storage Devices
				• Windows Clusters

MO.SerialNumber	Machine Serial Number	Physical platforms: Yes; Others: No	String: Serial number	• BladeCenter Chassis • Level 0: Agentless Systems • Level 1: Core Services • Logical Platforms • Scalable Partitions • Physical Platforms • Platforms
MO.ServiceID	SMI-S Provider Service ID	No	String: A unique string distinguishing the SMI-S CIM Server from other CIM Servers (for example, ″LSISSI:09.03.c2.cc″)	• SMI-S Storage Devices
MO.SLIMPort	Management Module SLIM Port Number	No	Integer: port number	• BladeCenter Chassis
MO.SNMPPort	Management Module SNMP Port Number	No	Integer: port number	• BladeCenter Chassis
MO.SSHFingerprint	SSH Host Key Fingerprint		String	• Level 0: Agentless Systems • Level 1: Core Services
MO.SSHPort	SSH Port			• Level 0: Agentless Systems • Level 1: Core Services
MO.SSHVersion	SSH Version			• Level 0: Agentless Systems • Level 1: Core Services

Attribute key	Name	Writable	Values	Managed-object types

MO.state	System State	No	String: One of the following values:
			`Indeterminate`
			`NotValid`
			`Offline`
			`OfflineError`
			`Online`
			`OnlineError`
			`Unknown`
			`Unlicensed`
				• BladeCenter Chassis
				• Chassis
				• Clusters
				• Level 0: Agentless Systems
				• Level 1: Core Services
				• Level 2: IBM Director Agents
				• Logical Platforms
				• Physical Platforms
				• Platforms
				• Racks
				• Remote I/O Enclosures
				• RMON Devices
				• Scalable Partitions
				• Scalable Systems
				• SMI-S Storage Devices
				• SNMP Devices
				• SNMP Printers
				• Storage Devices
				• Windows Clusters

MO.StorageDeviceType	Registered Profile	No	String: The main SMI-S profile used to represent the device (for example, ″SNIA:Array 1.0.2″)	• SMI-S Storage Devices
MO.Subscription	Subscrip- tions	No	String: A list of the CIM indication subscriptions that were created for this device to monitor health (for example, ″Overall Operational Status, Alerts″)	• SMI-S Storage Devices

MO.unsecureclient	Agent Unsecured	No	String: One of the following values:
			`false`
			`true`
				• Level 2: IBM Director Agents

MO.URL	WEB URL	No	String: Web address of the management module	• BladeCenter Chassis • Level 1: Core Services
MO.URLs	SLP Service Name	No	String: Web address used to connect to the SMI-S CIM Server supporting the device (for example, ″http://9.3.194.204:5988″)	• SMI-S Storage Devices
MO.UUID	System UUID	Physical platforms: Yes; Others: No	String: Universal unique identifier (UUID) 16-byte hex value	• Remote I/O Enclosures • Logical Platforms • Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents • Scalable Partitions • Scalable Systems • Physical Platforms • Platforms
MO.WBEMPort	WBEM Port		Integer	• Level 1: Core Services
MSCSClusterMO.nameresid	Cluster Resource ID	No	String: Cluster-resource ID	• Windows Clusters
MSCSClusterMO.version	Cluster Services Version	No	String: Version, in the form *major.minor.build*	• Windows Clusters

NativeMO.agentProdVersion	Agent Version	No	String: Values include but are not limited to the following:
			`3.10`
			`4.00`
			`4.20`
			`5.10`
				• Level 1: Core Services
				• Level 2: IBM Director Agents

NativeMO.allpaths	Available IPC Paths	No	String: Comma-delimited list of IPC locations in the form *protocol::address*, for example `'TCPIP::192.168.0.1', 'TCPIP::192.168.0.3'`	• Level 2: IBM Director Agents • Windows Clusters
NativeMO.path	Preferred IPC Path	No	String: IPC Location in the form *protocol::address*, for example, `TCPIP::192.168.0.1`	• Level 2: IBM Director Agents • Windows Clusters
NativeMO.uniqueid	Unique System ID	No	String: Universal system identifier 8 byte hex value	• Level 0: Agentless Systems • Level 1: Core Services • Level 2: IBM Director Agents
Partition.boot_flags	Boot Flags	No	Integer: Always equals `0`.	• Scalable Partitions
Partition.boot_path	Boot Path	No	String: Always an empty string (″″).	• Scalable Partitions

Attribute key	Name	Writable	Values	Managed-object types
Partition.complex_oid	Scalable System	No	String: Scalable-system name	• Scalable Partitions
Partition.expectedChassis	Expected Scalable Nodes	No	Integer: The expected number of scalable nodes in this scalable partition. This is the scalable system chassis count from the partition descriptor in nonvolatile random-access memory (NVRAM) on the service processor of a supported server.	• Scalable Partitions
Partition.id	Partition ID	No	String: Universal unique identifier (UUID) 16-byte hex value. This is the partition UUID from the partition descriptor in NVRAM on the service processor of a supported server.	• Scalable Partitions
Partition.primary_node_oid	Primary Scalable Node	No	String: Primary scalable-node name	• Scalable Partitions
Partition.state	State	No	Integer: One of the following values: 0 = Null or unknown 1 = Powered on 2 = Powering off 3 = Powered off 4 = Powering on 5 = Resetting	• Scalable Partitions
PhysicalPlatform.members	Host Membership	No	Integer: Managed-object identifier	• Physical Platforms
Platform.members	Host Membership	No	Integer: Managed-object identifier	• Logical Platforms • Platforms • Scalable Partitions
RackMO.location	Rack Description		String: Description of rack	• Racks
RackMO.slots	Rack Slots		Integer: Number of rack slots	• Racks
RackMO.type	Rack Type		String: Rack type descriptor	• Racks
RIOEnclosure.attachedPPMOs	Physical Platforms	No	String: The names of the physical platforms attached to this remote I/O enclosure.	• Remote I/O Enclosures
RIOEnclosure.slots	Slots	No	Integer: The number of slots in the remote I/O enclosure. One of the following values: 6 = Expansion Kit A 12 = Expansion Kits A and B	• Remote I/O Enclosures
SMISMO.cimNamespace	Namespace	No	String: The CIM namespace used for communication with the registered profile of the device (for example, ″root/lsissi″)	• SMI-S Storage Devices
SNMPMO.engineID	Engine ID		String	• RMON Devices • SNMP Devices • SNMP Printers
SNMPMO.sysContact	System Contact (MIB2)	No	String	• RMON Devices • SNMP Devices • SNMP Printers
SNMPMO.sysDescr	System Description (MIB2)	No	String	• RMON Devices • SNMP Devices • SNMP Printers
SNMPMO.sysLocation	System Location (MIB2)	No	String	• RMON Devices • SNMP Devices • SNMP Printers
SNMPMO.sysName	System Name (MIB2)	No	String:Name	• RMON Devices • SNMP Devices • SNMP Printers
SNMPMO.sysObjectID	System Object ID (MIB2)	No	String	• RMON Devices • SNMP Devices • SNMP Printers
SNMPMO.sysUpTime	System Uptime (MIB2)	No	String	• RMON Devices • SNMP Devices • SNMP Printers

[1] The attributes MO.OpSys, MO.OpSysMajVer, and MO.OpSysMinVer combine to indicate the specific Microsoft Windows version. Refer to the following table.

Microsoft Windows version	**MO.OpSys**	**MO.OpSysMajVer**	**MO.OpSysMinVer**
Microsoft Windows 2000	WINDOWS_NT	5	0
Microsoft Windows XP	WINDOWS_NT	5	1
Microsoft Windows Server 2003	WINDOWS_NT	5	2
Microsoft Windows 95	WINDOWS_9X	4	0
Microsoft Windows 98	WINDOWS_9X	4	10
Microsoft Windows Me	WINDOWS_9X	4	90

Appendix L. Resource-monitor attributes

You can use the Resource Monitors task to monitor critical system resources on managed systems. The resources that you can monitor are different depending on the operating system that is installed on the managed system. Use these resource-monitor attributes tables to identify the resource-monitor attributes that you want to monitor if you are planning your IBM Director installation or configuration or adjusting your resource-monitoring strategy.

Resource monitor data-collection rates vary depending on the managed system or device. In general, using the default settings, data collections occur every 5 to 10 seconds, and the display refreshes every 10 to 20 seconds.

To view the resource-monitor attributes that are available for a managed system or device, see viewing all resource-monitor thresholds.

AIX resource-monitor attributes

These resource-monitor attributes are for the AIX operating system.

Resource monitor	Attributes
CPU	• CPU utilization • Process count
Disk	**Note:** The disk drive monitor attributes are repeated for each local nonremovable logical drive that is found. • Blocks available • Blocks used • Inodes available • Inodes used • Percentage blocks available • Percentage block used • Percentage Inodes available • Percentage Inodes used • Percentage space available • Percentage space used • Space available (MB) • Space used (MB)

Resource monitor	Attributes
File	**Notes:** 1. File-monitor attributes can be files or directories. 2. For compatible file-system types, the “Directory exists” or “File exists” attribute (depending on which is applicable) is always valid data. • **Directory** – Directory exists – Last modified – Directory attributes – Directory owner – Directory size (bytes) – Object type – • **File** – Checksum – File exists – Last modified – File attributes – File owner – File size (bytes) – Object type
Memory	• Available (bytes) • Used (bytes) • Total memory
Process	**Note:** The number of applications or executable files that a process monitor checks can vary. The IBM Director user configures the processes that are monitored using the Process Monitor task in IBM Director Console. Each of the process-monitor attributes is displayed for each executable file that is monitored. • Current active processes • Maximum running at once • Maximum running yesterday • New executions counted • Times failed to start • Time started • Time stopped • Total execution time • Yesterday’s execution time • Yesterday’s new executions
UNIX system	• CPU monitors • Disk monitors • Memory monitors

i5/OS resource-monitor attributes

These resource-monitor attributes are for the i5/OS operating system.

Resource monitor	Attributes
File	**Notes:** 1. File-monitor attributes can be files or directories. 2. For compatible file-system types, the “Directory exists” or “File exists” attribute (depending on which is applicable) is always valid data. 3. If there are additional directories, additional subelements are displayed. 4. A directory can contain hundreds of subelements. If it does, a directory might take 5 seconds or longer to open. 5. QSYS.LIB can contain thousands of subelements. If a timeout occurs, reopening the directory after a timeout increases the timeout value, and may increase the timeout value sufficiently for the operation to complete. • **Directory** – Directory exists – Last modified – Directory attributes – Directory owner – Directory size (bytes) – Object type – • **File** – Checksum – File exists – Last modified – File attributes – File owner – File size (bytes) – Object type
File system	**Note:** The file system monitor attributes for specific directories are provided for typical i5/OS directories. If one of these directories does not exist, the attribute is not displayed. • / • /bin • /dev • /etc • /home • /lib • /tmp • /usr • /var
List of directory contents	• Directory attributes • Directory exists • Directory owner • Directory size (bytes) • Last modified • Object type

Resource monitor	Attributes
Process	**Note:** The number of applications or executable files that a process monitor checks can vary. The IBM Director user configures the processes that are monitored using the Process Monitor task in IBM Director Console. Each of the process-monitor attributes is displayed for each executable file that is monitored. • Current active processes • Maximum running at once • Maximum running yesterday • New executions counted • Times failed to start • Time started • Time stopped • Total execution time • Yesterday’s execution time • Yesterday’s new executions
I/O processors	• Auxiliary IOP Use % • IOP All Comm. Use % • IOP Disk Use % • IOP LAN Use % • IOP Memory Free (KB) • IOP Operational Status • IOP Optical Use % • IOP SDLC Use % • IOP System Function Use % • IOP Tape Use % • IOP Twinaxial Use % • IOP X.25 Use % • Primary IOP Use %
Job queues	• Job Queue Status • Jobs in Queue
Job statistics	• Batch Jobs Ended, Output Waiting • Batch Jobs Ending • Batch Jobs Held on Job Queue • Batch Jobs Held while Running • Batch Jobs on Held Job Queue • Batch Jobs on Unassigned Job Queue • Batch Jobs Running • Batch Jobs Waiting for Messages • Batch Jobs Waiting to Run • Jobs on System

Resource monitor	Attributes
NetServer™ statistics	• Average Response Time (Milliseconds) • File Opens/Minute • Kbytes Received per Minute • Kbytes Sent per Minute • Password Violations • Print Jobs Queued/Minute • Session Starts/Minute
Physical disks	• Disk Arm Utilization % • Disk Average Queue Length • Disk Mirroring Status • Disk Operational Status • Disk Processor Utilization % • Disk Read Commands/Minute • Disk Read Kbytes/Minute • Disk Space Free (MB) • Disk Space Used % • Disk Write Commands/Minute • Disk Write Kbytes/Minute
Storage pools	• Active to Ineligible (Transitions/Minute) • Active to Wait (Transitions/Minute) • Database Faults per Second • Database Pages per Second • Non-database Faults per Second • Non-database Pages per Second • Wait to Ineligible (Transitions/Minute)
Subsystems	• Subsystem % of Job Limit • Subsystem Active Jobs • Subsystem Status
System statistics	• CPU Utilization % • Current Temp Storage Used (MB) • Max Temp Storage Used (MB) • Permanent Addresses Used % • System ASP Used % • Temporary Addresses Used %
User statistics	• Users Disconnected • Users Signed Off, Output Waiting • Users Signed On • Users Suspended by Group Jobs • Users Suspended by System Request

Linux resource-monitor attributes

These resource-monitor attributes are for the Linux on xSeries and Linux on System z9 and zSeries operating systems.

Note: For resource-attribute information for Linux on POWER, see "Linux on POWER resource-monitor attributes" on page 748.

Resource monitor	**Attributes**
CPU	• CPU utilization • Process count
Disk	**Notes:** 1. The disk drive monitor attributes are repeated for each local nonremovable logical drive that is found. 2. The list of file-system attributes is displayed first; then, the disk monitor attributes are displayed under each file system. • Blocks available • Blocks used • Inodes available • Inodes used • Percentage blocks available • Percentage block used • Percentage Inodes available • Percentage Inodes used • Percentage space available • Percentage space used • Space available (MB) • Space used (MB) • Volume SYS: space remaining • Volume SYS: space used
File	**Notes:** 1. File-monitor attributes can be files or directories. 2. For compatible file-system types, the "Directory exists" or "File exists" attribute (depending on which is applicable) is always valid data. 3. If there are additional directories, additional subelements are displayed. 4. A directory can contain hundreds of subelements. If it does, a directory might take 5 seconds or longer to open. • **Directory** – Directory exists – Last modified – Directory attributes – Directory owner – Directory size (bytes) – Object type – • **File** – Checksum – File exists – Last modified – File attributes – File owner – File size (bytes) – Object type

Resource monitor	Attributes
File system	**Note:** The file system monitor attributes for specific directories are provided for typical Linux directories. If one of these directories does not exist, the attribute is not displayed. • / • /bin • /dev • /etc • /home • /lib • /lost+found • /sbin • /tmp • /usr • /var
List of directory contents	• Directory attributes • Directory exists • Directory owner • Directory size (bytes) • Last modified • Object type
Memory	• Available (bytes) • Used (bytes) • Total memory • Unused non-cached (MBytes)
Process	**Note:** The number of applications or executable files that a process monitor checks can vary. The IBM Director user configures the processes that are monitored using the Process Monitor task in IBM Director Console. Each of the process-monitor attributes is displayed for each executable file that is monitored. • Current active processes • Maximum running at once • Maximum running yesterday • New executions counted • Times failed to start • Time started • Time stopped • Total execution time • Yesterday’s execution time • Yesterday’s new executions
UNIX system	• CPU monitors • Disk monitors • Disk performance monitors • Memory monitors • Network monitors

Resource monitor	Attributes
CIM	**Note:** The attributes for CIM monitors can vary depending on the features and functions that you have configured on the managed system. • Namespaces • Classes • Instances • Properties

Linux on POWER resource-monitor attributes

These resource-monitor attributes are for the Red Hat Enterprise Linux AS for IBM POWER and SUSE LINUX Enterprise Server 9 for IBM POWER operating systems.

Note: For resource-attribute information for other versions of Linux, see "Linux resource-monitor attributes" on page 745.

Resource monitor	Attributes
CPU	• CPU utilization • Process count
Disk	**Notes:** 1. The disk drive monitor attributes are repeated for each local nonremovable logical drive that is found. 2. The list of file-system attributes is displayed first; then, the disk monitor attributes are displayed under each file system. • Blocks available • Blocks used • Inodes available • Inodes used • Percentage blocks available • Percentage block used • Percentage Inodes available • Percentage Inodes used • Percentage space available • Percentage space used • Space available (MB) • Space used (MB)

Resource monitor	Attributes
File	**Notes:** 1. File-monitor attributes can be files or directories. 2. For compatible file-system types, the “Directory exists” or “File exists” attribute (depending on which is applicable) is always valid data. 3. If there are additional directories, additional subelements are displayed. 4. A directory can contain hundreds of subelements. If it does, a directory might take 5 seconds or longer to open. • **Directory** – Directory exists – Last modified – Directory attributes – Directory owner – Directory size (bytes) – Object type –

Resource monitor	Attributes
File system	**Note:** The file system monitor attributes for specific directories are provided for typical Linux directories. If one of these directories does not exist, the attribute is not displayed. • / • /bin • /boot • /dev • /etc • /home • /lib • /lib64 • /media • /mnt • /opt • /proc • /root • /sbin • /srv • /sys • /tmp • /usr • /var The following attributes are available only on Red Hat Linux: • /autofsck • /O • /abb.srv • /afile • /AZ • /initrd • /lost+found • /mary • /mary_redir • /misc • /redir • /remote_pkg • /rpms • /selinux • /tftpboot • /ux
List of directory contents	• Directory attributes • Directory exists • Directory owner • Directory size (bytes) • Last modified • Object type

Resource monitor	Attributes
Memory	• Available (bytes) • Used (bytes) • Total memory • Unused non-cached (MBytes)
UNIX system	• CPU monitors • Disk monitors • Disk performance monitors • Memory monitors • Network monitors
CIM	**Note:** The attributes for CIM monitors can vary depending on the features and functions that you have configured on the managed system. • CIMV2 • ibmsd • pg_internal • pg_interop

NetWare resource-monitor attributes

These resource-monitor attributes are for the NetWare operating system.

Resource monitor	Attributes
CPU	• CPU utilization • CPU ’x’ utilization (on SMP devices • Process count • Thread count
Disk	**Note:** The disk drive monitor attributes are repeated for each local nonremovable logical drive that is found. • Volume SYS: space remaining • Volume SYS: space used
File	**Notes:** 1. File-monitor attributes can be files or directories. 2. For compatible file-system types, the “Directory exists” or “File exists” attribute (depending on which is applicable) is always valid data. • **Directory** – Directory exists – Last modified – • **File** – Checksum – File exists – Last modified – File size (bytes)
Memory	• Cache blocks in use • Percent of cache in use

Resource monitor	Attributes
Process	**Note:** The number of applications or executable files that a process monitor checks can vary. The IBM Director user configures the processes that are monitored using the Process Monitor task in IBM Director Console. Each of the process-monitor attributes is displayed for each executable file that is monitored. • Current active processes • Maximum running at once • Maximum running yesterday • New executions counted • Times failed to start • Time started • Time stopped • Total execution time • Yesterday's execution time • Yesterday's new executions

UNIX resource-monitor attributes

These resource-monitor attributes are for the UNIX operating system.

Resource monitor	Attributes
CPU	• CPU utilization • Process count
Disk	**Notes:** 1. The disk drive monitor attributes are repeated for each local nonremovable logical drive that is found. 2. The list of file-system attributes is displayed first; then, the disk monitor attributes are displayed under each file system. • Blocks available • Blocks used • Inodes available • Inodes used • Percentage blocks available • Percentage block used • Percentage Inodes available • Percentage Inodes used • Percentage space available • Percentage space used • Space available (MB) • Space used (MB) • Volume SYS: space remaining • Volume SYS: space used

Resource monitor	Attributes
File	**Notes:** 1. File-monitor attributes can be files or directories. 2. For compatible file-system types, the “Directory exists” or “File exists” attribute (depending on which is applicable) is always valid data. 3. If there are additional directories, additional subelements are displayed. 4. A directory can contain hundreds of subelements. If it does, a directory might take 5 seconds or longer to open. • **Directory** – Directory exists – Last modified – Directory attributes – Directory owner – Directory size (bytes) – Object type – • **File** – Checksum – File exists – Last modified – File attributes – File owner – File size (bytes) – Object type
File system	**Note:** The file system monitor attributes for specific directories are provided for typical UNIX directories. If one of these directories does not exist, the attribute is not displayed. • / • /bin • /dev • /etc • /home • /lib • /lost+found • /sbin • /tmp • /usr • /var
List of directory contents	• Directory attributes • Directory exists • Directory owner • Directory size (bytes) • Last modified • Object type
Memory	• Available (bytes) • Used (bytes)

Resource monitor	Attributes
Process	**Note:** The number of applications or executable files that a process monitor checks can vary. The IBM Director user configures the processes that are monitored using the Process Monitor task in IBM Director Console. Each of the process-monitor attributes is displayed for each executable file that is monitored. • Current active processes • Maximum running at once • Maximum running yesterday • New executions counted • Times failed to start • Time started • Time stopped • Total execution time • Yesterday’s execution time • Yesterday’s new executions
UNIX system	• CPU monitors • Disk monitors • Disk performance monitors • Memory monitors • Network monitors
CIM	**Note:** The attributes for CIM monitors can vary depending on the features and functions that you have configured on the managed system. • Namespaces • Classes • Instances • Properties

Windows resource-monitor attributes

These resource-monitor attributes are for the Windows operating system.

Note: The attributes for the following resource monitors can vary depending on the features and functions that you have configured on the managed system:
- CIM monitors
- Device, performance, and service monitors
- Registry monitors

Resource monitor	Attributes
CPU	• CPU utilization • CPU ’x’ utilization (on SMP devices) • Process count

Resource monitor	Attributes
Disk	**Note:** The disk drive monitor attributes are repeated for each local nonremovable logical drive that is found. • Disk 1 workload • Drive C: % space used • Drive C: Space remaining • Drive C: Space used
File	**Notes:** 1. File-monitor attributes can be files or directories. 2. For compatible file-system types, the “Directory exists” or “File exists” attribute (depending on which is applicable) is always valid data. • **Directory** – Directory exists – Last modified • **File** – Checksum – File exists – Last modified – File size (bytes)
Memory	• Locked memory • Memory usage
TCP/IP	• Interface *x* - Broadcast packets received • Interface *x* - Broadcast packets sent • Interface *x* - Bytes received • Interface *x* - Bytes sent • Interface *x* - Unicast packets received • Interface *x* - Unicast packets sent • IP packets received • IP packets received with errors • IP packets sent • TCP connections • UDP datagrams received • UDP datagrams sent

Resource monitor	Attributes
Process	**Note:** The number of applications or executable files that a process monitor checks can vary. The IBM Director user configures the processes that are monitored using the Process Monitor task in IBM Director Console. Each of the process-monitor attributes is displayed for each executable file that is monitored. • Current active processes • Maximum running at once • Maximum running yesterday • New executions counted • Times failed to start • Time started • Time stopped • Total execution time • Yesterday’s execution time • Yesterday’s new executions
CIM	• Namespaces • Classes • Instances • Properties

Appendix M. Security

This topic provides information about the security provided by cipher suites.

You can use Secure Sockets Layer (SSL) to protect data flowing between IBM Director Server and IBM Director Console.

IBM Director supports the following cipher suites:

- SSL_anon_WITH_AES_128_CBC_SHA
- SSL_DH_anon_EXPORT_WITH_DES40_CBC_SHA
- SSL_DH_anon_EXPORT_WITH_RC4_40_MD5
- SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
- SSL_DH_anon_WITH_DES_CBC_SHA
- SSL_DH_anon_WITH_RC4_128_MD5
- SSL_DHE_DSS_EXPORT_WITH_DES40_CBC_SHA
- SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA
- SSL_DHE_DSS_WITH_AES_128_CBC_SHA
- SSL_DHE_DSS_WITH_DES_CBC_SHA
- SSL_DHE_DSS_WITH_RC4_128_SHA
- SSL_DHE_RSA_EXPORT_WITH_DES40_CBC_SHA
- SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_DHE_RSA_WITH_AES_128_CBC_SHA
- SSL_DHE_RSA_WITH_DES_CBC_SH
- SSL_RSA_EXPORT_WITH_DES40_CBC_SHA
- SSL_RSA_EXPORT_WITH_RC2_CBC_40_MD5
- SSL_RSA_EXPORT_WITH_RC4_40_MD5
- SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_FIPS_WITH_DES_CBC_SHA
- SSL_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_WITH_AES_128_CBC_SHA
- SSL_RSA_WITH_DES_CBC_SHA
- SSL_RSA_WITH_NULL_MD5
- SSL_RSA_WITH_NULL_SHA
- SSL_RSA_WITH_RC4_128_MD5
- SSL_RSA_WITH_RC4_128_SHA

See the documentation for the operating system on which you are running IBM Director to determine which cipher suites are supported by the operating system.

Appendix N. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION ″AS IS″ WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
MW9A/050
5600 Cottle Road
San Jose, CA 95193
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX
AIX 5L
Alert on LAN
Asset ID

BladeCenter
DB2
DB2 Universal Database
DirMaint
Electronic Service Agent
Enterprise Storage Server
eServer
eServer logo
FlashCopy
HiperSockets
i5/OS
IBM
IBM logo
ibm.com
IntelliStation
iSeries
Netfinity
NetServer
NetView
OS/400
POWER
Predictive Failure Analysis
pSeries
RACF
Redbooks
ServeProven
SurePOS
System p5
System z9
Tivoli
Tivoli Enterprise
Tivoli Enterprise Console
Virtualization Engine
Wake on LAN
xSeries
z/VM
zSeries

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Red Hat and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Abbreviations, Acronyms, and Glossary

Abbreviation and acronym list

This topic lists abbreviations and acronyms used in the IBM Director documentation.

Table 43. Abbreviations and acronyms used in IBM Director documentation

Abbreviation or acronym	Definition
AES	advanced encryption standard
APAR	authorized program analysis report
ASF	Alert Standard Format
ASM	Advanced System Management
ASM PCI Adapter	Advanced System Management PCI Adapter
BIOS	basic input/output system
CEC	Central Electronics Complex
CIM	Common Information Model
CIMOM	Common Information Model Object Manager
CP	control program
CRC	cyclic redundancy check
CSM	IBM Cluster Systems Management
CSV	comma-separated value
DASD	direct access storage device
DBCS	double-byte character set
DES	data encryption standard
DHCP	Dynamic Host Configuration Protocol
DIMM	dual inline memory module
DMI	Desktop Management Interface
DMTF	Distributed Management Task Force
DNS	Domain Name System
DSA	Digital Signature Algorithm
EEPROM	electrically erasable programmable read-only memory
FRU	field-replaceable unit
FTMI	fault tolerant management interface
FTP	file transfer protocol
GB	gigabyte
Gb	gigabit
GMT	Greenwich Mean Time
GUI	graphical user interface
GUID	globally unique identifier
HMC	Hardware Management Console
HTML	hypertext markup language
IIS	Microsoft Internet Information Server
I/O	input/output
IP	Internet protocol
IPC	interprocess communication
IPMI	Intelligent Platform Management Interface
IPX	internetwork packet exchange
ISDN	integrated services digital network
ISMP	integrated system management processor
JVM	Java Virtual Machine
JCE	Java Cryptography Extension
JDBC	Java Database Connectivity
JFC	Java Foundation Classes
JRE	Java Runtime Environment
KB	kilobyte
Kb	kilobit
kpbs	kilobits per second
KVM	keyboard/video/mouse
LAN	local area network
LED	light-emitting diode
LPAR	logical partition
MAC	media access control

Table 43. Abbreviations and acronyms used in IBM Director documentation (continued)

Abbreviation or acronym	Definition
MB	megabyte
Mb	megabit
Mbps	megabits per second
MD5	message digest 5
MDAC	Microsoft Data Access Control
MHz	megahertz
MIB	Management Information Base
MIF	Management Information Format
MMC	Microsoft Management Console
MPA	Management Processor Assistant
MPCLI	Management Processor Command-Line Interface
MSCS	Microsoft Cluster Server
MST	Microsoft software transformation
NIC	network interface card
NNTP	Network News Transfer Protocol
NTP	network time protocol
NVRAM	nonvolatile random access memory
ODBC	Open DataBase Connectivity
OID	object ID
PCI	peripheral component interconnect
OSA	Open Systems Adapter
PCI-X	peripheral component interconnect-extended
PDF	Portable Document Format
PFA	Predictive Failure Analysis
POST	power-on self-test
PTF	program temporary fix
RAM	random access memory
RDM	Remote Deployment Manager
RPM	(1) Red Hat Package Manager (2) revolutions per minute
RSA	Rivest-Shamir-Adleman
RXE	Remote Expansion Enclosure

Table 43. Abbreviations and acronyms used in IBM Director documentation (continued)

Abbreviation or acronym	Definition
SAS	Serial Attached SCSI
SATA	Serial ATA
SCSI	Small Computer System Interface
SFS	shared file system
SHA	Secure Hash Algorithm
SI	Solution Install
SID	(1) security identifier (2) Oracle system identifier
SLP	service location protocol
SLPD	service location protocol daemon
SMBIOS	System Management BIOS
SMI	System Management Information
SMP	symmetric multiprocessor
SMS	Systems Management Server
SMTP	Simple Mail Transfer Protocol
SMART	Self-Monitoring, Analysis, and Reporting Technology
SMI-S	Storage Management Initiative Specification
SNMP	Simple Network Management Protocol
SPB	software package block
SQL	Structured Query Language
SSH	Secure Shell
SSL	Secure Sockets Layer
TAP	Telocator Alphanumeric Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
TTL	time to live
UDP	User Datagram Protocol
UID	unique ID
UIM	upward integration module
UNC	universal naming convention
USB	Universal Serial Bus
UUID	universal unique identifier
VPD	vital product data

Table 43. Abbreviations and acronyms used in IBM Director documentation (continued)

Abbreviation or acronym	Definition
VMRM	Virtual Machine Resource Manager
VRM	voltage regulator module
WAN	wide area network
WfM	Wired for Management
WINS	Windows Internet Naming Service
WMI	Windows Management Instrumentation
WQL	Windows Management Instrumentation Query Language
XML	extensible markup language

Glossary

This glossary includes terms and definitions from:

- The *American National Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 1430 Broadway, New York, New York 10018. Definitions are identified by the symbol (A) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Committee (ISO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.
- The *IBM Dictionary of Computing*, New York: McGraw-Hill, 1994.
- The *IBM Glossary of Computing Terms*, 1999.

To view other IBM glossary sources, see IBM Terminology at www.ibm.com/ibm/terminology.

A

Advanced Encryption Setting (AES)
A block cipher algorithm, also known as Rijndael, used to encrypt data transmitted between managed systems and the management server, which employs a key of 128, 192, or 256 bits. AES was developed as a replacement for DES.

Advanced System Management (ASM) interconnect
A feature of IBM service processors that enables users to connect up to 24 servers to one service processor, thus eliminating the need for multiple modems, telephones, and LAN ports. It provides such out-of-band management functions as system power control, service-processor event-log management, firmware updates, alert notification, and user profile configuration.

Advanced System Management (ASM) interconnect network
A network of IBM servers created by using the ASM interconnect feature. The servers are connected through RS-485 ports. When servers containing integrated system management processors (ISMPs) and ASM processors are connected to an ASM interconnect network, IBM Director can manage them out-of-band.

Advanced System Management (ASM) PCI adapter
An IBM service processor that is built into the Netfinity 7000 M10 and 8500R servers. It also was available as an option that could be installed in a server that contained an ASM processor. When an ASM PCI adapter is used with an ASM processor, the ASM PCI adapter acts as an Ethernet gateway, while the ASM processor retains control of the server. When used as a gateway service processor, the ASM PCI adapter can communicate with other ASM PCI adapters and ASM processors only.

Advanced System Management (ASM) processor
A service processor built into the mid-range Netfinity and early xSeries servers. IBM Director can connect out-of-band to an ASM processor located on an ASM interconnect; an ASM PCI adapter, a Remote Supervisor Adapter, or

a Remote Supervisor II must serve as the gateway service processor.

alert A message or other indication that identifies a problem or an impending problem.

alert forwarding
Alert forwarding can ensure that alerts are sent, even if a managed system experiences a catastrophic failure, such as an operating-system failure.

alert-forwarding profile
A profile that specifies where remote alerts for the service processor should be sent.

alert standard format (ASF)
A specification created by the Distributed Management Task Force (DMTF) that defines remote-control and alerting interfaces that can best serve a client system in an environment that does not have an operating system.

anonymous command execution
Execution of commands on a target system as either *system account* (for managed systems running Windows) or *root* (for managed systems running Linux). To restrict anonymous command execution, disable this feature and always require a user ID and password.

ASF See *alert standard format*.

ASM interconnect gateway
See *gateway service processor*.

association
(1) A way of displaying the members of a group in a logical ordering. For example, the Object Type association displays the managed objects in a group in folders based on their type. (2) A way to display additional information about the members of the group. For example, the Event Action Plans association displays any event action plans applied to the managed objects in the group in an Event Action Plan folder.

B

basic input/output system (BIOS)
The code that controls basic hardware operations, such as interactions with diskette drives, hard disk drives, and the keyboard.

BIOS See *Basic Input/Output System*.

blade server
An IBM @server BladeCenter server. A high-throughput, two-way, Intel Xeon-based server on a card that supports symmetric multiprocessors {SMP}.

BladeCenter chassis
A BladeCenter unit that acts as an enclosure. This 7-U modular chassis can contain up to 14 blade servers. It enables the individual blade servers to share resources, such as the management, switch, power, and blower modules.

bottleneck
A place in the system where contention for a resource is affecting performance.

C

chassis
The metal frame in which various electronic components are mounted.

chassis detect-and-deploy profile
A profile that IBM Director automatically applies to all new BladeCenter chassis when they are discovered. The profile settings include management module name, network protocols, and static IP addresses. If Remote Deployment Manager (RDM) is installed on the management server, the chassis detect-and-deploy profile also can include deployment policies.

CIM See *Common Information Model*.

Common Information Model (CIM)
An implementation-neutral, object-oriented schema for describing network management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

component association
In the IBM Director Rack Manager task, a function that can make a managed system or device rack-mountable when the inventory collection feature of IBM Director does not recognize the managed system or device. The function associates the system or device with a predefined component.

D

Data Encryption Standard (DES)
A cryptographic algorithm designed to encrypt and decrypt data using a private key.

database server
The server on which the database application and database used with IBM Director Server are installed.

deployment policy
A policy that associates a specific bay in a BladeCenter chassis with an RDM noninteractive task. When a blade server is added to or replaced in the bay, IBM Director automatically runs the RDM task.

DES See *Data Encryption Standard*.

Desktop Management Interface (DMI)
A protocol-independent set of application programming interfaces (APIs) that were defined by the Distributed Management Task Force (DMTF). These interfaces give management application programs standardized access to information about hardware and software in a system.

Diffie-Hellman key exchange
A public, key-exchange algorithm that is used for securely establishing a shared secret over an insecure channel. During Phase II negotiations, the Diffie-Hellman group prevents someone who intercepts your key from deducing future keys that are based on the one they have.

digital signature algorithm (DSA)
A security protocol that uses a pair of keys (one public and one private) and a one-way encryption algorithm to provide a robust way of authenticating users and systems. If a public key can successfully decrypt a digital signature, a user can be sure that the signature was encrypted using the private key.

discovery
The process of finding resources within an enterprise, including finding the new location of monitored resources that were moved.

DMI See *Desktop Management Interface*.

E

enclosure
A unit that houses the components of a storage subsystem, such as a control unit, disk drives, and power source.

event An occurrence of significance to a task or system, such as the completion or failure of an operation. There are two types of events: alert and resolution.

event action
The action that IBM Director takes in response to a specific event or events.

event-action plan
A user-defined plan that determines how IBM Director will manage certain events. An event action plan comprises one or more event filters and one or more customized event actions.

event-data substitution variable
A variable that can be used to customize event-specific text messages for certain event actions.

event filter
A filter that specifies the event criteria for an event action plan. Events must meet the criteria specified in the event filter in order to be processed by the event action plan to which the filter is assigned.

extension
See *IBM Director extension*.

F

field-replaceable unit (FRU)
An assembly that is replaced in its entirety when any one of its components fails. In some cases, a FRU may contain other FRUs.

file-distribution server
In the Software Distribution task, an intermediate server that is used to distribute a software package when the redirected-distribution method is used.

forecast
A function that can provide a prediction of future performance of a managed system using past data collected on that managed system.

FRU See *field-replaceable unit*.

G

gateway service processor
A service processor that relays alerts from service processors on an Advanced

System Management (ASM) interconnect network to IBM Director Server.

group A logical set of managed objects. Groups can be dynamic, static, or task-based.

GUID See *Universal Unique Identifier*.

I

IBM Director Agent
A component of IBM Director software. When IBM Director Agent is installed on a system, the system can be managed by IBM Director. IBM Director Agent transfers data to the management server using several network protocols, including TCP/IP, NetBIOS, and IPX.

IBM Director Console
A component of IBM Director software. When installed on a system, it provides a graphical user interface (GUI) for accessing IBM Director Server. IBM Director Console transfers data to and from the management server using TCP/IP.

IBM Director database
The database that contains the data stored by IBM Director Server.

IBM Director environment
The complex, heterogeneous environment managed by IBM Director. It includes systems, BladeCenter chassis, software, SNMP devices.

IBM Director extension
A tool that extends the functionality of IBM Director. Some of the IBM Director extensions are Capacity Manager, ServeRAID Manager, Remote Deployment Manager, Software Distribution.

IBM Director Server
The main component of IBM Director software. When installed on the management server, it provides basic functions such as discovery of the managed systems, persistent storage of configuration and management data, an inventory database, event listening, security and authentication, management console support, and administrative tasks.

IBM Director Server service
A service that runs automatically on the management server, and provides the server engine and application logic for IBM Director.

IBM Director service account
The Windows operating-system account associated with the IBM Director Server service.

in-band communication
Communication that occurs through the same channels as data transmissions. An example of in-band communication is the interprocess communication that occurs between IBM Director Server, IBM Director Agent, and IBM Director Console.

integrated system management processor (ISMP)
A service processor built into the some xSeries servers. The successor to the Advanced System Management (ASM) processor, the ISMP does not support in-band communication in systems running NetWare. For IBM Director Server to connect out-of-band to an ISMP, the server containing the ISMP must be installed on an ASM interconnect network. A Remote Supervisor Adapter or a Remote Supervisor Adapter II must serve as the gateway service processor.

interprocess communication (IPC)
1) The process by which programs communicate data to each other and synchronize their activities. Semaphores, signals, and internal message queues are common methods of interprocess communication. 2) A mechanism of an operating system that allows processes to communicate with each other within the same computer or over a network. It also is called in-band communication

inventory-software dictionary
A file that tracks the software installed on managed systems in a network.

IPC See *interprocess communication*.

ISMP See *integrated system management processor*.

J

job A separately executable unit of work defined by a user, and run by a computer.

L

Level-0 managed system
An IBM or non-IBM server, desktop computer, workstation, or mobile computer, that can be managed by IBM Director but does not have any IBM Director software installed on it.

Level-1 managed system
An IBM or non-IBM server, desktop computer, workstation, and mobile computer that has IBM Director Core Services installed. IBM Director uses IBM Director Core Services to communicate with and administer the Level-2 managed system. IBM Director Core Services includes the SLP instrumentation, the IBM Director Agent SLP service type, and Common Information Model (CIM).

Level-2 managed system
An IBM or non-IBM server, desktop computer, workstation, or mobile computer that has IBM Director Agent installed. IBM Director Agent provides managed systems with the full complement of IBM Director Agent function that is used to communicate with and administer the Level-2 managed system. The function of a Level-2 managed system varies depending on the operating system and platform.

light path diagnostics
A technology that provides a lighted path to failed or failing components to expedite hardware repairs.

M

MAC address
See media access control (MAC) address.

managed group
A group of systems or objects managed by IBM Director.

managed object
An item managed by IBM Director. In IBM Director Console, a managed object is represented by an icon that shows its type (such as chassis, cluster, system, or scalable system, for example).

managed object ID
A unique identifier for each managed object. It is the key value used by IBM Director database tables.

managed system
A system that is being controlled by a given system management application, for example, a system managed by IBM Director.

management console
A system (server, desktop computer, workstation, or mobile computer) on which IBM Director Console is installed.

management module
The BladeCenter component that handles system-management functions. It configures the chassis and switch modules, communicates with the blade servers and all I/O modules, multiplexes the keyboard/video/mouse (KVM), and monitors critical information about the chassis and blade servers.

management server
The server on which IBM Director Server is installed.

media access control (MAC) address
In a local area network, the protocol that determines which device has access to the transmission medium at a given time.

N

nonvolatile random-access memory (NVRAM)
Random access memory (storage) that retains its contents after the electrical power to the machine is shut off.

notification
See *alert*.

NVRAM
See *nonvolatile random-access memory*.

O

out-of-band communication
Communication that occurs through a modem or other asynchronous connection, for example, service processor alerts sent through a modem or over a LAN. In an IBM Director environment, such communication is independent of the operating system and interprocess communication (IPC).

P

partition
See *scalable partition*.

PCI See *Peripheral Component Interconnect*.

PCI-X See *Peripheral Component Interconnect-X*.

Peripheral Component Interconnect (PCI)
A standard for connecting attached devices to a computer.

Peripheral Component Interconnect-X (PCI-X)
An enhancement to the Peripheral Component Interconnect (PCI) architecture. PCI-X enhances the Peripheral Component Interconnect (PCI) standard by doubling the throughput capability and providing additional adapter-performance options while maintaining backward compatibility with PCI adapters.

PFA See *Predictive Failure Analysis*.

physical platform
An IBM Director managed object that represents a single physical chassis or server that has been discovered through the use of the Service Location Protocol (SLP).

plug-in
A software module, often written by a third party, that adds function to an existing program or application such as a Web browser. See *IBM Director extension*.

POST See *power-on self-test*.

power-on self-test
A series of internal diagnostic tests activated each time the system power is turned on.

Predictive Failure Analysis (PFA)
A scheduled evaluation of system data that detects and signals parametric degradation that might lead to functional failures.

private key
1) In secure communication, an algorithmic pattern used to encrypt messages that only the corresponding public key can decrypt. The private key is also used to decrypt messages that were encrypted by the corresponding public key. The private key is kept on the user's system and is protected by a password. 2) The secret half of a cryptographic key pair that is used with a public key algorithm. Private keys are typically used to digitally sign data and to decrypt data that has been encrypted with the corresponding public key.

public key
1) In secure communication, an algorithmic pattern used to decrypt messages that were encrypted by the corresponding private key. A public key is also used to encrypt messages that can be decrypted only by the corresponding private key. Users broadcast their public keys to everyone with whom they must exchange encrypted messages. 2) The non-secret half of a cryptographic key pair that is used with a public key algorithm. Public keys are typically used to verify digital signatures or decrypt data that has been encrypted with the corresponding private key.

R

redirected distribution
A method of software distribution that uses a file-distribution server.

remote I/O enclosure
An IBM Director managed object that represents an expansion enclosure of Peripheral Component Interconnect-X (PCI-X) slots, for example, an RXE-100 Remote Expansion Enclosure. The enclosure consists of one or two expansion kits.

Remote Supervisor Adapter
An IBM service processor. It is built into some xSeries servers and available as an optional adapter for use with others. When used as a gateway service processor, the Remote Supervisor Adapter can communicate with all service processors on the Advanced System Management (ASM) interconnect.

resolution
The occurrence of a correction or solution to a problem.

resource-monitor threshold
The point at which a resource monitor generates an event.

RXE Expansion Port
The dedicated high-speed port used to connect a remote I/O expansion unit, such as the RXE-100 Remote Expansion Enclosure, to a server.

S

scalable node
A physical platform that has at least one SMP Expansion Module. Additional attributes are assigned to a physical platform when it is a scalable node. These additional attributes record the number of SMP Expansion Modules, SMP Expansion Ports, and RXE Expansion ports on the physical chassis.

scalable object
An IBM Director managed object that is used with Scalable Systems Manager. Scalable objects include scalable nodes, scalable systems, scalable partitions, and remote I/O enclosures that are attached to scalable nodes.

scalable partition
An IBM Director managed object that defines the scalable nodes that can run a single image of the operating system. A scalable partition has a single, continuous memory space and access to all associated adapters. A scalable partition is the logical equivalent of a physical platform. Scalable partitions are associated with scalable systems and comprise only the scalable nodes from their associated scalable systems.

scalable system
An IBM Director managed object that consists of scalable nodes and the scalable partitions that are composed of the scalable nodes in the scalable system. When a scalable system contains two or more scalable nodes, the servers that they represent must be interconnected through their SMP Expansion Modules to make a multinode configuration, for example, a 16-way xSeries 455 server made from four scalable nodes.

Secure Sockets Layer (SSL)
A security protocol that provides communication privacy. SSL enables client/server applications to communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

Service Location Protocol (SLP)
In the Internet suite of protocols, a protocol that identifies and uses network hosts without having to designate a specific network host name.

service processor
A generic term for Remote Supervisor Adapters, Advanced System Management processors, Advanced System Management PCI adapters, and integrated system management processors (ISMPs). These hardware-based management processors used in IBM Netfinity and xSeries servers work with IBM Director to provide hardware status and alert notification.

SLP See *Service Location Protocol*.

SMBIOS
See *systems management BIOS*.

SMP Expansion Module
An IBM xSeries hardware option. It is a single module that contains microprocessors, disk cache, random access memory, and three SMP Expansion Port connections. Two SMP Expansion Modules can fit in a chassis.

SNMP Access and Trap Forwarding
An IBM Director Agent feature that enables SNMP to access managed-system data. When installed on a managed system, this feature enables SNMP-based managers to poll the managed system and receive its alerts. If System Health Monitoring is installed on the managed system also, hardware alerts can be forwarded as SNMP traps.

SNMP device
A network device, printer, or computer that has an SNMP device installed or embedded.

SQL See *Structured Query Language*

SSL See *Secure Sockets Layer*.

static partition
A view-only scalable partition.

sticky key
An input method that enables the user to press and release a series of keys sequentially (for example, Ctrl+Alt+Del), yet have the keys behave as if they were pressed and released at the same time. This method can be used for those who require special-needs settings to make the keyboard easier to use.

Structured Query Language (SQL)
A standardized language for defining and manipulating data in a relational database.

switch module
The BladeCenter component that provides network connectivity for the BladeCenter chassis and blade servers. It also provides interconnectivity between the management module and blade servers.

system
The computer and its associated devices and programs.

System Health Monitoring
An IBM Director Agent feature that provides active monitoring of critical system functions, including system temperatures, voltages, and fan speeds. It also handles in-band alert notification for managed systems running Windows and some managed systems running Linux.

system variable
A user-defined keyword and value pair that can be used to test and track the status of network resources. System variables can be referred to wherever event-data substitution is allowed.

systems management BIOS (SMBIOS)
A key requirement of the Wired for Management (WfM) 2.0 specification. SMBIOS extends the system BIOS to support the retrieval of management data required by the WfM specification. To run IBM Director Agent, a system must support SMBIOS, version 2.2 or later.

T

target system
A managed system on which an IBM Director task is performed.

time to live (TTL)
A technique used by best-effort delivery protocols to inhibit endlessly looping packets. The packet is discarded if the TTL counter reaches 0.

triple data encryption standard (DES)
A block cipher algorithm that can be used to encrypt data transmitted between managed systems and the management server. Triple DES is a security enhancement of DES that employs three successive DES block operations.

TTL See *time to live*.

U

universal unique identifier (UUID)
A 128-bit character string guaranteed to be globally unique and used to identify components under management.

uptime
The time during which a system is working without failure.

upward integration
The methods, processes and procedures that enable lower-level systems-management software, such as IBM Director Agent, to work with higher-level systems-management software, such as Tivoli Enterprise™ or Microsoft SMS.

upward integration module
Software that enables higher-level systems-management software, such as Tivoli Enterprise or Microsoft Systems Manager Server (SMS), to interpret and display data provided by IBM Director Agent. This module also can provide enhancements that start IBM Director Agent from within the higher-level systems-management console, as well as collect IBM Director inventory data and view IBM Director alerts.

UUID See *universal unique identifier*.

V

vital product data (VPD)
Information that uniquely defines the system, hardware, software, and microcode elements of a processing system.

VPD See *vital product data*.

W

Wake on LAN
A technology that enables a user to remotely turn on systems for off-hours maintenance. A result of the Intel-IBM Advanced Manageability Alliance and part of the Wired for Management Baseline Specification, users of this technology can remotely turn on a server and control it across the network, thus

saving time on automated software installations, upgrades, disk backups, and virus scans.

walk An SNMP operation that is used to discover all object instances of management information implemented in the SNMP agent that can be accessed by the SNMP manager.

Windows Management Instrumentation (WMI)
An application programming interface (API) in the Windows operating system that enables devices and systems in a network to be configured and managed. WMI uses the Common Information Model (CIM) to enable network administrators to access and share management information.

WMI See *Windows Management Instrumentation*.

WMI Query Language (WQL)
A subset of the Structured Query Language with minor semantic changes to support Windows Management Instrumentation.

WQL See *WMI Query Language*.

Index

Readers' Comments — We'd Like to Hear from You

IBM Systems
IBM Director
Systems Management Guide
Version 5.10

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	□	□	□	□	□

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	□	□	□	□	□
Complete	□	□	□	□	□
Easy to find	□	□	□	□	□
Easy to understand	□	□	□	□	□
Well organized	□	□	□	□	□
Applicable to your tasks	□	□	□	□	□

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? □ Yes □ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name	Address
Company or Organization	
Phone No.	

Readers' Comments — We'd Like to Hear from You

Cut or Fold
Along Line

Fold and Tape **Please do not staple** Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Dept. CGFA
PO Box 12195
Research Triangle Park, NC 27709-9990

Fold and Tape **Please do not staple** Fold and Tape

Cut or Fold
Along Line

Printed in USA