IBM Director 4.11

IBM

Scalable Systems Manager 4.11 Installation and User's Guide

> **Note**
>
> Before using this information and the product it supports, read the general information in Appendix D, "Notices", on page 99.

Contents

Figures

Tables

Preface

This book describes IBM® Scalable Systems Manager 4.11 and how to use it to view, create, and manage hardware partitions for supported IBM @server™ xSeries™ servers.

How this book is organized

Chapter 1, "Introducing Scalable Systems Manager", on page 1, contains an overview of Scalable Systems Manager, including its components and features, its integration with IBM Director, and its icons. This chapter also describes the features of IBM Director that are relevant to Scalable Systems Manager.

Chapter 2, "Installing Scalable Systems Manager", on page 13, details the system requirements for Scalable Systems Manager and how to install and uninstall Scalable Systems Manager.

Chapter 3, "Scalable objects", on page 19 describes IBM Director physical-platform discovery, the types of scalable objects, the supported configurations for scalable systems, and how to save and restore scalable information in IBM Director Server.

Chapter 4, "Using the Scalable Systems Discovery task", on page 35 describes how to automatically create and update scalable systems according to their hardware configuration.

Chapter 5, "Using the Scalable Systems Validation task", on page 39 describes how to validate the cabling of all scalable nodes in a scalable system.

Chapter 6, "The Scalable Systems Configuration task", on page 43 describes how to start Scalable Systems Configuration and explains how to use the graphical user interface.

Chapter 7, "Managing scalable objects", on page 49, describes how to use Scalable Systems Configuration task to perform operations on scalable objects, such as creating scalable partitions, performing power operations, and checking for problems.

Chapter 8, "Viewing details about scalable objects in Scalable Systems Configuration", on page 71 describes how to use Scalable Systems Configuration to display details about scalable objects and remote I/O enclosures.

Chapter 9, "Solving Scalable Systems Manager problems", on page 87 lists some of the problem symptoms and suggested solutions for Scalable Systems Manager.

Appendix A, "Error codes", on page 91 describes the error codes that Scalable Systems Manager can return through error messages displayed in Scalable Systems Configuration.

Appendix B, "Terminology summary and abbreviation list", on page 95 contains a summary of Scalable Systems Manager terminology and a list of abbreviations that are used in Scalable Systems Manager publications.

Appendix C, "Getting help and technical assistance", on page 97 contains information about accessing IBM Support Web sites for help and technical assistance.

Appendix D, "Notices", on page 99 contains product notices and trademarks.

The glossary on page 101 provides definitions for terms used in Scalable Systems Manager publications.

Notices that are used in this book

This book contains the following notices 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 is placed just before the instruction or situation in which damage could occur.

Scalable Systems Manager publications

The Web site for Scalable Systems Manager provides the following publication in Adobe Acrobat Portable Document Format (PDF):
- *Scalable Systems Manager 4.11 Installation and User's Guide*

You can also obtain this document from the IBM Support Web site at http://www.ibm.com/pc/support/ .

Scalable Systems Manager resources on the World Wide Web

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

IBM Support page
> http://www.ibm.com/pc/support/
>
> This is the IBM Support Web site for IBM hardware and systems-management software. For systems-management software support, click **Systems management**.

IBM Systems Management Software: Download/Registration page
> http://www.ibm.com/pc/us/eserver/xseries/systems_management/dwnl.html
>
> Use this Web page to download IBM systems-management software, including IBM Director.

IBM xSeries Systems Management page
> http://www.ibm.com/pc/ww/eserver/xseries/systems_management/index.html
>
> This Web page presents an overview of IBM systems management and IBM Director. Click **IBM Director 4.1** for the latest information and publications.

IBM ServerProven® page
> http://www.ibm.com/pc/us/compat/index.html
>
> This Web page provides information about IBM hardware compatibility with IBM systems-management software.

Chapter 1. Introducing Scalable Systems Manager

IBM Scalable Systems Manager is an extension to IBM Director. It communicates out-of-band with the service processors in supported servers to manage hardware partitions. These hardware partitions are either *static partitions* that consists of xSeries 440 16-way servers or *scalable partitions* that consists of xSeries 445 servers; they can each run a single image of the operating system. Static partitions are presented as view-only scalable partitions in Scalable Systems Manager. The servers defined in a scalable partition have at least one SMP Expansion Module and at least one RXE Expansion Port and are referred to as *scalable nodes*. A *scalable system* consists of one or two scalable nodes and the scalable partitions that are made from those scalable nodes.

These IBM Director managed objects are referred to as *scalable objects* throughout this documentation. Scalable objects also include the remote I/O enclosures that are attached to scalable nodes. For detailed explanations of scalable objects and how they are used in Scalable Systems Manager, see Chapter 3, "Scalable objects", on page 19. For definitions of these individual objects, see Appendix B, "Terminology summary and abbreviation list", on page 95.

Scalable Systems Manager provides support for specific hardware as follows:

- You can view information about predefined scalable systems and static partitions that were saved in NVRAM by the Configuration/Setup Utility program (in the BIOS code) on xSeries 440 servers.
- You can view information about predefined scalable systems and their scalable partitions that were saved in NVRAM by the Configuration/Setup Utility program on xSeries 445 servers.
- You can create, configure, and manage scalable systems and scalable partitions that are created from xSeries 445 servers.

  Note: You cannot manage scalable systems and static partitions that are created from xSeries 440 servers. See "Differences between manageable and view-only scalable systems" on page 23.

- You can configure RXE-100 Remote Expansion Enclosures that are attached to xSeries 445 servers.

Because Scalable Systems Manager communicates out-of-band with the supported servers, the primary task that you can perform with Scalable Systems Manager is to configure a scalable partition for an offline server before installing its operating system.

Important: Scalable Systems Manager uses the communication methods provided by Management Processor Assistant (MPA) to communicate out-of-band with service processors on supported servers. For this reason, IBM Director and Scalable Systems Manager cannot communicate out-of-band unless the following conditions are met:

- Service-processor networks must either have static IP address assignments or have Dynamic Host Configuration Protocol (DHCP) configured to maintain consistent IP addresses for the service processors. To do so, create reservations in DHCP that identify explicit IP addresses for the service processors.
- The IP addresses that are assigned to the service processors cannot change after the servers are discovered in IBM Director. This

limitation is true whether you are using DHCP or statically assigning IP addresses. Chapter 9, "Solving Scalable Systems Manager problems", on page 87 describes suggested actions to take if the IP address of a service processor has changed since it was discovered by IBM Director.

If Scalable Systems Manager cannot communicate with a service processor, use MPA to make sure that out-of-band communications with the service processor are working properly. For more information about communicating with service processors in an IBM Director environment, see the *IBM Director 4.11 Installation and Configuration Guide*.

After you power-on a scalable partition out-of-band with Scalable Systems Manager, if 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. This managed system can be managed in-band in IBM Director as any other managed system can. For example, you can use Capacity Manager to monitor and forecast managed-system performance or Software Distribution to distribute software packages to the managed system.

Note: The illustrations in this document might differ slightly from those shown by Scalable Systems Manager.

Overview of scalable partition configuration

Figure 1 shows an overview of how to configure scalable partitions.

```
                    ┌─────────────────────────────┐
                    │      Cable hardware          │
                    └─────────────────────────────┘
                                 │
                                 ▼
             ┌────────────────────────────────────────┐
             │  Configure Remote Supervisor Adapters   │
             │     for out-of-band communication       │
             └────────────────────────────────────────┘
                                 │
                                 ▼
                    ┌─────────────────────────────┐
                    │   Start IBM Director Console │
                    └─────────────────────────────┘
                                 │
                                 ▼
         ┌──────────────────────────────────────────────┐
         │        Discover physical platforms:           │
         │  Click Tasks ▶ Discover Systems ▶ Physical     │
         │                Platforms                       │
         └──────────────────────────────────────────────┘
```

Did discovered physical platforms unlock?

No → Request access to discovered physical platforms:
• Interrogates server NVRAM for partition descriptor
• Discovers attached remote I/O enclosures

Yes

Were scalable systems previously configured by BIOS?

No → Run the Scalable Systems Discovery task to discover scalable systems not defined by BIOS

Yes

Run the Scalable Systems Validation task to verify scalability cabling

Run the Scalable Systems Configuration task to configure scalable partitions:
• Create scalable partitions
• Insert scalable nodes into scalable partitions
• Set primary scalable nodes
• Configure remote I/O enclosures

Power on the scalable partitions from Scalable Systems Configuration

Scalable partitions are set up and ready for use in IBM Director Console

Figure 1. Configuring scalable partitions overview

Scalable Systems Manager tasks

The console component of Scalable Systems Manager is an extension to IBM Director that provides the following console tasks:

- Scalable Systems Configuration
- Scalable Systems Discovery
- Scalable Systems Validation

The following illustration shows IBM Director Console window with Scalable Systems Manager and its tasks in the Tasks pane.
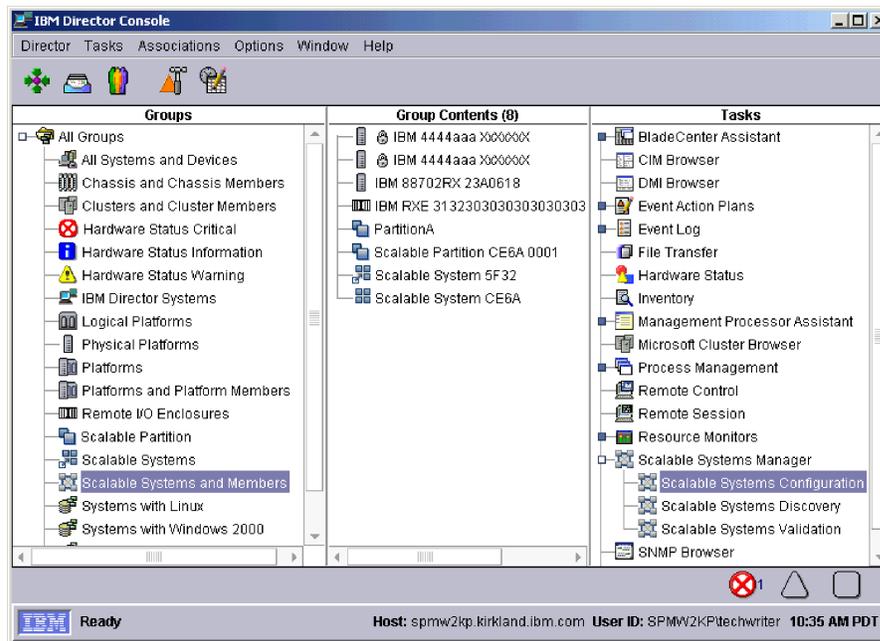
Figure 2. Displaying scalable systems and members in IBM Director Console

Scalable Systems Configuration task

You can use the Scalable Systems Configuration task to view and manage scalable objects in the graphical user interface provided by Scalable Systems Manager. From the interface, you can perform the following tasks:

- View detailed information about scalable objects
- Create and delete scalable objects
- Add and delete unassigned scalable nodes
- Insert and remove scalable nodes into and from a scalable partition
- Configure the remote I/O enclosures in a scalable partition
- Perform power operations on scalable partitions, such as power-on, shut down, and power-off immediately

Scalable Systems Discovery task

For servers that are cabled together, you can use the Scalable Systems Discovery task to perform the following tasks automatically:

- Create a manageable scalable system of the unassigned scalable nodes that represent servers that are cabled together

- Update a manageable scalable system to add unassigned scalable nodes that are cabled to the scalable nodes that are already defined in that manageable scalable system

The Scalable Systems Discovery task does not create or update view-only scalable systems.

The Scalable Systems Discovery task must target one or more unassigned scalable nodes or one or more scalable systems. The Scalable Systems Discovery task is a non-interactive task that you can schedule or run immediately.

Scalable Systems Validation task

You can use the Scalable Systems Validation task to validate the cabling of one or more manageable scalable systems. The Scalable Systems Validation task is a non-interactive task that you can schedule or run immediately.

For each targeted scalable system, Scalable Systems Manager communicates with the service processor of each scalable node in that scalable system. Scalable Systems Manager issues a ping command to each scalability port on each scalable node and waits for the responses. Scalable Systems Manager determines whether the resulting responses correspond to a supported configuration of a manageable scalable system. The Scalable Systems Validation task does not validate the cabling for view-only scalable systems.

Icons in Scalable Systems Manager

Scalable Systems Manager uses icons from IBM Director and its own unique icons to depict information about scalable objects and their status.

Scalable object icons

Table 1 describes the icons that are used for various scalable objects in IBM Director and Scalable Systems Configuration.

Table 1. Scalable object icons

Icon	Description
	Scalable Systems Manager and its subtasks. This icon is also used to identify the Scalable Systems and Members group in IBM Director Console and to identify each help topic that is provided by Scalable Systems Manager.
	A manageable scalable system. This icon is also used to identify the Scalable Systems group in IBM Director Console. See "Scalable systems" on page 23 for a description of manageable scalable systems.
	A view-only scalable system. See "Scalable systems" on page 23 for a description of view-only scalable systems.
	A scalable node that is part of a powered-on scalable partition. The startup disk icon is displayed beside the scalable-node icon when it is the primary scalable node. IBM Director Console uses this same icon to depict all physical platforms, including those that are not scalable nodes and those that are not in powered-on scalable partitions.
	A scalable node that is unassigned or is part of a scalable partition that is powered-off (inactive). This icon is not used by IBM Director Console.
	A scalable partition that is powered-on (active). This icon is not used by IBM Director Console.

Table 1. Scalable object icons (continued)

Icon	Description
	A scalable partition that is powered-off (inactive). IBM Director Console uses this icon to depict all scalable partitions whether they are powered-on or powered-off. Furthermore, IBM Director Console uses additional icons with this icon to indicate the state of a scalable partition. For more information, see "States of scalable partitions" on page 83.
	A remote I/O enclosure that has both PCI-X expansion kit A and B attached to the corresponding scalable node. This icon is also used to identify all remote I/O enclosures in IBM Director Console, regardless of the expansion kits the enclosure is configured to use.
	A remote I/O enclosure that has PCI-X expansion kit A attached to the corresponding scalable node. This icon is not used by IBM Director Console.
	A remote I/O enclosure that has PCI-X expansion kit B attached to the corresponding scalable node. This icon is not used by IBM Director Console.

Other icons

Scalable Systems Manager also uses icons within the Scalable Systems Configuration window. For details about these icons, see "The Scalable Systems Configuration window" on page 44 and "Invoking Scalable Systems Configuration functions" on page 47.

IBM Director features relevant to scalable objects

IBM Director includes several features that are relevant to scalable objects.

Groups used with scalable objects

IBM Director provides several default groups of managed objects in the Groups pane for easier management of these objects. The following default groups are relevant to scalable objects.

Table 2. IBM Director groups used with scalable objects

Group name	Includes these 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.
Remote I/O Enclosures	Only remote I/O enclosures.
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.

User Administration

You can use the security features of IBM Director Console to configure or restrict the users that can run the Scalable Systems Manager tasks and that can perform specific operations. For example, you can create an operator user that cannot perform any power-on or power-off operations in IBM Director or Scalable Systems Manager.

To configure user defaults, click **Options** → **User Administration** in IBM Director Console, which causes the User Administration window to open. From the User Administration window, select a user from the list and click **User** → **User Defaults** to open the User Defaults Editor window.

From the User Defaults Editor window in IBM Director, you can:

- Control the privileges of users as they apply to Scalable Systems Manager. The following user default privileges are relevant to Scalable Systems Manager:
  – Allow discovery requests
  – Allow power down of systems
  – Allow power on of systems
  – Allow shutdown of systems
  – Allow system create/modify/delete operations
- Limit user access to specific IBM Director groups, such as Scalable Systems and Members.
- Limit user access to the Scalable Systems Manager extension and its tasks.

For more information about user administration, see the *IBM Director 4.11 System Management Guide*.

Associations used with scalable objects

You can use IBM Director associations to display the scalable objects of a group in the Group Contents pane in a logical ordering, and if applicable, in a tree structure. The following associations are relevant to scalable objects.

Table 3. IBM Director associations used with scalable objects

Association name	Alphabetically orders objects in group by
Physical Platform - Remote I/O Enclosures	Physical platforms and scalable nodes. If the scalable node has an attached remote I/O enclosure, a tree structure depicts the attached remote I/O enclosure.
Platform Membership	Platforms, which includes both logical platforms, such as scalable partitions, and physical platforms, such as scalable nodes. If the platform is associated with a managed system, a tree structure depicts the related managed system.
Scalable Partitions Membership	Scalable partitions and scalable nodes and their associations with any managed systems.
Scalable Systems Membership	Scalable systems, each with a tree structure that lists the scalable nodes and the scalable partitions that are associated with the scalable system. For each scalable partition in the scalable system, the scalable nodes in the scalable partition and their attached remote I/O enclosures are listed in a tree structure. If the scalable partition is powered-on, the resulting managed system is shown as a child of the scalable partition.
Status	Status folders. The Scalable Partition Power Status folder groups scalable partitions by their current power state.

Event filters and actions for use with scalable objects

Scalable Systems Manager provides several events and one event action for use with scalable partitions. For example, with the events and event actions that are provided by Scalable Systems Manager, you could create an event action plan that automatically powers-on or shuts down and powers-off a scalable partition when a scalable partition fails. Management Processor Assistant (MPA) also provides events that are relevant to scalable nodes and scalable partitions.

Important: Before using event filters and actions with scalable objects, use the Management Processor Configuration subtask of MPA to configure an alert-forwarding profile for the servers that are represented by the scalable nodes that you plan to target. In the alert-forwarding profile, the connection type must be set to "IBM Director Comprehensive" and the IP address must be set to the IP address of the management server that is being used to manage scalable objects.

For more information about event filters, event actions, and the Event Filter Builder, see the IBM Director help, the *IBM Director 4.11 System Management Guide*, and the *IBM Director 4.1 Events Reference*.

Events provided by Management Processor Assistant

Management Processor Assistant (MPA) provides events for both scalable nodes and scalable partitions. MPA events are under the **MPA** event type in the Event Filter Builder. For more information about MPA events, see the *IBM Director 4.1 Events Reference*.

The MPA events for a scalable node correlate to the state of a scalable node and the mode of a scalable node. These events are under the **Component** subcategory and its **Server** subcategory and under the **Platform** subcategory and its **Scalable node** subcategory. For more information, see "States of scalable nodes" on page 75 and "Modes of scalable nodes" on page 76.

The MPA events for a scalable partition relate to the state of the scalable partition. These events are under the **Platform** subcategory and its **Scalable partition** subcategory. For more information, see "States of scalable partitions" on page 83.

Events provided by Scalable Systems Manager

The events that are provided by Scalable Systems Manager are for detecting events related to scalable partitions. These events are under the **SSM** event type in the Event Filter Builder and its **Scalable partition** subcategory.

The following Event Filter Builder window shows the Scalable Systems Manager events.
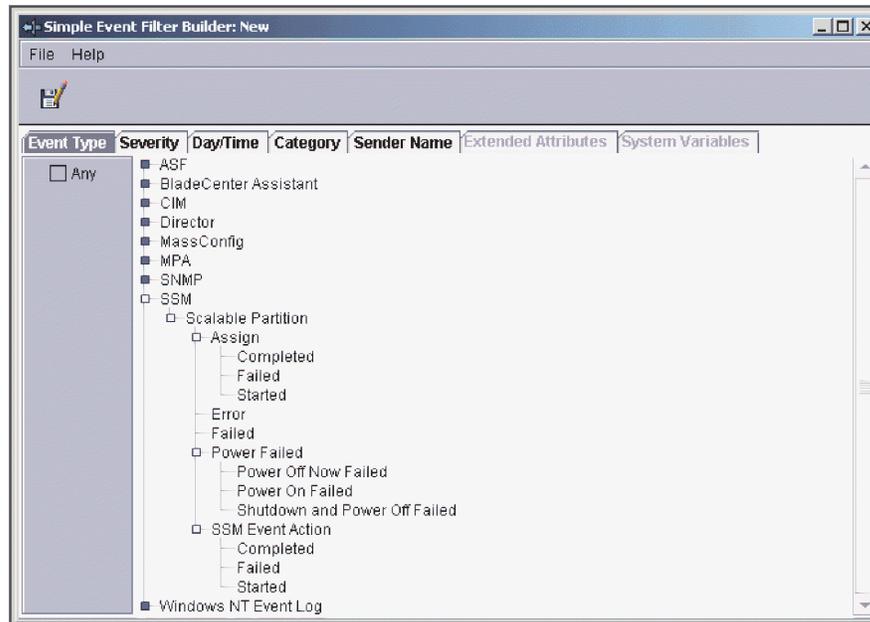
Figure 3. "Simple Event Filter Builder" window showing Scalable System Manager events

Table 4 defines the events that are provided with Scalable Systems Manager.

Table 4. Scalable Systems Manager events

Scalable partition event subcategory	Additional event subcategory	Invoked when
Assign	Completed	Scalable Systems Manager has successfully written the partition descriptor for the scalable partition to NVRAM on the service processor of the primary scalable node.
	Failed	Scalable Systems Manager could not write the partition descriptor for the scalable partition to NVRAM on the service processor of the primary scalable node.
	Started	Scalable Systems Manager has started to write the partition descriptor for the scalable partition to NVRAM on the service processor of the primary scalable node.
Error		One of the following errors occurs for a scalable partition: • A scalable node attempts to be the primary scalable node, but its scalable partition recognizes a different scalable node as primary. • A scalable node attempts to be a secondary scalable node, but its scalable partition recognizes a different scalable node as secondary. • The Escape key is pressed on the keyboard of the primary or secondary scalable node while it is starting so that BIOS can be set-up or flashed. • The primary scalable node could not start the operating system on the scalable partition, which means that the scalable node is now in stand-alone reset mode.
Power Failed	Power Off Now Failed	A scalable partition could not be successfully powered-off, for example, because a server in the scalable partition would not turn off.

Table 4. Scalable Systems Manager events (continued)

	Power On Failed	A scalable partition could not be successfully powered-on. This can happen when the service processor fails or the Ethernet connection to the service processor fails. In both cases, the power-on request cannot be communicated to IBM Director Server. The power-on request can also fail when the partition descriptor on IBM Director Server contains bad data.
	Shut Down and Power Off Failed	A scalable partition could not be successfully shut down and powered-off, for example, because the operating system would not shut down. Another example is when a server in the scalable partition would not turn off.
SSM Event Action	Completed	The "Power on/shut down and power off scalable partition" action has been completed.
	Failed	The "Power on/shut down and power off scalable partition" action has failed.
	Started	The "Power on/shut down and power off scalable partition" action has started.

By default, when you use these events in event filters, IBM Director triggers these events for all scalable partitions. To trigger these events only for certain scalable partitions, click the **Extended Attributes** tab of the Event Filter Builder. From this page, you can designate the names of one or more specific scalable partitions by entering each name in the **Values** field and then clicking **Add**. Repeat as necessary until you have selected all the scalable partitions that you want to add.

By using these scalable-partition events and the Scalable Systems Manager action "Power on/shut down and power off scalable partition," you can create event action plans that can be used to automatically power-on or shut down and power-off a scalable partition according to a particular event such as a failed scalable partition.

Using event actions with scalable partitions

Scalable Systems Manager adds one event action template to IBM Director Event Action Plan Builder. This template, "Power on/shut down and power off scalable partition," is used to create custom event actions. To use this template, right-click the template name; then, click **Customize**.

By creating custom event actions, you can specify which action you want IBM Director to take as a result of the occurrence of an event that is triggered by an event filter. You create event filters by using Scalable Systems Manager events. After you have created custom event actions and event filters, you can create an event action plan that contains specific filters and their associated actions. For example, you can create an event action plan that automatically powers-on a scalable partition when another scalable partition fails.

The "Power on/shut down and power off scalable partition" event action template has the following two fields that must be defined when you create a custom event action:

Scalable partition name
> Designates the scalable partition to be acted upon by the operation that is selected in the **Function** field. From the list of available scalable partitions, click the name of the scalable partition that you want to use.

Function
      Designates whether to power-on or shut down and power-off the scalable
      partition that you selected in the first field.

To use a custom event action, you must add it to an event filter that is already in an
event action plan.

Chapter 2. Installing Scalable Systems Manager

This chapter provides information about system requirements, supported systems, and instructions for installing and uninstalling Scalable Systems Manager.

System requirements

In addition to the requirements that are applicable to IBM Director 4.11, Scalable Systems Manager has further requirements applicable to the hardware and software with which it can be used.

Hardware requirements for management servers and consoles

Scalable Systems Manager can be installed on any management server or management console that supports IBM Director 4.11. However, Scalable Systems Manager has these additional hardware requirements beyond those of IBM Director 4.11:

- Management servers need 2.0 MB of additional disk space
- Management consoles need 1.5 MB of additional disk space

Operating systems for management servers and consoles

Scalable Systems Manager 4.11 is supported for use only with IBM Director 4.11 management servers and management consoles that are running the following operating systems:

- Microsoft® Windows® XP Professional (Service Pack 1 recommended)
- Windows 2000 Professional, Server, and Advanced Server (Service Pack 3 required)
- Windows Server 2003 (Standard, Enterprise, and Web Editions)

Note: In an IBM Director environment that has multiple management servers, use only one management server to manage scalable objects with Scalable Systems Manager. Using multiple management servers to manage scalable objects causes unpredictable results, such as being unable to power-on scalable partitions.

Requirements for out-of-band servers

Scalable partition operations in Scalable Systems Manager are supported for these out-of-band servers and operating systems:

- xSeries 445 (Intel Xeon MP Microprocessors only)
  - Windows 2000 Server and Advanced Server (Service Pack 3 required)
  - Windows 2000 Datacenter Server
  - Windows Server 2003 (Standard, Enterprise, and Datacenter Editions)
  - Red Hat Enterprise Linux AS 2.1
  - VMware ESX Server 2

  For more information, see the table of operating systems and the configurations that they support in the *xSeries 445 User's Guide* on the IBM *xSeries Documentation* CD.
- xSeries 440 (16-way configurations only)
  - Windows 2000 Datacenter Server
  - Windows Server 2003 Datacenter Edition

Additional requirements for out-of-band servers in 16-way configurations

The following additional requirements apply to out-of-band servers in 16-way configurations:

- Both servers 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 needed so that the service processors can communicate and perform the necessary functions for the two servers to merge as one combined server or unmerge as two separate servers. This connection is also required for out-of-band communications with IBM Director and Scalable Systems Manager.

Installation prerequisites

Before installing Scalable Systems Manager, make sure that the following prerequisites are met:

- IBM Director Server is installed on the management server and IBM Director Console is installed on any systems from which an administrator will access IBM Director.
- The operating-system account that is used to install Scalable Systems Manager must have at least local Administrator authority.

If these prerequisites are not met, the Scalable Systems Manager installation program will not install Scalable Systems Manager.

Installing the Scalable Systems Manager extension

The Scalable Systems Manager installation program detects which IBM Director software components (IBM Director Server and IBM Director Console) are installed on a system and automatically installs the matching Scalable Systems Manager components.

Complete the following steps to install Scalable Systems Manager:

1. Go to http://www.ibm.com.
2. Click **Support & Downloads**.
3. Click **All downloads & drivers**.
4. In the **Enter search terms** field, type `Scalable Systems Manger 4.11` and click Submit.
5. Click **Systems Management - Scalable Systems Manger 4.11**.
6. Click the file link for the executable file to download the file.

   Additionally, you can download the readme.txt file and the *Scalable Systems Manager Installation and User's Guide* from this Web page.
7. If necessary, copy the downloaded file to a shared network drive so that you can access the file from the out-of-band systems.
8. If IBM Director Console is running, close it.
9. From the system on which you want to install Scalable Systems Manager, run the executable file that you downloaded.

   The Scalable Systems Manager Setup program starts, and the "Welcome to the InstallShield Wizard for IBM Scalable Systems Manager" window opens.
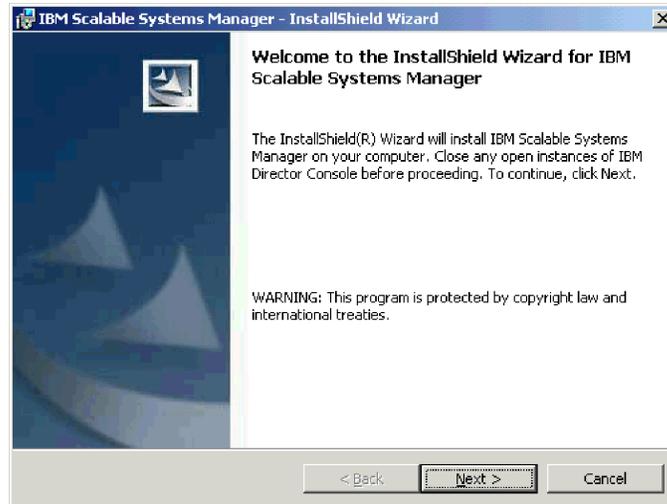
Figure 4. Installing Scalable Systems Manager: "Welcome to the InstallShield wizard" window

10. Click **Next**. The "License Agreement" window opens.
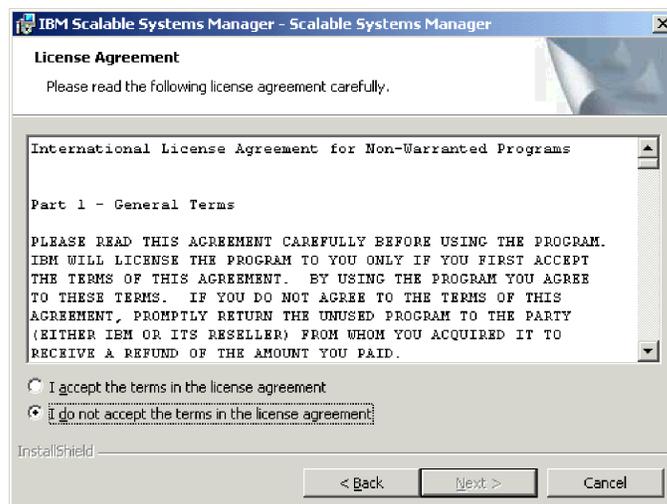
Figure 5. Installing Scalable Systems Manager: "License Agreement" window

11. Read the license agreement; then, click **I accept the terms in the license agreement** and click **Next**. The "Destination Folder" window opens.
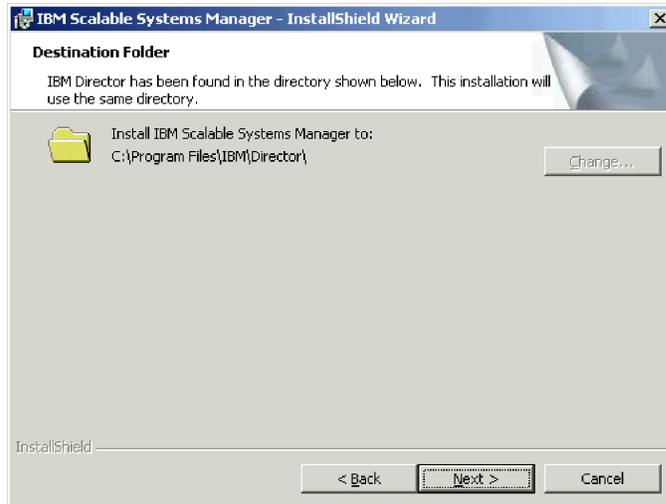
Figure 6. Installing Scalable Systems Manager: "Destination Folder" window

12. Click **Next**. The "Ready to Install the Program" window opens.
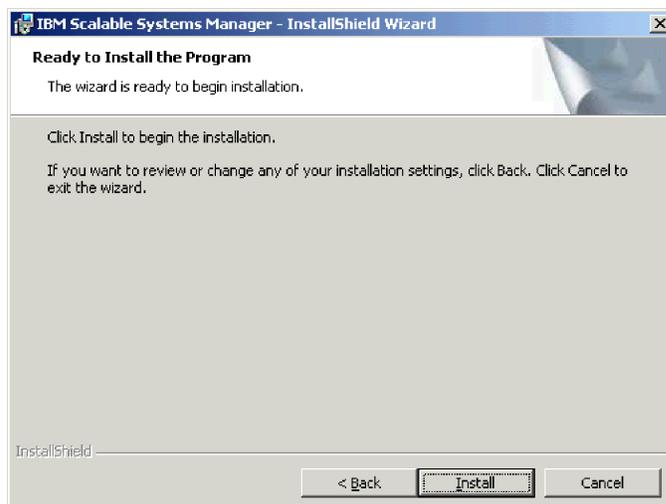
Figure 7. Installing Scalable Systems Manager: "Ready to Install the Program" window

13. Click **Install**. The "Installing IBM Scalable Systems Manager" window opens. The status bar displays the progress of the installation.

Figure 8. Installing Scalable Systems Manager: "Installing" window

When the installation completes, the "InstallShield Wizard Completed" window opens.
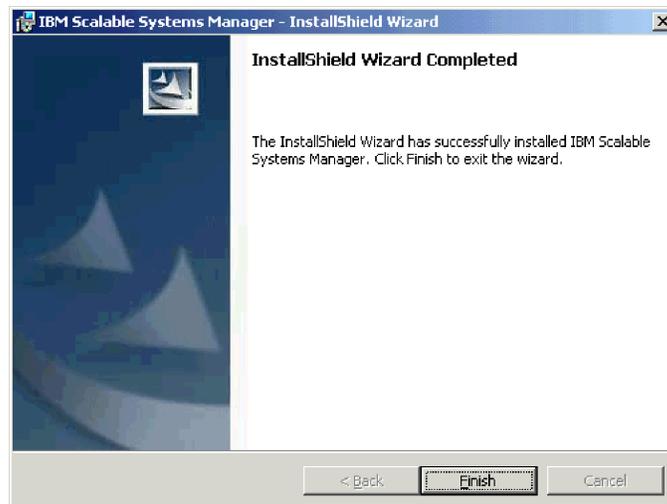
Figure 9. Installing Scalable Systems Manager: "InstallShield Wizard Completed" window

14. Click **Finish**.

Uninstalling Scalable Systems Manager

You can uninstall Scalable Systems Manager by itself, or if you uninstall IBM Director while Scalable Systems Manager is still installed, the IBM Director setup program will uninstall Scalable Systems Manager silently.

Complete the following steps to uninstall Scalable Systems Manager:

1. If necessary, close IBM Director Console.
2. Click **Start** → **Settings** → **Control Panel**. The "Control Panel" window opens.
3. Double-click **Add/Remove Programs**. The "Add/Remove Programs" window opens.
4. Click **IBM Scalable Systems Manager**.

5. Click **Remove**.
6. When prompted `Are you sure you want to remove IBM Scalable Systems Manager from your computer`, click **Yes**.

Chapter 3. Scalable objects

This chapter describes IBM Director physical-platform discovery, the types of scalable objects, the supported configurations for scalable systems, and how to save and restore scalable information.

Discovering physical platforms

A *physical platform* is an IBM Director managed object that represents a system that can support an operating system.

To discover physical platforms, from IBM Director Console, click **Tasks** → **Discover Systems** → **Physical Platforms**. IBM Director uses out-of-band communication to discover physical platforms. Using service location protocol (SLP), IBM Director queries the service processor in each server on the network. Depending on the size of your network, this discovery can take several minutes to complete.

When IBM Director discovers a physical platform, it creates a physical-platform managed object to represent the system in the Group Contents pane of IBM Director Console. The physical-platform managed object is identified by a ▯ icon and the name is a concatenation of the machine type, machine model, and machine serial number, for example, IBM 88702RX 23A0615.

When IBM Director initially creates a physical platform, it is generally locked, and you must request access to it before you can use it with Scalable Systems Manager. For more information, see "Unlocking physical platforms" on page 20.

When you add xSeries 445 servers to an IBM Director environment, you can use Scalable Systems Manager to create physical-platform managed objects that represent these servers without initiating physical-platform discovery for all servers on the network. For details, see "Manually creating unassigned scalable nodes" on page 49.

Types of scalable objects

Scalable objects are the managed objects that you can use with Scalable Systems Manager. The scalable objects are as follows:
- Scalable nodes, including unassigned scalable nodes
- Scalable systems
- Scalable partitions
- Remote I/O enclosures that are attached to scalable nodes

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.

Important: Scalable Systems Manager cannot communicate with the service processors on the servers that are represented as scalable nodes in IBM Director unless the following conditions are met:

- Service-processor networks must either have static IP address assignments or have DHCP configured to maintain consistent IP addresses for the service processors. To do so, create reservations in DHCP that identify explicit IP addresses for the service processors.
- The IP addresses that are assigned to the service processors do not change after the servers are discovered in IBM Director. This limitation is true whether you are using DHCP or statically assigning IP addresses. Chapter 9, "Solving Scalable Systems Manager problems", on page 87 describes suggested actions to take if the IP address of a service processor has changed since it was discovered by IBM Director.

In Scalable Systems Manager 4.11, the xSeries 440 and xSeries 445 are the only IBM servers that are defined as scalable nodes. Support is as follows:

- Scalable nodes that represent xSeries 445 servers in supported configurations are used in manageable scalable systems. For details about which configurations are supported, see "Manageable scalable system configurations" on page 25.
- Scalable nodes that represent xSeries 440 servers in a 16-way configuration are used in view-only scalable systems.

Note: You can create scalable partitions only from scalable nodes that are used in manageable scalable systems.

Scalable Systems Manager uses the universal unique identifier (UUID) to identify the scalable node. The UUID is not the same as the host name; the host name is an attribute of the service processor, and the UUID is an attribute of the server system board. You can remove and replace the service processor without affecting the UUID (but the host name is affected).

In IBM Director Console, scalable nodes are identified with the ▯ icon, which is the same icon that is used for all physical platforms. To determine whether a physical platform has the additional attributes that are to be considered 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 the server that is represented by the physical-platform managed object are displayed. The attributes for SMP Expansion Modules and RXE Expansion Ports are at the end of the list.

Note: Scalable nodes cannot be defined in a scalable partition until they are first part of a scalable system.

Unlocking physical platforms

When IBM Director first creates a physical-platform managed object to represent a server, the 🔒 icon is usually displayed beside the physical-platform icon. The lock means that access through the service processor has not been granted. However, if the service processor has the default login ID and password, IBM Director automatically unlocks the service processor and does not display the padlock icon.

Locked physical platforms that are scalable nodes can be added to scalable systems and scalable partitions, and can be viewed in Scalable Systems Configuration. However, you cannot start scalable partitions that contain a locked scalable node. If the login ID and password have been changed from the default values, you must request access to the service processor.

To request access, from the Group Contents pane, right-click the physical-platform icon and click **Request Access**. By providing a valid login ID and password that has read/write access to the service processor, you can unlock the physical platform. To create a login ID, see the documentation provided with the service processor.

After the service processor has been unlocked, IBM Director can communicate with it out-of-band. After a physical platform that represents a scalable node is unlocked, any remote I/O enclosures that are attached to the server are displayed in IBM Director.

Interrogation for partition descriptors

When a scalable node is unlocked, IBM Director performs additional discovery for the server that it represents. This discovery determines whether the NVRAM of the service processor contains a partition descriptor that was stored by the Configuration/Setup Utility program (in the BIOS code). If it does, IBM Director uses the partition-descriptor information to create a scalable system and a scalable partition. This process can create manageable or view-only scalable systems.

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

The primary server becomes the primary scalable node of the scalable partition. (The primary server is the server from which the Configuration/Setup Utility program was run to create the partition descriptor.)

The interrogation of NVRAM to locate a BIOS-created partition descriptor is performed by IBM Director, regardless of whether Scalable Systems Manager is installed.

Notes:

1. Scalable systems that are created for 16-way xSeries 440 servers are view-only scalable systems and are represented by the ⊞ icon. View-only scalable systems have limited functionality. For more information about view-only scalable systems, see "Scalable systems" on page 23.

2. Interrogation of NVRAM to locate a BIOS-created partition descriptor is the only process that can result in the creation of a scalable system and a scalable partition that contain scalable nodes to represent a 16-way xSeries 440 server.

This automatic creation of a scalable system and scalable partition enable you to use IBM Director to view and manage scalable objects without first using the Scalable Systems Configuration task. For example, you immediately can use Management Processor Assistant (MPA) and Scalable Systems Configuration to power-on a scalable partition and start an operating system on it.

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.

Viewing information about scalable nodes

You can use IBM Director and Scalable Systems Configuration to view information about all scalable nodes. This includes those in manageable scalable systems, those in view-only scalable systems, and those that are unassigned. You can view information as follows:

- Display scalable objects in IBM Director Console
- Display details about scalable nodes in Scalable Systems Configuration
- Refresh information displayed about scalable nodes

Managing scalable nodes

Scalable Systems Manager can manage only scalable nodes that represent xSeries 445 servers.

Note: You cannot manage scalable nodes in view-only scalable systems (those that are represented by the ⊞ icon). View-only scalable systems have limited functionality.

To manage a scalable node in a scalable partition, you must either discover a scalable system that uses the scalable node or create a scalable system. Then, create a scalable partition and insert the scalable nodes from the scalable system into the scalable partition. After you set the primary scalable node for the scalable partition, you can power it on.

You can use Scalable Systems Configuration to perform the following tasks:

- For scalable nodes in a scalable system that are not yet in a scalable partition, you can manage them as follows:
  - Update the scalable node selection for an existing scalable system
  - Insert a scalable node into a scalable partition
- You can perform these additional tasks on scalable nodes in a scalable partition:
  - Set the primary scalable node in a scalable partition
  - Remove a scalable node from a scalable partition
- You can create and name scalable nodes as follows:
  - Create a new unassigned scalable node
  - Rename a scalable node

Unassigned scalable nodes

Unassigned scalable nodes are physical-platform managed objects that represent xSeries 445 servers that have been discovered by IBM Director but are not part of a manageable scalable system. Because IBM Director cannot determine which scalable nodes are interconnected through their SMP Expansion Modules and so form a valid scalable system, you must either use the Scalable Systems Discovery task to automatically create manageable scalable systems from unassigned scalable nodes, or you must use Scalable Systems Configuration to create or update manageable scalable systems.

The unassigned scalable nodes are displayed in a tree hierarchy in the Unassigned pane of the "Scalable Systems Configuration" window.

The actual contents that are shown in the Unassigned pane depends on how you started the Scalable Systems Configuration task:

- If you double-click the **Scalable Systems Configuration** task to start it, the Unassigned pane shows all discovered scalable nodes that are not yet part of a scalable system.
- If you select scalable systems from the IBM Director Group Contents pane and start Scalable Systems Configuration just for those scalable systems, no unassigned scalable nodes are displayed.
- If you select unassigned scalable nodes from the IBM Director Group Contents pane and start Scalable Systems Configuration, only the selected unassigned scalable nodes are shown.

You can perform the following operations on unassigned scalable nodes:
- Discover scalable systems that use unassigned scalable nodes
- Create a scalable system from unassigned scalable nodes
- Update the scalable node selection for an existing scalable system
- Delete an unassigned scalable node

Scalable systems

A scalable system is an IBM Director managed object that consists of scalable nodes and the scalable partitions that are made from the scalable nodes in the scalable system. When a scalable system contains two scalable nodes, the servers that they represent must be interconnected through their SMP Expansion Modules to make a 16-way configuration, for example, a 16-way xSeries 445 server.

Differences between manageable and view-only scalable systems

Table 5 describes manageable and view-only scalable systems and the icon that is displayed beside each type of scalable system.

Table 5. Descriptions of manageable and view-only scalable systems

Icon	Description
	A manageable scalable system. Scalable Systems Manager supports several different server configurations for use in manageable scalable systems. They can be managed with all the features in Scalable Systems Manager. For more information about supported configurations, see a "Manageable scalable system configurations" on page 25. A single scalable node in a manageable scalable system is the smallest manageable unit of physical partitioning. As such, the server that it represents is the smallest physical configuration that can be partitioned with Scalable Systems Manager to run a single instance of an operating system.

Table 5. Descriptions of manageable and view-only scalable systems (continued)

Icon	Description
	A view-only scalable system. Scalable Systems Manager supports only 16-way configurations of xSeries 440 servers as view-only scalable systems. These scalable systems are created only from partition descriptors saved to NVRAM by the Configuration/Setup Utility program (in the BIOS code). You can use Scalable Systems Configuration to display detailed information about view-only scalable systems. The limitations of view-only scalable systems are as follows: • You cannot discover view-only scalable systems with Scalable Systems Discovery. • You cannot create view-only scalable systems in Scalable Systems Configuration. • You cannot update the scalable node selection for view-only scalable systems. • You cannot use Scalable Systems Validation with view-only scalable systems. • You cannot use the RXE Configuration wizard to configure a remote I/O enclosure that is attached to an xSeries 440 server in a view-only scalable system. • You cannot perform power operations on the scalable partition that is defined for a view-only scalable system.

Creation of scalable systems

IBM Director creates scalable systems in the following ways:

• Automatically interrogating the service processor NVRAM for partition descriptors

When a scalable node is unlocked, IBM Director automatically interrogates the NVRAM of the service processor to determine whether it contains a partition descriptor that was stored by the Configuration/Setup Utility program. If it does, IBM Director uses the partition descriptor information to create a scalable system and a scalable partition that contain the scalable nodes that are defined in the partition descriptor. This process can create manageable or view-only scalable systems.

• Running the Scalable Systems Discovery task

You can run or schedule the Scalable Systems Discovery task to create manageable scalable systems from unassigned scalable nodes that represent servers that are cabled together. This task also automatically updates manageable scalable systems with unassigned scalable nodes that are cabled to the scalable nodes that are already defined in the scalable system. The Scalable Systems Discovery task does not create or update view-only scalable systems.

• Using the Scalable Systems Configuration task

You can use the Scalable Systems Configuration task to create a manageable scalable system. When you create a scalable system using the Scalable Systems Configuration task, you use a graphical user interface to define the name of the scalable system and select the unassigned scalable nodes that are to be assigned to the scalable system. You cannot create view-only scalable systems.

Viewing information about scalable systems

You can use IBM Director and Scalable Systems Configuration to view information about manageable and view-only scalable systems. You can view information as follows:

• Display scalable objects in IBM Director Console

- Display details about scalable systems in Scalable Systems Configuration
- Refresh information about scalable systems

Managing scalable systems

Scalable Systems Manager can manage only scalable systems that consist of scalable nodes that represent xSeries 445 servers.

Note: You cannot manage view-only scalable systems (those that are represented by the ⊞ icon). View-only scalable systems have limited functionality.

You can perform the following tasks on manageable scalable systems:

- Automatically discover or complete a scalable system
- Create a scalable system
- Rename a scalable system
- Update the scalable-node selection for an existing scalable system
- Validate the cabling of a scalable system against supported configurations
- Create a scalable partition for a scalable system
- Delete a scalable partition from a scalable system
- Check for problems with scalable systems
- Delete a scalable system

Manageable scalable system configurations

In Scalable Systems Manager 4.11, xSeries 445 servers are the only server models that can be used in manageable scalable systems. Scalable Systems Manager 4.11 supports several server configurations for use in manageable scalable systems.

When you are creating or updating a manageable scalable system, Scalable Systems Configuration displays configuration graphics to depict the scalable nodes, or server chassis, that you have selected for the manageable scalable system.

Server configurations used in manageable scalable systems

A server that is represented as a scalable node in a manageable scalable system can contain either one or two SMP Expansion Modules. Each SMP Expansion Module has as many as four microprocessors, its own system cache, its own memory, and dedicated high-speed ports that are used to interconnect the modules. The following configurations of servers can be used in Scalable Systems Manager to form manageable scalable systems:

4-way One server with one SMP Expansion Module that has up to four microprocessors

8-way One server with two SMP Expansion Modules that each have four microprocessors

16-way
One server with two SMP Expansion Modules that each have four microprocessors
Two 8-way servers that are cabled together through their SMP Expansion Modules

Information provided in the configuration graphics

The configuration graphics show the number of SMP Expansion Modules in each server and the number of cables that connect the two scalable nodes. The graphics also depict whether and how a remote I/O enclosure is attached to the scalable nodes in the scalable system.

The following table describes the notation that is used in the configuration graphics to display information about each server and enclosure.

Table 6. Notations used in graphics for manageable scalable systems

Depicted in graphic	Notation used
Scalable nodes that are defined in the scalable partition	Blue rectangles
Remote I/O enclosures that are attached to the scalable nodes in the scalable system	Green rectangles
RXE Management cables and their connection to RXE Management Ports	Purple solid lines
RXE Expansion cables and their connection to RXE Expansion Ports	Orange solid lines
Optional RXE Expansion cables and their connection to RXE Expansion Ports	Orange dashed lines
SMP Expansion cables and their connection to SMP Expansion Ports	Red solid lines

For information about cabling, see the hardware documentation that comes with your servers.

Supported 4-way configurations

A 4-way configuration can have one scalable partition.

Table 7. Supported scalable systems for 4-way servers

Scalable system configuration	Description
	One 4-way xSeries 445 server. This server has one SMP Expansion Module that has up to four microprocessors.
	One 4-way xSeries 445 server that is attached to expansion kit A of a remote I/O enclosure. **Note:** At the time of publication, Scalable Systems Manager does not support a scalable partition created from this scalable system configuration that enables both expansion kits of the remote I/O enclosure. For more information, see "Remote I/O enclosures" on page 30.
	One 4-way xSeries 445 server that is attached to expansion kit B of a remote I/O enclosure. **Note:** At the time of publication, Scalable Systems Manager does not support a scalable partition created from this scalable system configuration that enables both expansion kits of the remote I/O enclosure. For more information, see "Remote I/O enclosures" on page 30.

Supported 8-way configurations

An 8-way configuration can have one scalable partition.

Table 8. Supported scalable systems for 8-way servers

Scalable system configuration	Description
	One 8-way xSeries 445 server. This server has two SMP Expansion Modules that each have four microprocessors. The SMP Expansion Modules are connected to each other through SMP Expansion Ports on each SMP Expansion Module.
	One 8-way xSeries 445 server that is attached only to expansion kit B of a remote I/O enclosure. **Note:** At the time of publication, Scalable Systems Manager does not support a scalable partition created from this scalable system configuration that enables both expansion kits of the remote I/O enclosure. For more information, see "Remote I/O enclosures" on page 30.
	One 8-way xSeries 445 server that is attached to both expansion kit A and B of a remote I/O enclosure. **Note:** At the time of publication, Scalable Systems Manager does not support a scalable partition created from this scalable system configuration that enables both expansion kits of the remote I/O enclosure. For more information, see "Remote I/O enclosures" on page 30.
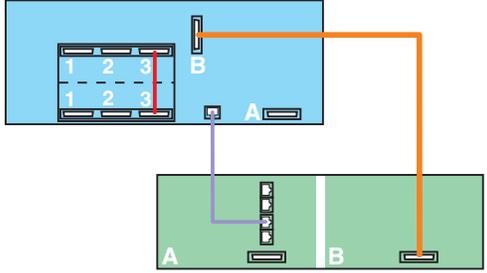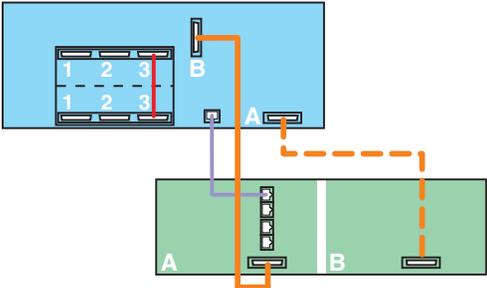
Supported 16-way configurations

A 16-way configuration can have three scalable partitions: one scalable partition per server and one scalable partition that consists of both servers.

Table 9. Supported scalable systems for 16-way servers
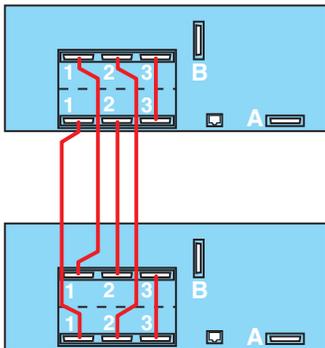
Scalable system configuration	Description
	Two 8-way xSeries 445 servers that are connected together through three SMP Expansion ports on each SMP Expansion Module. This configuration is called a 16-way configuration.

Table 9. Supported scalable systems for 16-way servers (continued)

Scalable system configuration	Description
	A 16-way configuration of xSeries 445 servers that is attached to both expansion kit A and B of a remote I/O enclosure.

Scalable partitions

A scalable partition is an IBM Director managed object that is defined for a scalable system. It contains one or two scalable nodes from that scalable system and can run a single image of an operating system. A scalable partition also includes any remote I/O enclosures that are attached to the scalable nodes that are contained in that scalable partition.

Note: You cannot create scalable partitions for unassigned scalable nodes until they have been defined into a manageable scalable system.

Scalable partitions can:
- Be powered-on and powered-off individually
- 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

Note: After creating a scalable partition, you must designate one scalable node as the primary scalable node before you can start an operating system on that scalable partition.

When you power-on a scalable partition, Scalable Systems Manager communicates out-of-band with the service processor on the scalable nodes in the scalable partition, for example, to physically power-on those scalable nodes and then start the operating system on the primary scalable node.

Note: The first time you start a scalable partition that you created, you must use Scalable Systems Manager to start it. Do not use the power-control button of a server in the scalable partition the first time the scalable partition is started.

After the scalable partition has been started through Scalable Systems Manager, you can use the power-control button of the server or Scalable

Systems Manager to perform power operations. Generally, the power-control button of a server will start the last scalable partition that was started.

IBM Director Console identifies all scalable partitions with the ⬜ icon whether they are powered-on or powered-off. IBM Director Console uses additional icons with this scalable partition icon to indicate the state of a scalable partition. For more information, see "States of scalable partitions" on page 83.

Scalable Systems Configuration identifies a scalable partition that is powered-on (active) with the ⬜ icon and a scalable partition that is powered-off (inactive) with the ⬜ icon. In the "Scalable Systems Configuration" window, a scalable partition is displayed in the Topology pane under the name of its parent scalable system. The scalable nodes that the scalable partition contains are listed as children of the scalable partition.

Determining the primary scalable node for a scalable partition

After you create a scalable partition and before you power-on (start an operating system on) the scalable partition, you must designate one scalable node as the primary scalable node. Scalable Systems Configuration does not set the primary scalable node automatically, even for a scalable partition with only one scalable node. (Scalable partitions that are created from interrogation of partition descriptors automatically have a designated primary scalable node.)

IBM Director uses the primary scalable node to power-on the scalable partition. The primary scalable node controls the resources that are assigned to the two servers and retains the error logs. CD-ROM, DVD-ROM, and diskette drives are active on the designated primary scalable node only.

For example, if a scalable system contains scalable node A and scalable node B, you can define the following scalable partitions. These configurations assume that an operating system is installed on each scalable node. Although you can simultaneously create all four scalable partitions for the scalable system, you can power-on only those scalable partitions that do not contain scalable nodes that are already powered-on in a different scalable partition.

Table 10. Scalable partitions for a 16-way configuration

Scalable partition name	Scalable nodes in the scalable partition	Primary scalable node for the scalable partition
Partition1	Scalable node A and scalable node B	Scalable node A
Partition2	Scalable node A and scalable node B	Scalable node B
Partition3	Scalable node A	Scalable node A
Partition4	Scalable node B	Scalable node B

You must use Scalable Systems Configuration to set the primary node for a scalable partition.

Updating partition descriptors in NVRAM

When a scalable partition is starting, Scalable Systems Manager writes the definition of that scalable partition to the service processor of the server corresponding to the primary scalable node. Then, the service-processor firmware

writes the definition to the NVRAM of all scalable nodes in the scalable partition. This definition is called a *partition descriptor*.

This partition descriptor must be successfully written to NVRAM before the scalable partition actually starts.

Creation of managed systems for scalable partitions
When you power-on a scalable partition, the scalable nodes are physically powered on, and the operating system is started on the scalable partition.

If IBM Director Agent is installed on a powered-on scalable partition, it is also started. When IBM Director 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. IBM Director Server communicates in-band with this managed system, which means that it can be managed in IBM Director as any other managed system can.

Viewing information about scalable partitions
You can use IBM Director and Scalable Systems Configuration to view information about scalable partitions in both manageable and view-only scalable systems. You can view information as follows:

- Display scalable objects in IBM Director Console
- Display details about scalable partitions in Scalable Systems Configuration
- Refresh information displayed about scalable partitions.

Managing scalable partitions
Scalable Systems Manager can manage only scalable partitions that consist of scalable nodes that represent xSeries 445 servers.

Note: You cannot manage scalable partitions in view-only scalable systems (those that are represented by the ⊞ icon). View-only scalable systems have limited functionality.

To manage an existing scalable partition, you must run the Scalable Systems Configuration task on the scalable system that contains that scalable partition.

You can perform the following tasks for scalable partitions:

- Create a scalable partition in a scalable system
- Rename a scalable partition
- Delete a scalable partition
- Insert scalable nodes into a scalable partition
- Set the primary scalable node for a scalable partition
- Configure a remote I/O enclosure for a scalable partition
- Check for problems with scalable partitions
- Perform power operations on a scalable partition
- Remove a scalable node from a scalable partition

Remote I/O enclosures

Remote I/O enclosures, for example, RXE-100 Remote Expansion Enclosures, are represented as IBM Director managed objects. When a physical-platform managed object is unlocked, IBM Director discovers any remote I/O enclosures that are attached to the server that it represents and creates managed objects for the

remote I/O enclosures. IBM Director uses the UUIDs to identify the remote I/O enclosures. By default, IBM Director configures both expansion kits of an enclosure for use.

To discover an attached remote I/O enclosure, IBM Director uses the RXE Management cables connected to a server as IBM Director cannot detect the RXE Expansion cabling. IBM Director assumes the RXE Expansion cables are connected to the server in the same manner as are the RXE Management cables. For example, if an RXE Management cable is attached to port A of the server, IBM Director assumes that an RXE Expansion cable is attached to port A of the server and to expansion kit A of the remote I/O enclosure. This determination affects the remote I/O configuration choices that you can select in the RXE Configuration wizard. For more information, see "Configuring a remote I/O enclosure" on page 55.

At the time of publication, Scalable Systems Manager does not support management of a scalable partition that consists of a stand-alone server connected to a remote I/O enclosure that has both expansion kits enabled. Such configurations are discovered by IBM Director, and will not display as unsupported, but you cannot use Scalable Systems Manager to manage these configurations as discovered. After discovery, you must use the RXE Configuration wizard to disable at least one of the expansion kits. For information about the RXE Configuration wizard, see "Configuring a remote I/O enclosure" on page 55.

Note: To configure a scalable partition consisting of two stand-alone servers sharing one remote I/O enclosure, you must create the scalable partition through the Configuration/Setup Utility program (in the BIOS code) on the server. This configuration is unsupported by Scalable Systems Manager; it can be managed only with IBM Director.

Each remote I/O enclosure can contain one or two PCI-X expansion kits, which each contain six slots for adapters. When IBM Director discovers a remote I/O enclosure that is attached to two servers in a 16-way configuration, IBM Director Console shows PCI-X expansion kit A as being attached to one scalable node and shows PCI-X expansion kit B as being attached to the other scalable node.

Remote I/O enclosures can be attached in the following ways:
* Both PCI-X expansion kit A and B can be attached to the same scalable node, as shown by the ▦ icon.
* PCI-X expansion kit A can be attached to one scalable node, as shown by the ▦ icon.
* PCI-X expansion kit B can be attached to a different scalable node, as shown by the ▦ icon.

When a scalable node is part of a scalable partition, you can use the RXE Configuration wizard to configure a remote I/O enclosure that is attached to that scalable node.

A PCI-X expansion kit in a remote I/O enclosure that is attached to a supported server can be used only by a scalable partition that contains that scalable node.

Information about remote I/O enclosures that are attached to scalable nodes is displayed in the ″Scalable Systems Configuration″ window.

Saving and restoring scalable-object information in IBM Director Server

When you use Scalable Systems Manager to create scalable systems and scalable partitions, IBM Director stores this scalable-object information in persistent storage.

If you perform a twgreset command on IBM Director Server, the IBM Director persistent storage, including the scalable-object information, is removed. This information can be preserved during an upgrade or uninstallation/reinstallation, but is removed if you answer **Yes** to the prompt ″Do you wish to delete configuration data?.″ Although physical platforms can be rediscovered, all the information about scalable systems and scalable partitions, including the configuration of remote I/O enclosures is lost.

If you want to preserve scalable-system and scalable-partition information, you must make a backup copy of IBM Director persistent storage before you issue the twgreset command or remove configuration data during an uninstall/reinstall of IBM Director Server.

Saving IBM Director Server persistent storage

Complete the following steps to back up the IBM information that is saved in IBM Director Server:

1. To stop IBM Director Server, from a command prompt, type the following command and press Enter:

   d:\net stop twgipc

   where *d* is the drive letter of the hard disk drive where IBM Director is installed.

2. To save IBM Director persistent storage, type the following command and press Enter:

   d:\twgsave

   where *d* is the drive letter of the hard disk drive where IBM Director is installed.

   Issuing this command copies all the files in the IBM Director data directory to a backup directory. The directory is named *director*.save.*number*, where *director* is the name of the directory where IBM Director is installed and *number* is the number of the backup.The directory is in the directory where IBM Director Server is installed.

Restoring IBM Director Server persistent storage

After IBM Director persistent storage has been reset with a twgreset command or configuration data has been removed during an uninstallation, you can restore the previously backed-up IBM Director persistent storage. This will restore any Scalable Systems Manager scalable systems and scalable partitions that were previously saved.

Complete the following steps to restore IBM Director persistent storage:

1. From a command prompt, type the following command and press Enter:

   d:\twgrestore "*path**directory_name*"

   where
   - *d* is the drive letter of the hard disk drive where IBM Director is installed
   - *path* is the path of the backup directory, for example, Program Files\IBM.

- *directory_name* is the name of the backup directory

This command copies all the files from the backup directory to the IBM Director data directory and asks whether you want to overwrite each file.

2. To start IBM Director Server, type the following command and press Enter:

   `d:\net start twgipc`

   where *d* is the drive letter of the hard disk drive where IBM Director is installed.

Chapter 4. Using the Scalable Systems Discovery task

You can manually invoke the Scalable Systems Discovery task to:

- Automatically create a manageable scalable system from unassigned scalable nodes that represent servers that are cabled together
- Automatically update a manageable scalable system to add unassigned scalable nodes that are cabled to the scalable nodes already defined in that manageable scalable system

Note: Scalable Systems Discovery does not create or update view-only scalable systems (those that are represented by the ⊞ icon). For example, Scalable Systems Discovery does not create or update scalable systems that are defined with scalable nodes that represent xSeries 440 servers.

The Scalable Systems Discovery task is a non-interactive task that you can schedule or run immediately. For more information on scheduling jobs, see the *IBM Director 4.11 System Management Guide*.

After using IBM Director to discover all physical platforms, you can run the Scalable Systems Discovery task for any unassigned scalable nodes. It is also useful to run this task for a scalable system that is created from an NVRAM partition descriptor when the number of expected scalable nodes defined by the partition descriptor is more than the actual number defined in the scalable system. The Scalable Systems Discovery task must target one or more unassigned scalable nodes, or one or more scalable systems.

Operation of Scalable Systems Discovery

The Scalable Systems Discovery task determines which unassigned scalable nodes belong in a scalable system by communicating with the service processor on each targeted scalable node.

The Scalable Systems Discovery task performs the following operations depending on the scalable object that is targeted:

- Creates one scalable system for unassigned scalable nodes that are cabled together through their SMP Expansion Ports.
- Creates one scalable system for each unassigned scalable node that is not cabled to any other scalable node.
- Updates a scalable system to add unassigned scalable nodes that were not available when the scalable system was created. (These unassigned scalable nodes are cabled to scalable nodes in the scalable system).
- For a scalable system that was created from an NVRAM partition descriptor, determines whether the expected number of scalable nodes that are defined in the partition descriptor matches the actual number of scalable nodes defined in the scalable system. The Scalable Systems Discovery task makes this determination by comparing the value of the "Expected scalable nodes" property for the scalable system to the actual cabling configuration that it determines.
  - If these two numbers of scalable nodes match, the Scalable Systems Discovery task performs no further action as it does not need to update the scalable-system definition.
  - If the two numbers of scalable nodes do not match, the Scalable Systems Discovery task issues ping commands to the scalability ports on the scalable

nodes to determine the current cabling configuration. Scalable Systems Discovery task updates the scalable-system definition according to its findings.

Depending on the number of scalable systems and scalable nodes that are involved, and the complexity of the cabling for each scalable node, the discovery process can take several minutes to complete.

Note: For performance reasons, the Scalable Systems Discovery task issues ping commands only to the minimum number of scalability ports when determining how to create or update a scalable system. To issue ping commands to every scalability port on every scalable node in a scalable system, use the Scalable Systems Validation task described in Chapter 5, "Using the Scalable Systems Validation task", on page 39.

The Scalable Systems Discovery task does not discover remote I/O enclosures. Rather, IBM Director discovers remote I/O enclosures that are attached to a server when their associated physical platform is automatically or manually unlocked.

The Scalable Systems Discovery task does not create scalable partitions for the scalable systems that it creates. Instead, use Scalable Systems Configuration to create scalable partitions from the scalable nodes in a scalable system.

Starting Scalable Systems Discovery

The Scalable Systems Discovery task is under the **Scalable Systems Manager** task in the Tasks pane of IBM Director Console. You can start the **Scalable Systems Discovery** task from IBM Director Console or from Scalable Systems Configuration. You can run this task immediately or as a scheduled job.

Important: Before using Scalable Systems Discovery, you must use the Management Processor Configuration subtask of MPA to configure an alert-forwarding profile for the servers that are represented by the scalable nodes that you plan to target. In the alert-forwarding profile, the connection type must be set to "IBM Director Comprehensive" and the IP address must be set to the IP address of the management server that is being used to manage scalable objects with Scalable Systems Manager.

To start the Scalable Systems Discovery task in IBM Director Console, drag it onto one or more of the following targets:
- Scalable system in the Group Contents pane
- Scalable node in the Group Contents pane
- A group in the Groups pane that contains scalable systems or scalable nodes

Note: The targeted scalable nodes must be unlocked. Otherwise, the discovery is skipped for that scalable node and any scalable systems that contain that scalable node. For information about unlocking scalable nodes, see "Unlocking physical platforms" on page 20.

You can start this task for more than one scalable system or scalable node by selecting multiple scalable systems or scalable nodes in the Group Contents pane before you drag the Scalable Systems Discovery task onto one of the selected scalable objects. If you run multiple Scalable Systems Discovery tasks at the same time, IBM Director queues the tasks so that only one is running at a time.

You can also start Scalable Systems Discovery from Scalable Systems Configuration by right-clicking the name of a scalable system or scalable node; then, clicking **Scalable system discovery**.

After you start the task, Scalable Systems Manager displays a window asking whether you want to schedule a job for the Scalable Systems Discovery task or run it immediately. The window lists the scalable systems and scalable nodes on which you want to run the discovery task, either in the window title bar when only one scalable system or scalable node is selected, or in a list in the window when multiple scalable systems or scalable nodes are selected.

- Click **Schedule** to create an IBM Director scheduled job. For more information about scheduled jobs, click **Help** in the New Scheduled Job window.
- Click **Execute Now** to start the Scalable Systems Discovery task immediately.

In addition to invoking the Scalable Systems Discovery task by dragging it in IBM Director Console, you can also select it from the Scheduler when you create a new scheduled job. For information about the Scheduler, see the help in IBM Director Console.

Viewing the results from Scalable Systems Discovery

When the discovery task is started immediately, a job window opens providing details about the task execution. This window provides the status of the currently running discovery job and history information about discovery jobs that were previously run on that management server.
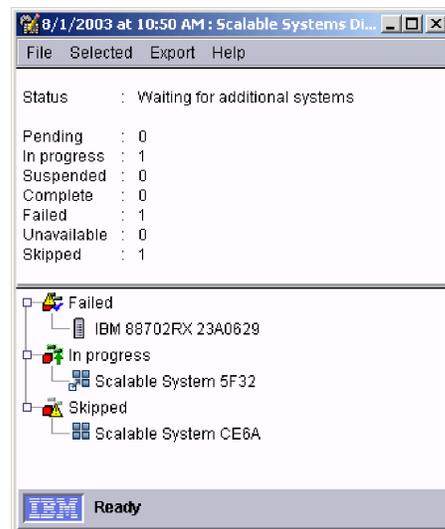
Figure 10. Viewing summary results from Scalable Systems Discovery

The lower portion of the window lists the scalable systems and scalable nodes that were used during the discovery task. Expand the tree structure for **Successful**, **In progress**, **Skipped**, or **Failed** to see the list of scalable systems and scalable nodes under each category.

A scalable node is skipped when the physical-platform managed object does not represent a server that is supported by Scalable Systems Discovery, for example,

an xSeries 440 server. A scalable node or scalable system is also skipped if the scalable node is already part of a scalable system or if the scalable system is already complete.

The Scalable Systems Discovery process is considered failed for a scalable node or scalable system when any of the ping commands to the scalability ports on the selected scalable nodes have failed.

For discovery-task details about a specific scalable system or scalable node, double-click the name of the scalable system or scalable node. A window opens that lists the start and completion times for the discovery task on that scalable object.
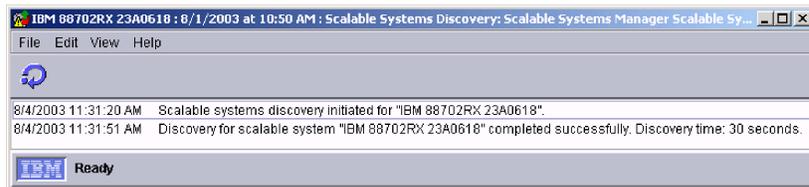
Figure 11. Viewing detailed results from Scalable Systems Discovery

Use the **View** → **Detail** menu function on the window to display low, medium, or high levels of detail. By default, a low level of detail is displayed. Use the **File** → **Clear All Messages** menu function to remove all messages from the discovery log file. This option removes all messages from the log, not just those for the scalable object that you were viewing.

Important: The Scalable Systems Configuration window might list the automatically-created scalable systems before the discovery task completes. You should not manage them until the discovery task completes.

Using the resulting scalable systems

You can use the Scalable Systems Configuration task to create scalable partitions for the scalable systems that are created by the Scalable Systems Discovery task.

You can also use Scalable Systems Configuration to update the scalable node selection for a scalable system. For example, the Scalable Systems Discovery task could create a scalable system that contains one scalable node. Later, you might cable a second server to the first server and create a new scalable node for that server. In this case, you would need to update the scalable system so that it contained the new scalable node.

For more information, see "Creating scalable partitions" on page 52 and "Updating a scalable system" on page 66.

Chapter 5. Using the Scalable Systems Validation task

Use the Scalable Systems Validation task to validate the cabling of all scalable nodes in one or more manageable scalable systems. The Scalable Systems Validation task is a non-interactive task that you can schedule or run immediately. For more information on scheduling jobs, see the *IBM Director 4.11 System Management Guide*.

When the validation task is started for a scalable system, Scalable Systems Manager communicates with the service processor of each scalable node in that scalable system. Scalable Systems Manager issues a ping command to each scalability port on each scalable node and waits for the responses. Scalable Systems Manager determines if the resulting responses correspond to a supported configuration of a manageable scalable system.

Note: Scalable Systems Validation does not validate the cabling for view-only scalable systems (those that are represented by the ⊞ icon). For example, Scalable Systems Validation cannot target scalable systems that are defined with scalable nodes that represent xSeries 440 servers.

This task is useful for running as a diagnostic tool, for example, when a scalable partition does not seem to start correctly, and as a preventive maintenance tool, for example, to schedule periodic validations of scalable systems to verify they are cabled correctly.

The Scalable Systems Validation task is also useful for verifying the cable configuration of automatically discovered scalable systems. This is applicable when the Scalable Systems Discovery task discovers a scalable system with more than one scalable node because that task does not query all ports on all scalable nodes when it creates a scalable system. For more information, see Chapter 4, "Using the Scalable Systems Discovery task", on page 35.

Depending on the number of scalable systems that are involved and the complexity of each scalable node in those scalable systems, the validation process can take several minutes to complete. Scalable Systems Manager waits five seconds for a response from each ping command. If there is no response in that time frame, the validation process continues with the next scalability port.

Starting Scalable Systems Validation

The Scalable Systems Validation task is under the **Scalable Systems Manager** task in the Tasks pane of IBM Director Console. You can start the Scalable Systems Validation task from IBM Director Console or from Scalable Systems Configuration. You can run this task immediately or as a scheduled job.

Important: Before using Scalable Systems Validation, you must use the Management Processor Configuration subtask of MPA to configure an alert-forwarding profile for the servers that are represented by the scalable nodes that you plan to target. In the alert-forwarding profile, the connection type must be set to "IBM Director Comprehensive" and the IP address must be set to the IP address of the management server that is being used to manage scalable objects with Scalable Systems Manager.

From IBM Director Console, start the Scalable Systems Validation task by dragging it onto a scalable system in the Group Contents pane. To start this task for multiple scalable systems, select the scalable systems that you want in the Group Contents pane before you drag the task onto one of the selected systems. You can also start this task by dragging it onto the Scalable Systems group in the Groups pane of IBM Director.

In addition to invoking the Scalable Systems Validation task by dragging it in IBM Director Console, you can also select it from the Scheduler when you create a new scheduled job. For information about the Scheduler, see the help in IBM Director Console.

From Scalable Systems Configuration, start the Scalable Systems Validation task by right-clicking the name of a scalable system; then, clicking **Scalable system validation**.

Note: The scalable nodes in the targeted scalable systems must be unlocked. Otherwise, the validation is skipped for the scalable system that contains that scalable node. For information about unlocking scalable nodes, see "Unlocking physical platforms" on page 20.

After you start the task, Scalable Systems Manager displays a window asking whether you want to schedule a job for the Scalable Systems Validation task or run it immediately. The window lists the scalable systems to be validated, either in the window title bar when only one scalable system is selected, or in a list in the window when multiple scalable systems are selected.

- Click **Schedule** to create an IBM Director scheduled job. For more information about scheduled jobs, click **Help** on the New Scheduled Job window.
- Click **Execute Now** to start the Scalable Systems Validation task immediately.

Results from Scalable Systems Validation

When the validation task is started immediately, a job window opens providing details about the task execution. This window provides the status of the currently running validation job and history information about validation jobs that were previously run on that management server.
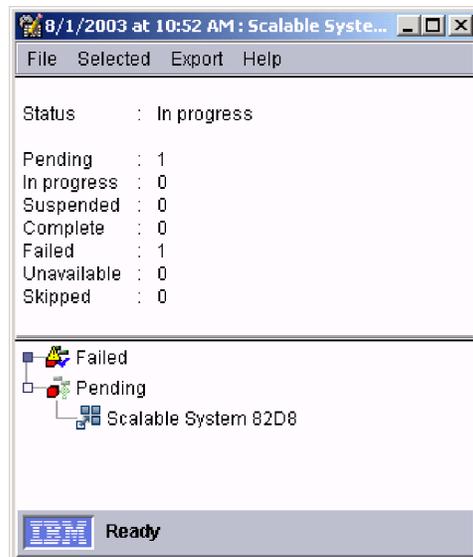
Figure 12. Viewing summary results from Scalable Systems Validation

The lower portion of the window lists the scalable systems that were validated or that failed to be validated. Expand the tree structure for **Successful**, **In progress**, or **Failed** to see the list of scalable systems under each category.

The Scalable Systems Validation process is considered failed for a scalable system when any of the ping commands to the scalability ports on the scalable nodes in the scalable system have failed.

Note: The ping failures returned to Scalable Systems Validation do not necessarily mean that the scalability ports are faulty. For example, if the server processor remote alert recipients are not configured properly, the ping command between some scalability ports will fail even though the ports might be fully functional.

For validation details about a specific scalable system, double-click the name of the scalable system. A window opens that lists the start and completion times for the validation task on that scalable system.
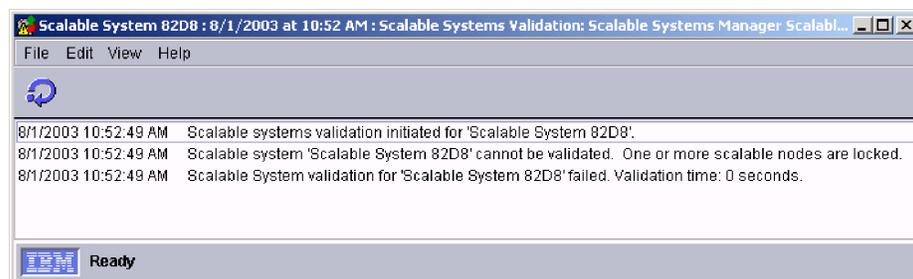
Figure 13. Viewing detailed results from Scalable Systems Validation

Use the **View** → **Detail** menu function on the window to display low, medium, or high levels of detail. By default, a low level of detail is displayed. Use the **File** → **Clear All Messages** menu function to remove all messages from the validation log file. This option removes all messages from the log, not just those for the scalable system that you were viewing.

Chapter 6. The Scalable Systems Configuration task

The Scalable Systems Configuration task presents the graphical user interface provided by Scalable Systems Manager for managing scalable systems, scalable partitions, scalable nodes, and remote I/O enclosures that are defined as managed objects on IBM Director Server. The "Scalable Systems Configuration" window presents this information through icons and detailed information tables.

Starting the Scalable Systems Configuration task

The Scalable Systems Configuration task displays information about scalable objects in a graphical user interface so you can create and manage scalable objects. The Scalable Systems Configuration task can be started for all scalable objects that are known to IBM Director Server, or it can be started for only specific scalable objects. Regardless, when you start the Scalable Systems Configuration task, it displays only those physical platforms that represent scalable nodes.

When you first start Scalable Systems Configuration, the status bar at the bottom of the Scalable Systems Configuration window indicates that it is initializing. During this stage, Scalable Systems Configuration is enumerating the scalable objects that are to be displayed. Depending on the number and complexity of scalable objects, this process can take several minutes. When the initialization completes, the status bar is updated to indicate that Scalable Systems Configuration is ready.

Starting Scalable Systems Configuration for all scalable objects

Double-click the **Scalable Systems Configuration** task in the Tasks pane of IBM Director Console to start Scalable Systems Configuration for all scalable objects.

The "Scalable Systems Configuration" window opens and displays a tree of all previously defined scalable systems, scalable partitions, scalable nodes, and attached remote I/O enclosures on IBM Director Server. When there are numerous scalable objects on the management server, this method can take several minutes for Scalable Systems Configuration to start. However, the advantage of starting Scalable Systems Configuration in this manner is that all scalable objects can be managed from one window.

Starting Scalable Systems Configuration for specific scalable objects

Select one or more scalable systems, scalable nodes, scalable partitions, or remote I/O enclosures in the Group Contents pane of IBM Director Console and drag the Scalable Systems Configuration task onto one of the selected scalable objects to start it only for the selected objects.

The Topology pane in Scalable Systems Configuration contains only the selected scalable systems and only those scalable systems that contain the scalable partitions, scalable nodes, or remote I/O enclosures that have also been selected. Scalable Systems Configuration determines the containing scalable system and displays the entire tree structure for that scalable system, not just for the scalable objects that you selected. The Unassigned pane does not contain any scalable nodes unless you selected unassigned scalable nodes from the Group Contents pane before you started Scalable Systems Configuration.

The advantage of starting Scalable Systems Configuration in this manner is that Scalable Systems Manager must initialize only a subset of scalable objects instead of initializing all scalable objects on IBM Director Server. For example, if you want

to create only a scalable system and the scalable nodes that you want are displayed in the Group Contents pane of IBM Director Console, you can multi-select the scalable nodes in the Group Contents pane and then drag the Scalable Systems Configuration task onto them. Scalable Systems Configuration starts more quickly when you use this method. You can then create a scalable system to contain the unassigned scalable nodes.

The disadvantage is that you can manage only those scalable objects that were selected when you started Scalable Systems Configuration. If you later need to manage other scalable objects, you must start another Scalable Systems Configuration task for those additional scalable objects.

The Scalable Systems Configuration window

After you start the Scalable Systems Configuration task, the "Scalable Systems Configuration" window opens and provides a graphical user interface for managing the selected scalable systems, scalable partitions, scalable nodes, and remote I/O enclosures on the management server. The following illustration shows a typical "Scalable Systems Configuration" window.
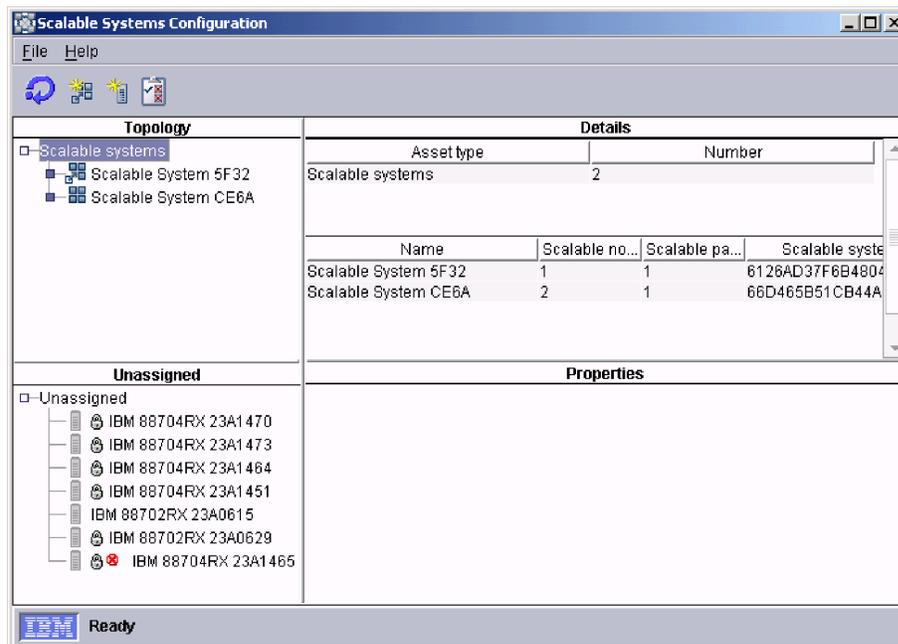
Figure 14. "Scalable Systems Configuration" window

The left side of the window contains the Topology and Unassigned panes, which contain tree structures that can be expanded and collapsed.

- The Topology pane displays the scalable systems that were selected when you started Scalable Systems Configuration. The root of the scalable systems tree structure is the word *Scalable systems*. When you expand an individual scalable system, the tree structure shows a Scalable nodes folder that identifies the nodes in the scalable system and a Scalable partitions folder that identifies the scalable partitions that are defined for that scalable system.

  Note: If IBM Director discovers scalable objects after you have started the Scalable Systems Configuration task, the discovered scalable objects are not displayed in Scalable Systems Configuration as soon as they have

been discovered. Instead, you must exit Scalable Systems Configuration and restart it for the newly discovered scalable objects.

For example, if Scalable Systems Configuration is running and IBM Director discovers a remote I/O enclosure when you unlock a scalable node, this remote I/O enclosure is not displayed in Scalable Systems Configuration until you restart the task.

- The Unassigned pane displays all supported scalable nodes that do not belong to a scalable system. The root of the unassigned tree structure is the word *Unassigned*.

The right side of the window contains the Details and Properties panes, which use tables to display information about the scalable objects that are currently selected in the tree structure.

Note: The Scalable Systems Configuration window displays only those scalable objects for which it was started, which might not be all scalable objects on the management server.

Both panes use icons to represent the various scalable objects (scalable system, scalable partition, scalable node, or remote I/O enclosure). The icon for a scalable partition reflects whether it is powered-on or powered-off, and the icon for a scalable node indicates whether it is part of a scalable partition that is powered-on. The icon for a remote I/O enclosure indicates the sides that are attached to the scalable node with which it is displayed. For samples of all the scalable object icons, see "Icons in Scalable Systems Manager" on page 5.

Note: Scalable Systems Configuration will update the icons and table information whenever events and scalable object operations occur. However, for this to happen, you must use the Management Processor Configuration subtask of MPA to configure an alert-forwarding profile for the servers that are represented by the scalable nodes that you plan to manage in Scalable Systems Manager. In the alert-forwarding profile, the connection type must be set to "IBM Director Comprehensive" and the IP address must be set to the IP address of the management server that is being used to manage scalable objects.

Additionally, you can force the window to refresh in different ways, as explained in "Refreshing the Scalable Systems Configuration window" on page 47.

Status icons in Scalable Systems Configuration

Scalable Systems Manager assigns hardware status to a scalable node or remote I/O enclosure object according to the status of the underlying hardware from the Hardware Status task. Scalable Systems Manager does not affect or control the status of this hardware other than reporting the status that is assigned by the Hardware Status task. Scalable Systems Configuration displays icons beside a scalable node or remote I/O enclosure when its corresponding hardware has a Warning or Critical status. These icons are described in Table 11.

Table 11. Status change icons

Icon	Description
	The startup disk icon is displayed beside the icon for a scalable node when it represents the primary scalable node.

Table 11. Status change icons (continued)

Icon	Description
⚠️	The Warning icon is displayed beside the icon for a scalable node or remote I/O enclosure when the underlying hardware has a warning status that requires attention. The Warning icon is displayed when something is wrong with the hardware but it is still usable.
⊗	The Critical icon is displayed beside the icon for a scalable node or remote I/O enclosure when the underlying hardware has a critical status that requires immediate attention. The Critical icon is displayed when the object is nearly or completely unusable. For example, a scalable node can be in a critical status when it has a problem with a microprocessor or memory.

Help text for Scalable Systems Configuration

The help pages that explain the Scalable Systems Configuration interface are displayed through the IBM Director help interface. This help system also documents all the features of Scalable Systems Manager, including Scalable Systems Discovery and Scalable Systems Validation. Each page is identified by the same icon that is used in IBM Director Console to identify the Scalable Systems Manager task. Use the Find icon in the IBM Director help interface to search within the current help document or the Search icon to search within all the help that is provided for IBM Director and its extensions. Use the links at the bottom of each help page to navigate the help that is provided for Scalable Systems Manager and its subtasks.

Displaying scalable object information in Scalable Systems Configuration window

The Details and Properties panes on the right side of the Scalable Systems Configuration window display various properties for the scalable systems, scalable nodes, scalable partitions, and remote I/O enclosures that are displayed in the Topology and Unassigned panes on the left side of the window.

- If you select a specific scalable system, scalable node, or scalable partition in the Topology or Unassigned panes, the top table in the Details pane displays the UUID for the selected object and important properties such as the object name and status. The table in the Properties pane displays additional general information about the selected object.

- If you click the words **Scalable systems**, **Scalable nodes**, **Scalable partitions**, or **Unassigned** in the Topology or Unassigned pane, this top table identifies the asset type and the number of assets of this type that are defined in IBM Director Server (for **Scalable systems** or **Unassigned**) or in the scalable system (for **Scalable nodes** or **Scalable partitions**).

  A second table in the Details pane contains a summary of the children for the selected tree node. For example, if the **Scalable systems** tree node is selected, the second table shows a list of scalable systems that are being managed in Scalable Systems Configuration and their primary property values. You can select one of the children in this table to display property values for the individual object in the Properties pane.

Details about the types of information that are displayed about each scalable object are in Chapter 8, "Viewing details about scalable objects in Scalable Systems Configuration", on page 71.

Invoking Scalable Systems Configuration functions

Operations are invoked in Scalable Systems Configuration primarily through the toolbar or a menu that is displayed when you right-click a word or object in the Topology or Unassigned pane. The menu items that are available depends on the word or object that is selected and its current state. For example, if you select a scalable partition that is powered-on, the menu item to power-off the scalable partition is available, but the menu item to power-on the scalable partition is unavailable. For details about the toolbar or menu operations, see Chapter 7, "Managing scalable objects", on page 49.

The toolbar contains the functions that are described in the following table:

Table 12. Toolbar icons and their actions

Icon	Action
	Refreshes the entire tree structure in both the Topology pane and the Unassigned pane.
	Opens the "Create scalable system" window.
	Opens the "Create a new scalable node" window.
	Checks for problems with scalable partitions and scalable systems and displays any problems in the "Problem Checker" window.

You can use the menu bar to run the problem checker to check for scalable-system problems and scalable-partition problems, access the Scalable Systems Configuration help, access an index to the IBM Director help system, and close the Scalable Systems Configuration window.

Refreshing the Scalable Systems Configuration window

You can use the Refresh function to request an update of the information that is shown in the "Scalable Systems Configuration" window. Scalable Systems Configuration obtains its information from persistent storage in IBM Director Server and then caches the information on the console for subsequent display in Scalable Systems Configuration. See Chapter 8, "Viewing details about scalable objects in Scalable Systems Configuration", on page 71 for details about the types of information that is displayed by Scalable Systems Configuration.

By default, the window view is refreshed automatically whenever an event occurs or when you perform any operations on scalable systems, scalable partitions, or scalable nodes. With the Refresh function, you can force a refresh request rather than waiting for an event or operation to occur.

You can refresh information for portions of the tree structure or for the entire tree structure in the Topology pane and Unassigned pane of Scalable Systems Configuration. The following table describes the ways in which you can refresh a window.

Table 13. Ways to refresh information in the Scalable Systems Configuration window

Refresh target	Action to take	Result
Entire tree structure in both the Topology pane and the Unassigned pane	In the toolbar, click the 🔄 icon.	Scalable Systems Configuration destroys the entire cached tree structure on the window and re-creates it from information that is stored on the management server. If the management server contains several scalable objects, the refresh operation can take several minutes. The hierarchy of each scalable system is collapsed.
Single scalable system	Right-click the scalable system; then, click **Refresh scalable system**.	Scalable Systems Configuration deletes the tree structure for the selected scalable system from the cache and re-creates it from information stored that is on the management server. The hierarchy of the scalable system is fully expanded.
Single scalable partition	Right-click the scalable partition; then, click **Refresh**.	Scalable Systems Configuration deletes the cached information for the properties of the scalable partition and any scalable nodes that it contains and re-creates the scalable partition by using information from the management server. The displayed structure of the scalable partition remains as it was before the refresh operation. For example, if before the refresh operation, the scalable partition hierarchy is expanded to show the scalable nodes in the scalable partition, after the refresh operation, it remains expanded.
Single scalable node in a scalable partition or scalable system	Right-click the scalable node; then, click **Refresh**.	Scalable Systems Configuration deletes the cached information for the properties of the scalable node and re-creates it from information that is stored on the management server. If the scalable node has an attached remote I/O enclosure, the displayed structure of the scalable node remains as it was before the refresh operation.
Single scalable node that is unassigned.	Right-click the scalable node that is unassigned; then, click **Refresh**.	Scalable Systems Configuration deletes the cached information for the properties of the unassigned scalable node and re-creates it from information that is stored on the management server. If the unassigned scalable node has an attached remote I/O enclosure, the displayed structure of the unassigned scalable node remains as it was before the refresh operation.

During the refresh operation, the status line of Scalable Systems Configuration displays the word *Refreshing* to indicate that the operation is still active. After the refresh operation is complete, the status line displays the word *Ready*.

Chapter 7. Managing scalable objects

You can manage scalable objects as follows:

- Create scalable objects
- Rename scalable objects
- Configure scalable partitions
- Perform pow3er operations for scalable partitions
- Update scalable node selections for a scalable system
- Check for problems with scalable systems and scalable partitions
- Delete scalable objects

Creating scalable objects

You can use the Scalable Systems Configuration task to create the following scalable objects:

- Unassigned scalable nodes
- Scalable systems
- Scalable partitions

IBM Director assigns each scalable object a UUID.

Manually creating unassigned scalable nodes

When IBM Director discovers physical platforms, servers that contain SMP Expansion Modules are represented as scalable nodes. This process can take several minutes to complete because IBM Director queries each server on the network.

If you later add scalable servers to the network, you can avoid invoking physical-platform discovery for all servers on the network by creating a physical-platform managed object. IBM Director runs physical-platform discovery for only the designated server. In Scalable Systems Manager, this feature is called *creating unassigned scalable nodes*.

Complete the following steps to create an unassigned scalable node:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Unassigned pane, right-click **Unassigned**; then, click **Create scalable node**. The "Create a new scalable node" window opens.
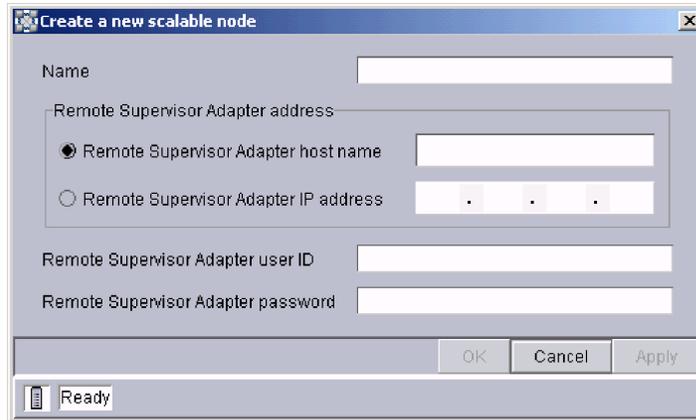
Figure 15. "Create a new scalable node" window

3. In the **Name** field, type the name of the scalable node. You can use the same name for more than one scalable node.

4. Use one of the following fields to provide information about the service processor in the scalable node:

   Remote Supervisor Adapter host name
   > Type the host name of the service processor.

   Remote Supervisor Adapter IP Address
   > Type the IP address of the service processor.

5. In the **Remote Supervisor Adapter user ID** field, type a valid user name.

6. In the **Remote Supervisor Adapter password** field, type the password for the user name that you typed in step 5.

7. Click **OK**.

   When the scalable node is created, the "Scalable Systems Configuration" window is refreshed to reflect the new scalable node. Because it can take several seconds, Scalable Systems Configuration displays a confirmation message after IBM Director has successfully created the scalable node.

Creating scalable systems

You can create a scalable system by using the Scalable Systems Configuration task. When you create a scalable system, you can name the scalable system and select the unassigned scalable nodes that are to be assigned to the scalable system. A scalable node can be part of one scalable system only. You can use the same name for more than one scalable system.

Complete the following steps to create a scalable system that contains scalable nodes:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.

2. In the Topology pane, right-click **Scalable systems**; then, click **New scalable system**. The "Scalable system name" window opens.
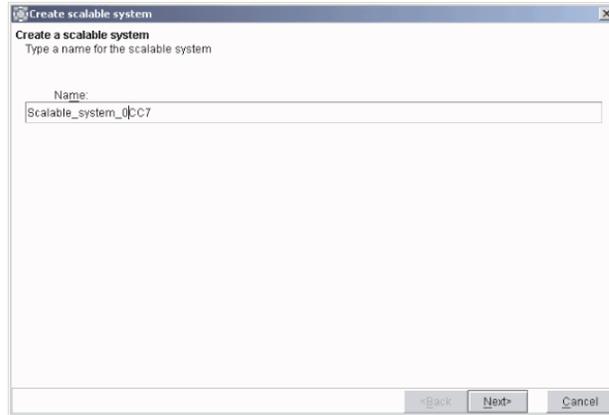
Figure 16. Creating a scalable system: "Scalable system name" window

3. In the **Name** field, type a scalable system name.
4. Click **Next**. The "Add scalable nodes" window opens.

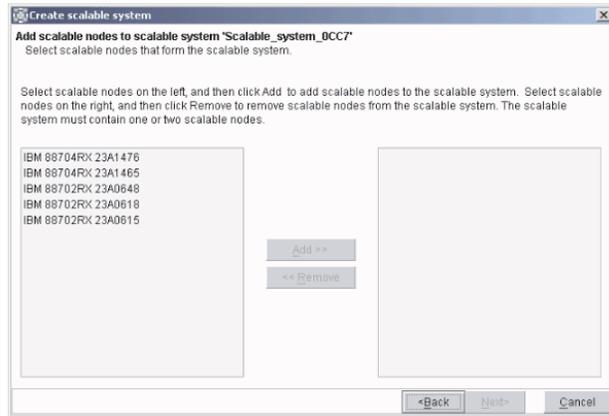*Figure 17. Creating a scalable system: "Add scalable nodes" window*

5. Select two scalable nodes:
   a. In the left list, click one scalable node.
   b. Press Control and click the second scalable node.
   c. Click **Add**.
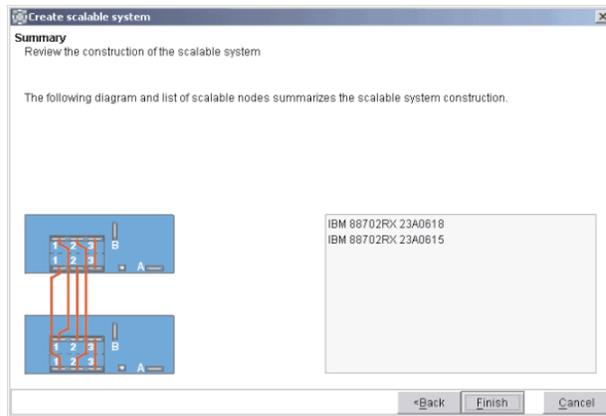6. Click **Next**. The "Summary" window opens.

Figure 18. Creating a scalable system: "Summary" window

7. Review the scalable system. The proposed layout should match how you think the servers and remote I/O enclosures are cabled. To make changes, click **Back**.

   Note: The ordering of the scalable nodes in the graphic does not correspond to the ordering of the scalable nodes as shown in the list on the left. For example, the first scalable node in the graphic is not necessarily the first scalable node in the list on the left.

8. Click **Finish**.

The scalable system is created. Because it can take several seconds, Scalable Systems Configuration displays a completion message after IBM Director has successfully created the scalable system.

The new scalable system does not contain any scalable partitions; you must create them.

Creating scalable partitions

You can create scalable partitions from Scalable Systems Configuration from the scalable nodes in a manageable scalable system.

Complete the following steps to create a scalable partition:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Topology pane, expand the tree structure for a scalable system.
3. Right-click **Scalable partitions**; then, click **New scalable partition**. The "Create a new scalable partition" window opens.
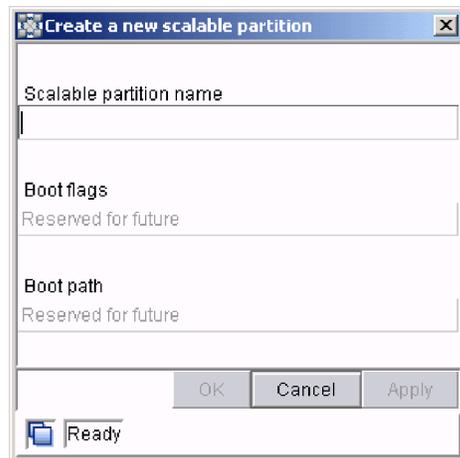
Figure 19. "Create a new scalable partition" window

4. Type the scalable partition name. You can use the same name for more than
    one scalable partition.

    Note: The scalable partition name must not contain more than 31 characters.
5. Click **OK**.

Because it can take several seconds, Scalable Systems Configuration displays a
completion message after IBM Director has successfully created the scalable
partition.

The new scalable partition does not contain any scalable nodes; you must add
scalable nodes and define the primary scalable node.

Renaming scalable objects

You can rename scalable objects (scalable systems, scalable partitions, scalable
nodes, or remote I/O enclosures) by using IBM Director Console. You cannot
rename scalable objects from within the "Scalable Systems Configuration" window.

Complete the following steps to rename scalable objects:
1. From IBM Director Console, in the Group Contents pane, right-click the scalable
    object; then, click **Rename**. The "Rename System" window opens.

Figure 20. "Rename System" window

2. Type the new name. You can use the same name for more than one scalable
    object.

Note: The name for a scalable partition must not contain more than 31 characters.

3. Click **OK**. The scalable object is renamed.

Configuring scalable partitions

This section contains procedures for performing the following tasks:

- Inserting a scalable node
- Setting the primary scalable node
- Removing a scalable node
- Configuring a remote I/O enclosure

Inserting scalable nodes

When you create a scalable partition, the scalable partition does not contain any scalable nodes. You must insert at least one scalable node into each scalable partition. You can insert scalable nodes into a scalable partition as follows:

- Use the Scalable Systems Configuration menus
- Drag scalable nodes onto scalable partitions

After you add scalable nodes to a scalable partition, you must set the primary scalable node for the scalable partition.

Complete the following steps to insert a scalable node into a scalable partition:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.

2. In the Topology pane, expand the tree structure for the scalable system that contains the scalable partition.

3. Right-click the scalable partition; then, click **Insert scalable node**. The scalable nodes that you can add to the scalable partition are displayed.
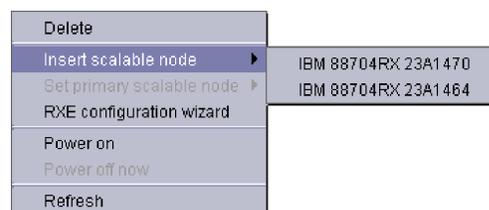

Figure 21. Insert scalable node menu option

4. Click the scalable node that you want to insert into the scalable partition.

Setting the primary scalable node

Before you can power-on a scalable partition, you must designate one scalable node as the primary scalable node. For information about how to determine the primary scalable node to use for a scalable partition, see "Determining the primary scalable node for a scalable partition" on page 29.

Important: Before you can designate a different primary scalable node in a scalable partition, you must shut down and power-off the affected scalable partition.

Complete the following steps to set the primary scalable node:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.

2. In the Topology pane, fully expand the tree structure for the scalable system that contains the scalable partition.

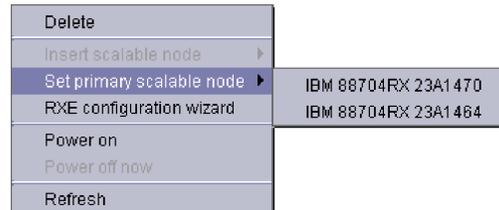3. Right-click the scalable partition; then, click **Set primary scalable node**.

Delete
Insert scalable node ▶
Set primary scalable node ▶	IBM 88704RX 23A1470
RXE configuration wizard	IBM 88704RX 23A1464
Power on
Power off now
Refresh

Figure 22. Set primary scalable node menu option

4. Click the scalable node that you want to set as the primary scalable node.

Alternatively, you can right-click the scalable node that you want to become the primary scalable node for the scalable partition. Then, click **Set primary scalable node**.

In the Topology pane of the "Scalable Systems Configuration" window, the   icon is displayed beside the primary scalable node.

Removing scalable nodes

After a scalable node has been removed from a scalable partition, you can remove it from the scalable system; then, the scalable node becomes unassigned.

Note: If you remove either the primary scalable node or all the scalable nodes from a scalable partition, you cannot perform power operations on the scalable partition.

Complete the following steps to remove a scalable node from a scalable partition:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.

2. In the Topology pane, expand the tree structure for the scalable system that contains the scalable partition.

3. Right-click the scalable node that you want to remove; then, click **Remove**. A window opens asking you to confirm that you want to remove the indicated scalable node from the scalable partition.

4. Click **Yes, remove**.

Configuring a remote I/O enclosure

After a scalable node with an attached remote I/O enclosure has been inserted into a scalable partition, you can use the RXE Configuration wizard in Scalable Systems Configuration to configure the enclosure for the scalable partition to which it belongs.

Note: The new enclosure configuration does not take affect until the next time that the scalable partition is started. For this reason, you do not have to shut down and power-off the scalable partition that contains the enclosure before you configure it.

By default, for a stand-alone server, IBM Director configures both expansion kits of an enclosure for use, which is not supported by Scalable Systems Manager at the time of publication. However, by using the RXE Configuration wizard, you can change the enclosure configuration to disable at least one expansion kit. You can also change which RXE Expansion cables and RXE Expansion Ports are used to link the enclosure to the server. The choices available depend on the cabling configuration that IBM Director initially discovers for a remote I/O enclosure.

For more information about remote I/O enclosure discovery, see "Remote I/O enclosures" on page 30.

Complete the following steps to use the RXE Configuration wizard:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. Right-click the scalable partition that contains the enclosure; then, click **RXE configuration wizard**. A window opens asking you to confirm that you want to update the enclosure configuration.
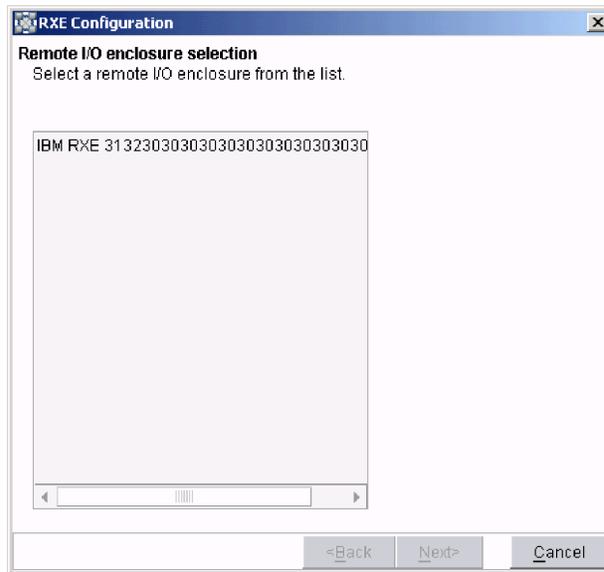3. Click **Yes, update**. The "Remote I/O enclosure selection" window opens.

Figure 23. Configuring a remote I/O enclosure: "Remote I/O enclosure selection" window listing the remote I/O enclosures in the scalable partition

4. From the list of enclosures that are part of the scalable partition, select the enclosure that you want to configure. The "Remote I/O enclosure selection" window is updated to display a graphic depicting the scalable partition and the current configuration of the enclosure that you selected.
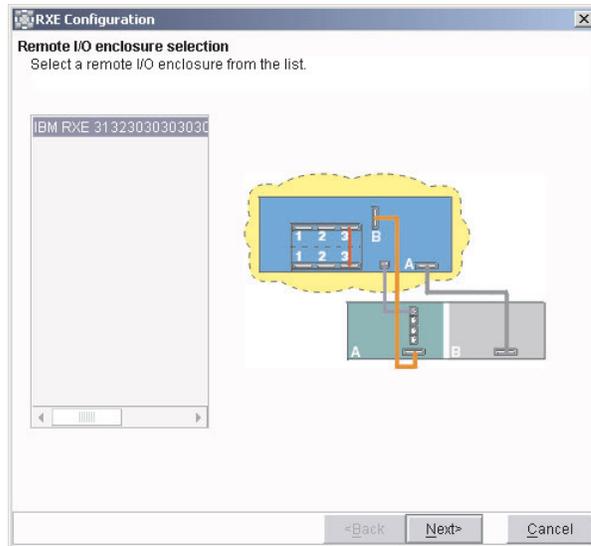
Figure 24. Configuring a remote I/O enclosure: "Remote I/O enclosure selection" window illustrating the current configuration of the selected remote I/O enclosure

By default, the graphic shows that both expansion kits of the enclosure are enabled for use. For more information about the graphics, see "Discovered remote I/O enclosure configurations" on page 59.

5. Click **Next**. The "Remote I/O enclosure configuration selection" window displays configuration selections that are applicable to the enclosure that you have selected. By default, the window selects the choice that corresponds to the current enclosure configuration. Choices are unavailable if they are not applicable to the selected enclosure.
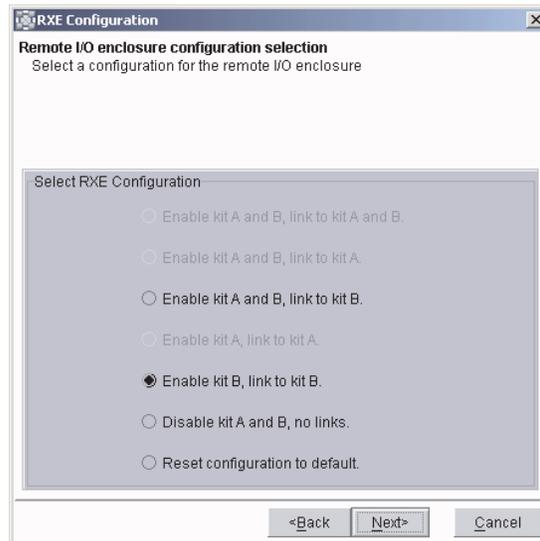

Figure 25. Configuring a remote I/O enclosure: "Remote I/O enclosure configuration selection" window

You can select the following choices, if they are available:

Table 14. Configuration choices in RXE Configuration wizard

Enclosure configuration	Description
Enable kit A and B, link to kit A and B	Enable both expansion kits A and B for use by the scalable partition. The scalable partition will use the RXE Expansion cable on both kit A and B as needed. **Note:** At the time of publication, this configuration choice is not supported when a scalable partition contains only one scalable node.
Enable kit A and B, link to kit A	Enable both expansion kits A and B for use by the scalable partition. The scalable partition will use the RXE Expansion cable to kit A. **Note:** At the time of publication, this configuration choice is not supported when a scalable partition contains only one scalable node.
Enable kit A and B, link to kit B	Enable both expansion kits A and B for use by the scalable partition. The scalable partition will use the RXE Expansion cable to kit B. **Note:** At the time of publication, this configuration choice is not supported when a scalable partition contains only one scalable node.
Enable kit A, link to kit A	Enable expansion kit A for use by the scalable partition. The scalable partition will use the RXE Expansion cable to kit A.
Enable kit B, link to kit B	Enable expansion kit B for use by the scalable partition. The scalable partition will use the RXE Expansion cable to kit B.
Disable kit A and B, no links	Disable both expansion kits A and B for use by the scalable partition. The scalable partition cannot use any RXE Expansion cable to the enclosure.
Reset configuration to default	Reset the enclosure configuration to the original configuration that was detected by IBM Director.

6. Click the enclosure configuration that you want.
7. Click **Next**. The "Summary" window opens with a graphic depicting the scalable partition and the new configuration of the enclosure.
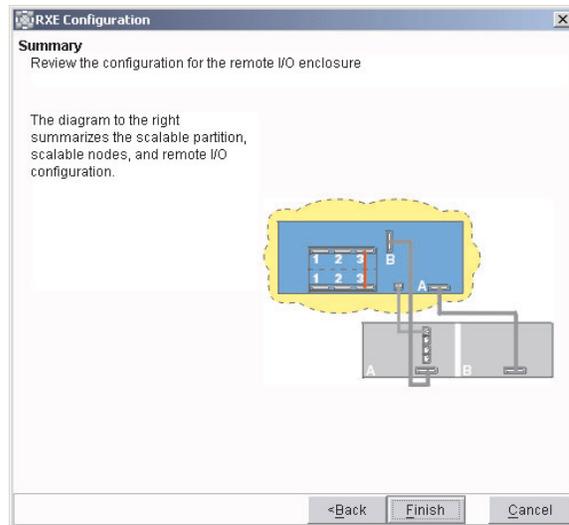
Figure 26. Configuring a remote I/O enclosure: "Summary" window

This graphic uses purple cables, orange cables, and green expansion kits to show the configuration elements that can be used by the scalable partition. It uses gray cables and gray expansion kits to show the configuration elements that cannot be used by the scalable partition. Other details of the graphic notation are explained in "Discovered remote I/O enclosure configurations" on page 59.

8. Click **Finish** if the graphic represents your choice. Otherwise, click **Back** to select a different enclosure configuration.

9. Click **OK**.

Discovered remote I/O enclosure configurations

The RXE Configuration wizard displays the following graphics that show the configurations of remote I/O enclosures that IBM Director will discover. The following table describes the information that these graphics depict about each server and enclosure.

Table 15. Notations used in graphics for RXE Configuration wizard

Depicted in graphic	Notation used
Scalable partition	Yellow bubble
Scalable nodes that are defined in the scalable partition	Blue rectangles
Expansion kits of the enclosure that are enabled for use by the scalable partition	Green rectangles
RXE Management cables and their connection to RXE Management Ports	Purple solid lines
RXE Expansion cables and their connection to RXE Expansion Ports	Orange solid lines
Optional RXE Expansion cables and their connection to RXE Expansion Ports	Orange dashed lines
SMP Expansion cables and their connection to SMP Expansion Ports	Red solid lines
The scalable node and scalable partition of an 8-way server in a 16-way configuration that is not applicable to the scalable partition whose enclosure is being configured	Gray shading

Table 15. Notations used in graphics for RXE Configuration wizard (continued)

Depicted in graphic	Notation used
The RXE Expansion cables, RXE Management cables, and expansion kits that are disabled for use by the scalable partition whose enclosure is being configured	Gray shading

For details about cabling, see the hardware documentation that comes with the server and enclosure.

Discovered enclosure configurations for scalable partitions on 4-way servers

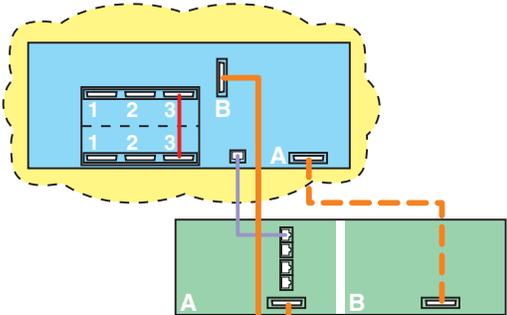Scalable Systems Manager can create one scalable partition for each 4-way configuration.

Table 16. Discovered enclosure configurations for scalable partitions on 4-way servers

Enclosure configuration	Description
	A scalable partition that includes one 4-way server that is attached to expansion kit A of a remote I/O enclosure. **Note:** You must use the RXE Configuration wizard to enable only one expansion kit or disable both expansion kits in this scalable partition. As discovered, Scalable Systems Manager does not support this scalable partition.For information about the RXE Configuration wizard, see "Configuring a remote I/O enclosure" on page 55.
	A scalable partition that includes one 4-way server that is attached to expansion kit B of a remote I/O enclosure. **Note:** You must use the RXE Configuration wizard to enable only one expansion kit or disable both expansion kits in this scalable partition. As discovered, Scalable Systems Manager does not support this scalable partition. For information about the RXE Configuration wizard, see "Configuring a remote I/O enclosure" on page 55.

Discovered enclosure configurations for scalable partitions on 8-way servers

Scalable Systems Manager can create one scalable partition for each 8-way configuration.

Table 17. Discovered enclosure configurations for scalable partitions on 8-way servers

Enclosure configuration	Description
	A scalable partition that includes one 8-way server that is attached to only expansion kit B of a remote I/O enclosure. **Note:** You must use the RXE Configuration wizard to enable only one expansion kit or disable both expansion kits in this scalable partition. As discovered, Scalable Systems Manager does not support this scalable partition. For information about the RXE Configuration wizard, see "Configuring a remote I/O enclosure" on page 55.
	A scalable partition that includes one 8-way server that is attached to both expansion kit A and B of a remote I/O enclosure. **Note:** You must use the RXE Configuration wizard to enable only one expansion kit or disable both expansion kits in this scalable partition. As discovered, Scalable Systems Manager does not support this scalable partition. For information about the RXE Configuration wizard, see "Configuring a remote I/O enclosure" on page 55.

Discovered enclosure configurations for scalable partitions on 16-way servers

Scalable Systems Manager can create one scalable partition that consists of one 16-way server and one scalable partition per 8-way server, for a total of three unique partition configurations.

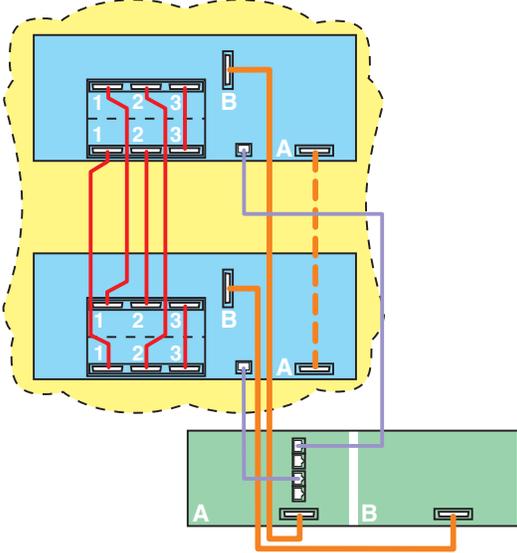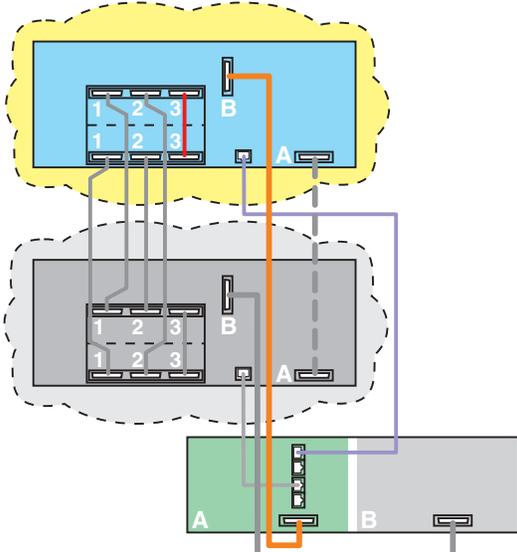*Table 18. Discovered enclosure configurations for scalable partitions on 16-way servers*
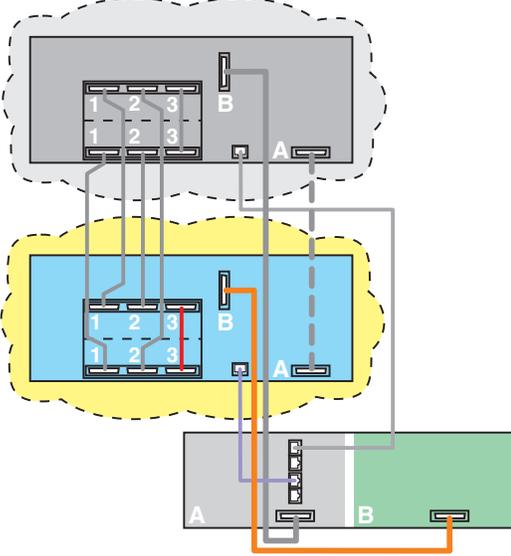
Enclosure configuration	Description
	A scalable partition that includes both scalable nodes of a 16-way configuration that is attached to both expansion kit A and B of a remote I/O enclosure.
	A scalable partition that includes one scalable node of a 16-way configuration that is attached to expansion kit A of a remote I/O enclosure. **Note:** The RXE Configuration wizard offers configuration choices that enable both kits of an enclosure when the scalable partition contains only one scalable node. At the time of publication, these configuration choices are not supported by Scalable Systems Manager. Do not use the RXE Configuration wizard to enable both expansion kits for a single scalable node. For information about the RXE Configuration wizard, see "Configuring a remote I/O enclosure" on page 55.

Table 18. Discovered enclosure configurations for scalable partitions on 16-way servers (continued)

Enclosure configuration	Description
	A scalable partition that includes one scalable node of a 16-way configuration that is attached to expansion kit B of a remote I/O enclosure. **Note:** The RXE Configuration wizard offers configuration choices that enable both kits of an enclosure when the scalable partition contains only one scalable node. At the time of publication, these configuration choices are not supported by Scalable Systems Manager. Do not use the RXE Configuration wizard to enable both expansion kits for a single scalable node. For information about the RXE Configuration wizard, see "Configuring a remote I/O enclosure" on page 55.

Performing power operations for scalable partitions

You can use the Scalable Systems Configuration subtask to use out-of-band communications to power-on and power-off scalable partitions. If an IBM Director managed-system object represents the scalable partition, you can also use in-band communication to power-off the scalable partition.

Alternatively, you can create event action plans that trigger automatic power operations to occur by using the "Power on/shut down and power off scalable partition" event action. For more information about this event action, see "Using event actions with scalable partitions" on page 10.

Scalable Systems Configuration uses different icons for each scalable partition to indicate whether it is powered-off (inactive, 🖥️) or powered-on (active, 🖥️) . Scalable Systems Configuration also uses different icons for the scalable nodes in a scalable partition to indicate whether they are powered-off (inactive, ▯) or powered-on (active, ▯) . These icons are updated each time a scalable partition and its scalable nodes are powered-on or powered-off.

Important: Scalable Systems Manager cannot communicate with the service processors on the servers in a scalable partition unless the following conditions are met:
- Service-processor networks must either have static IP address assignments or have DHCP configured to maintain consistent IP addresses for the service processors. To do so, create reservations in DHCP that identify explicit IP addresses for the service processors.
- The IP addresses that are assigned to the service processors do not change after the servers are discovered in IBM Director. This limitation is true whether you are using DHCP or statically assigning

IP addresses. Chapter 9, "Solving Scalable Systems Manager problems", on page 87 describes suggested actions to take if the IP address of a service processor has changed since it was discovered by IBM Director.

IBM Director Console identifies all scalable partitions with the icon whether they are powered-on or powered-off. However, IBM Director Console uses additional icons with this scalable-partition icon to indicate the state of a scalable partition. For more information, see "States of scalable partitions" on page 83. 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.

If you have a scalable partition that can no longer be powered-on or powered-off, see Chapter 9, "Solving Scalable Systems Manager problems", on page 87.

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.

Note: Before you can power-on a scalable partition, these conditions must be met:
- IBM Director must be able to access all scalable nodes in the scalable partition. If any scalable nodes are locked, you first must request access to the scalable nodes. For details, see "Unlocking physical platforms" on page 20.
- You must designate a primary scalable node for the scalable partition. For details, see "Setting the primary scalable node" on page 54.

You cannot power-on a scalable partition that contains a scalable node that is part of an already started scalable partition. Also, if the scalable partition that you want to power-on contains a scalable node with an attached remote I/O enclosure and that scalable node is in an already powered-on scalable partition, you will not have the use of the slots in the remote I/O enclosure. To access these slots, you must power-off the first scalable partition before powering-on the second scalable partition.

Complete the following steps to power-on a scalable partition:
1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Topology 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**.

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.

Shutting down and powering-off a scalable partition

When a scalable partition is powered-on, you can use the Scalable Systems Configuration task to shut down and power-off the scalable partition. Doing so causes IBM Director to notify 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 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 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 Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Topology pane, expand the tree structure for the scalable system that contains the scalable partition.
3. Right-click the active scalable partition that you want to power-off; then, click **Shutdown and power off**. A window opens and prompts you to confirm that you want to power-off the scalable partition.
4. Click **Yes, shutdown**.

Powering off a scalable partition

When a scalable partition is powered-on, you can use the Scalable Systems Configuration task to immediately power-off the scalable partition. Doing so causes IBM Director to notify 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.

Alternatively, you can shut down the operating system and application processes in a more orderly way before the servers in the scalable partition are powered-off. See "Shutting down and powering-off a scalable partition" on page 65 for details.

Complete the following steps to immediately power-off a scalable partition:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Topology pane, expand the tree structure for the scalable system that contains the scalable partition.
3. Right-click the scalable partition that you want to power-off; then, click **Power off now**. A window opens and prompts you to confirm that you want to power-off the scalable partition.

4. Click **Yes, power off**.

Updating a scalable system

You can use the Scalable Systems Configuration task to update a scalable system. You can add or remove scalable nodes, although the resulting scalable system must contain one or two scalable nodes.

Complete the following steps to update a scalable system:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.

2. In the Topology pane, expand the tree structure until you can see the scalable system that you want to update.

3. Right-click the scalable system that you want to update; then, click **Update scalable system**. A window opens and prompts you to confirm that you want to update the scalable system.

   Important: Before you remove scalable nodes that are part of a scalable partition, you must shut down and power-off the scalable partition.

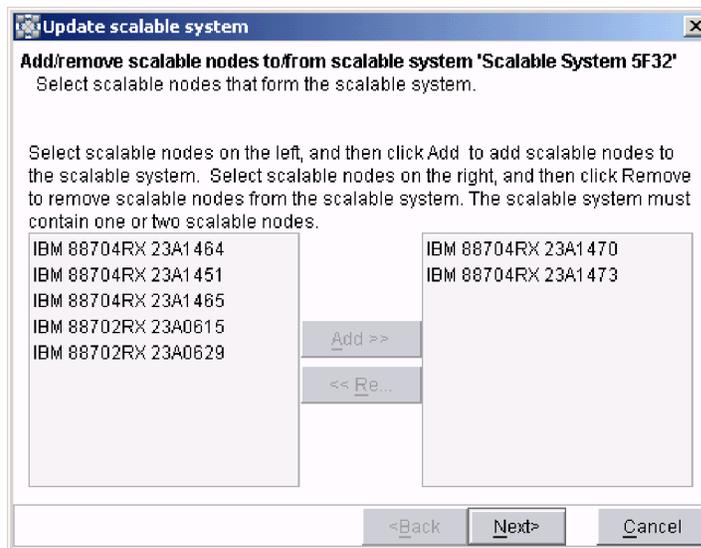4. Click **Yes, update**. The "Add/remove scalable nodes" window opens.

Figure 27. Updating a scalable system: "Add/remove scalable nodes" window

5. Change the scalable nodes that are part of the scalable system. To remove a scalable node, in the right list, select a scalable node and click **Remove**. To add a scalable node, in the left list, select a scalable node and click **Add**.
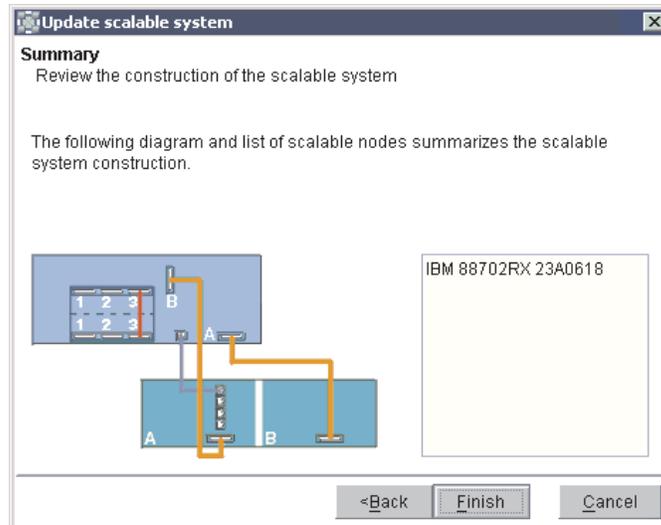
6. Click **Next**. The "Summary" window opens.

Figure 28. Updating a scalable system: "Summary" window

7. Review the scalable system. The proposed layout should match how the servers and remote I/O enclosures are cabled. To make changes, click **Back**.

    Note: The ordering of the scalable nodes in the graphic does not correspond to the ordering of the scalable nodes as shown in the list on the left. For example, in a scalable system with two scalable nodes, the first scalable node in the graphic is not necessarily the first scalable node in the list on the left.

8. Click **Finish**.

Scalable Systems Configuration updates the definition of the scalable system and refreshes the "Scalable Systems Configuration" window to reflect the changes to the scalable system. Scalable nodes that are added to a scalable system are no longer shown in the Unassigned pane, and scalable nodes that are removed from the scalable system are now shown as unassigned.

Checking for problems with scalable systems and scalable partitions

You can use the Scalable Systems Configuration task to check scalable systems and scalable partitions for problems. Complete the following steps to check for problems with scalable systems and scalable partitions:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.

2. Click **File → Problem Checker**.

The Problem Checker starts; if it finds one or more problems, it opens the "Problem Checker" window, which lists each scalable object and the detected problem. You must fix the problem before managing the scalable system or scalable partition.
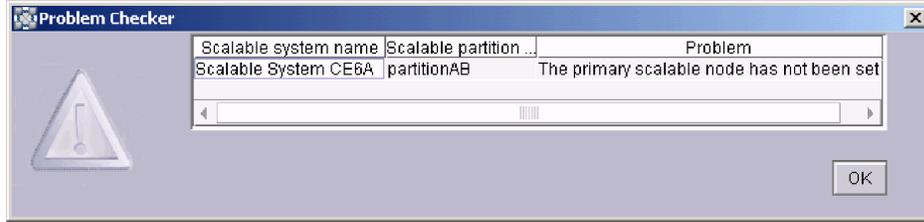
Figure 29. "Problem Checker" window

Possible problems with scalable systems

Table 19 lists the possible problems that the Problem Checker can identify for scalable systems. It also lists the corrective actions that you must take to solve the problem. The variables *n* and *m* represent numeric values that are returned by the Problem Checker.

Table 19. Possible problems with scalable systems

Problem	Corrective action
This scalable system was created from the BIOS Configuration/Setup Utility program menus. It has *n* nodes; *m* is expected.	Update the scalable-node selection for the scalable system so that it contains *m* scalable nodes. **Note:** This problem is not applicable for scalable systems that were discovered by Scalable Systems Discovery or created in Scalable Systems Configuration. These scalable systems do not use the "expected nodes" attribute so there is nothing for the Problem Checker to compare against.
This scalable system has *n* nodes; *m* is the minimum that is supported.	Update the scalable-node selection for the scalable system so it contains at least *m* scalable nodes.
This scalable system has *n* nodes; *m* is the maximum that is supported.	Update the scalable-node selection for the scalable system so it contains no more than *m* scalable nodes.

Possible problems with scalable partitions

Table 20 lists the possible problems that the Problem Checker can identify for scalable partitions. It also lists the corrective actions that you must take to solve the problem. The variables *n* and *m* represent numeric values that are returned by the Problem Checker.

Table 20. Possible problems with scalable partitions

Problem	Corrective action
This scalable partition was created from the BIOS Configuration/Setup Utility program menus. It has *n* nodes, *m* is expected.	Insert scalable nodes into the scalable partition until it has *m* scalable nodes. **Note:** This problem is not applicable for scalable partitions that were manually created in Scalable Systems Configuration. These scalable partitions do not use the 'expected nodes' attribute so there is nothing for the Problem Checker to compare against.
This scalable partition has *n* nodes; *m* is the minimum supported.	Insert scalable nodes into the scalable partition until it has at least *m* scalable nodes.

Table 20. Possible problems with scalable partitions (continued)

Problem	Corrective action
This scalable partition has *n* nodes; *m* is the maximum supported.	Remove scalable nodes from the scalable partition until it contains no more than *m* scalable nodes.
The primary node has not been set for this scalable partition.	Set the primary scalable node for the scalable partition.
The name of the scalable partition contains 32 or more characters.	Rename the scalable partition to contain 31 or fewer characters.

Deleting scalable objects

This section provides information about deleting scalable objects.

Deleting a scalable node

Complete the following steps to delete a scalable node:

1. If the scalable node is unassigned, go to step 2. Otherwise, remove the scalable node from the scalable partition; then, remove the scalable node from the scalable system. See "Removing scalable nodes" on page 55 and "Updating a scalable system" on page 66 for more information.
2. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
3. Right-click the unassigned scalable node; then, click **Delete**. A window opens prompting you to confirm the deletion.
4. Click **Yes, delete**.

Deleting a scalable partition

Notes:

1. If you create a scalable partition in Scalable Systems Configuration (instead of Scalable Systems Manager automatically creating it as described in "Interrogation for partition descriptors" on page 21) and then you delete this scalable partition, you cannot automatically restore the definition for it. Instead, you must create a new scalable partition.
2. If you delete a scalable partition that is powered-on, you will no longer be able to manage the scalable partition out-of-band with Scalable Systems Configuration. However, if a managed-system object represented the active scalable partition, you can continue to manage it in-band with IBM Director.
3. The scalable partition is deleted only from IBM Director; any corresponding partition descriptor in NVRAM is not deleted.

Complete the following steps to delete a scalable partition:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. Right-click the scalable partition that you want to delete; then, click **Delete**.

3. Use one of the following procedures, depending on whether the scalable partition contains scalable nodes:

- If the scalable partition contains scalable nodes, a window opens and a message warning you that the scalable partition cannot be managed out-of-band if you delete it is displayed. The message also indicates whether the scalable partition is currently powered-on.

  Click **Yes, delete**.

- If the scalable partition does not contain any scalable nodes, the scalable partition is deleted, and then a window opens and a message confirming the successful deletion is displayed.

  Click **OK**.

Other than causing loss of out-of-band management, deleting a powered-on scalable partition does not affect the operating system that is running on the scalable nodes that were in the scalable partition. It also does not affect the servers that are represented by the scalable nodes in the scalable partition.

Deleting a scalable system

Notes:

1. When you delete a scalable system, all scalable partitions in that scalable system are deleted. See "Deleting a scalable partition" on page 69 for information that you should consider when deleting scalable partitions.

2. When you delete a scalable system, all scalable nodes in that scalable system are unassigned. See "Unassigned scalable nodes" on page 22 for more information.

Complete the following steps to delete a scalable system:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.

2. Right-click the scalable system that you want to delete; then, click **Delete**. A window opens prompting you to confirm the deletion.

3. Click **Yes**.

Chapter 8. Viewing details about scalable objects in Scalable Systems Configuration

You can use Scalable Systems Configuration to display details about the following scalable objects:

- Scalable nodes
- Scalable systems
- Scalable partitions
- Remote I/O enclosures that are attached to scalable nodes

Displaying details about scalable nodes

You can use Scalable Systems Configuration to display information about scalable nodes as follows:

- For all scalable nodes that are defined in one scalable system
- For all unassigned nodes
- For one specific scalable node (either in a scalable system or unassigned)

Displaying information about all scalable nodes in one scalable system

Complete the following steps to display information about all scalable nodes in a scalable system:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Topology pane, expand the tree structure for a scalable system.
3. Click **Scalable Nodes**.

Scalable Systems Configuration updates the Details pane to display information about all scalable nodes in that scalable system. The Properties pane is not used when the word **Scalable Nodes** is selected in the Topology pane.
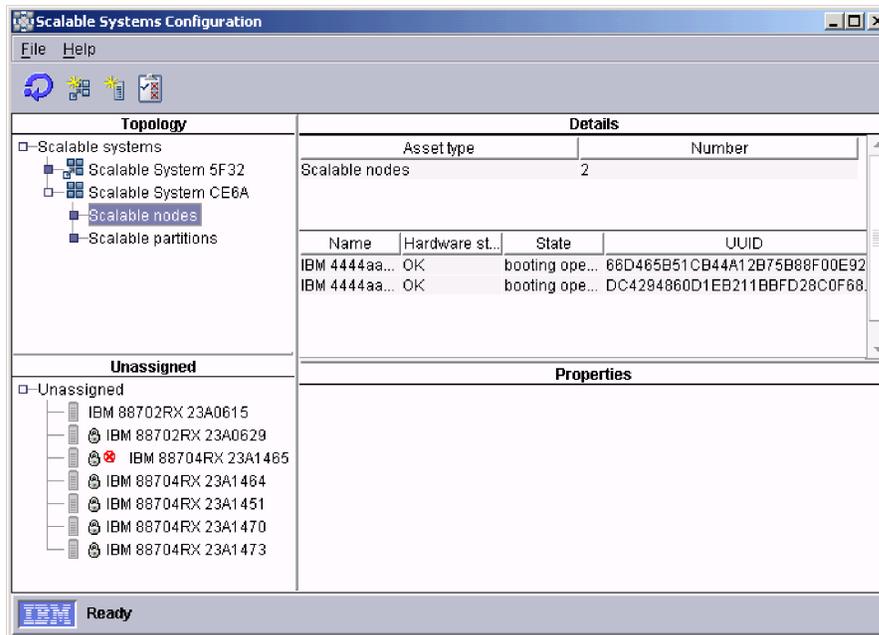
Figure 30. Displaying information about all scalable nodes in one scalable system

Scalable Systems Configuration displays the following information:

- The top table in the Details pane identifies the asset type (Scalable nodes in the example) and the number of scalable nodes in the scalable system.
- The bottom table in the Details pane provides information about each scalable node in the scalable system. For each scalable node, the table lists one row identifying its name, its hardware status, its state, and its UUID.

For more information, see "Hardware status of scalable nodes and remote I/O enclosures" on page 74 and "States of scalable nodes" on page 75.

Displaying information about all unassigned scalable nodes

Complete the following steps to display information about all unassigned scalable nodes:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Unassigned pane, click **Unassigned**.

Scalable Systems Configuration updates the Details pane to display information about all unassigned scalable nodes. The Properties pane is not used when the word **Unassigned** is selected in the Unassigned pane.
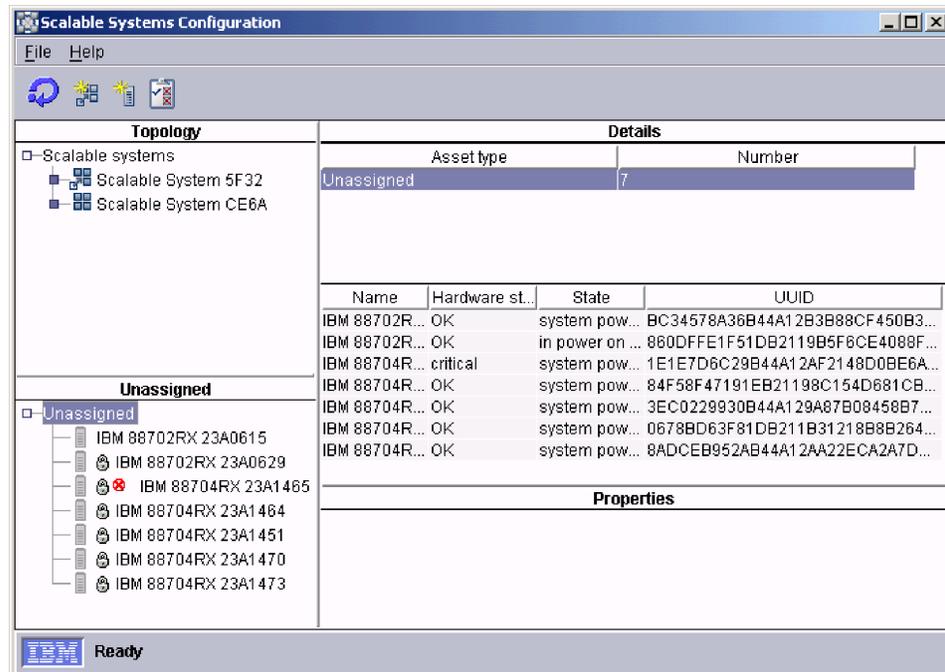
Figure 31. Displaying information about all unassigned scalable nodes

Scalable Systems Configuration displays the following information:

- The top table in the Details pane identifies the asset type (Unassigned in the example) and the number of unassigned scalable nodes.
- The bottom table in the Details pane provides information about each unassigned scalable node. For each scalable node, the table lists one row identifying its name, its hardware status, its state, and its UUID.

For more information, see "Hardware status of scalable nodes and remote I/O enclosures" on page 74 and "States of scalable nodes" on page 75.

Displaying information about one scalable node

Complete the following steps to display information about one scalable node, either in a scalable system or unassigned:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Topology pane, expand the tree structure until you can see the scalable node.
3. Click the scalable node.

Scalable Systems Configuration updates the Details and Properties panes to display information about the selected scalable node.
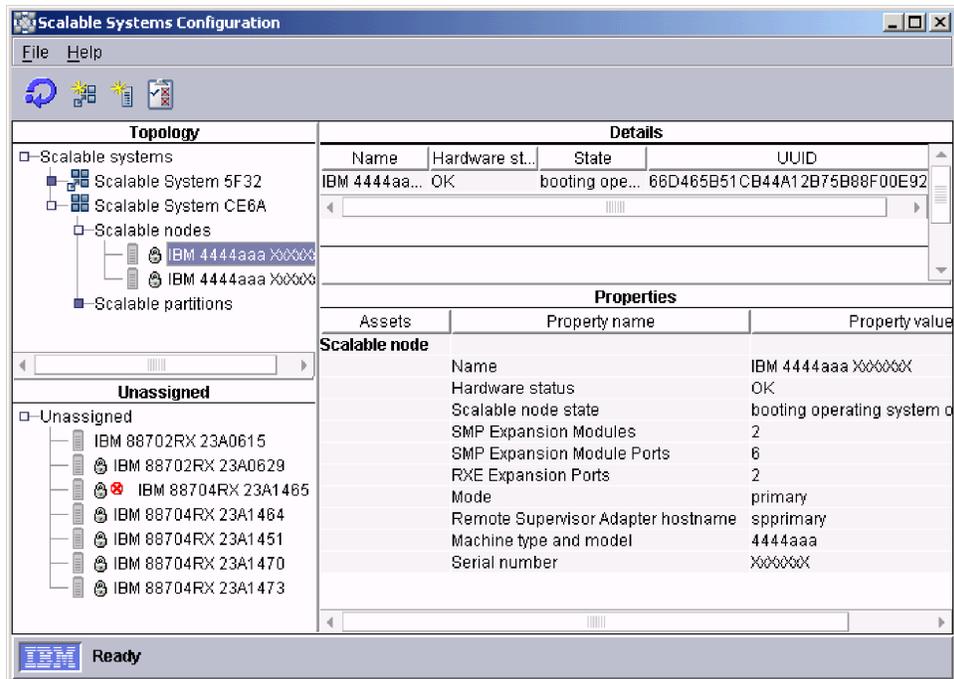
Figure 32. Displaying information about one scalable node

Scalable Systems Configuration displays the following information:

- The Details pane shows the name of the scalable node, its hardware status, its state, and its UUID.
- The Properties pane shows the following information about the scalable node:
  - The name of the scalable node
  - Its hardware status
  - Its state
  - Its number of SMP Expansion Modules
  - Its number of SMP Expansion Ports
  - Its number of RXE Expansion Ports
  - Its mode
  - The host name of its service processor
  - Its machine type and model
  - Its serial number

For more information, see "Hardware status of scalable nodes and remote I/O enclosures", "States of scalable nodes" on page 75, and "Modes of scalable nodes" on page 76.

Hardware status of scalable nodes and remote I/O enclosures

You can view the hardware-status value for a scalable node or remote I/O enclosure when you use Scalable Systems Configuration to display details about scalable nodes or display details about remote I/O enclosures. Scalable Systems Manager assigns hardware status to a scalable node or remote I/O enclosure object according to the status of the underlying hardware from the Hardware Status task. Scalable Systems Manager does not affect or control the status of this hardware other than reporting the status that is assigned by the Hardware Status task. Scalable Systems Manager displays icons beside the scalable node or remote I/O enclosure when its corresponding hardware has a Warning or Critical status.

Table 21 describes the hardware status values for scalable nodes and remote I/O enclosures.

Table 21. Hardware status values for scalable nodes and remote I/O enclosures

Hardware status value	Description
OK	The hardware that is represented by the scalable node or remote I/O enclosure object has no errors.
Warning	The hardware that is represented by the scalable node or remote I/O enclosure object has a Warning status. Run the Hardware Status task from IBM Director Console to examine the details of the Warning status for the affected device.
Critical	The hardware that is represented by the scalable node or remote I/O enclosure object has a Critical status. Run the Hardware Status task from IBM Director Console to examine the details of the Critical status for the affected device.
Unknown	The status of the hardware that is represented by the scalable node or remote I/O enclosure object could not be determined.

States of scalable nodes

When you use Scalable Systems Configuration to display details about scalable nodes, one of the details that is displayed is the current state of the scalable node. The state value of a scalable node reflects the state value of the server that is represented by the scalable node. State values are assigned to a server by its service processor. Scalable Systems Configuration does not affect or control the state of a server other than reporting the state that is assigned by the service processor.

Table 22 describes the various states that a scalable node can have.

Table 22. State values for scalable nodes

State value	Description
System powered off	The server that is represented by the scalable node is not powered-on, or the state is unknown.
System powered-on	The server that is represented by the scalable node is turned on, but power-on self-test (POST) has not started.
In power-on self-test	The server that is represented by the scalable node has begun POST, but it is not complete.
Power-on self-test error	An error was encountered during POST on the server that is represented by the scalable node.
Powered-on	POST has been completed, so the scalable node is powered-on, but there is no active operating system that is using the scalable node.
Booting operating system	The server that is represented by the scalable node is powered-on and has started the process of starting the operating system.
In operating system	The server that is represented by the scalable node is powered-on and is part of an active operating system. This state can occur only when the service-processor device driver has been installed on the operating system. Otherwise, the scalable node state remains as "booting operating system" even though the startup process might be complete.

Modes of scalable nodes

Scalable Systems Configuration displays the values of the mode attribute for scalable nodes. The mode value is set by Scalable Systems Configuration as the result of Management Processor Assistant (MPA) platform events that occur for scalable nodes. Table 23 describes the possible mode values for a scalable node.

Table 23. Mode values for scalable nodes

Mode value	Description
Null or unknown	The scalable node has not been set as the primary scalable node or a secondary scalable node. A scalable node is in this mode when it does not have a partition descriptor in NVRAM, it has not been configured into a scalable partition by Scalable Systems Configuration, and it has not been started as a stand-alone system.
Primary	The scalable node is the primary scalable node in a scalable partition.
Secondary	The scalable node is the secondary scalable node in the scalable partition.
Stand-alone	The scalable node is defined in a single-chassis scalable partition.
Stand-alone reset primary	The scalable node is the primary scalable node in a scalable partition, but instead of merging with any secondary scalable nodes in the scalable partition, it has started in stand-alone reset mode. This condition generally occurs when the Escape key is pressed on the keyboard of the primary scalable node while it is starting so that BIOS code can be set up or flashed. It can also happen when the primary scalable node could not start the operating system on the scalable partition because of an error. Possible errors include when the partition descriptor fails to timely merge, when the SMP Expansion Cables are defective, when the cables become detached from the server, or when other similar hardware problems occur. The stand-alone reset primary mode will automatically revert back to primary mode when the system is restarted. When the primary scalable node has entered stand-alone reset mode, the scalable partition containing that scalable node cannot start. When this scenario occurs, power-off the other scalable nodes in the scalable partition until the primary scalable node is available again.
Stand-alone reset secondary	The scalable node is a secondary scalable node in a scalable partition, but instead of merging with the primary scalable node, it has started in stand-alone reset mode. This condition generally occurs when the Escape key is pressed on the keyboard of the secondary scalable node while it is starting so that BIOS code can be set up or flashed. It can also happen when the primary scalable node could not start the operating system on the scalable partition because of an error. Possible errors include when the partition descriptor fails to timely merge, when the SMP Expansion Cables are defective, when the cables become detached from the server, or when other similar hardware problems occur. The stand-alone reset secondary mode will automatically revert back to secondary mode when the system is restarted. When a secondary scalable node has entered stand-alone reset mode, the scalable partition containing that scalable node cannot start. When this scenario occurs, power-off the other scalable nodes in the scalable partition until the secondary scalable node is available again.

Displaying details about scalable systems

You can use Scalable Systems Configuration to display information about all scalable systems on the management server or display information about one specific scalable system. You can display information for both manageable and view-only scalable systems.

In Scalable Systems Configuration, a scalable system is depicted in a tree structure in the Topology pane showing the name of the scalable system at the top of the hierarchy, followed by a listing of the scalable nodes and scalable partitions, also known as *members*, that are defined in that scalable system.

Displaying information about all scalable systems

Complete the following steps to display information about all scalable systems:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Topology pane, click **Scalable systems**.

Scalable Systems Configuration updates the Details pane to display information about all scalable systems. The Properties pane is not used when the word **Scalable systems** is selected in the Topology pane.
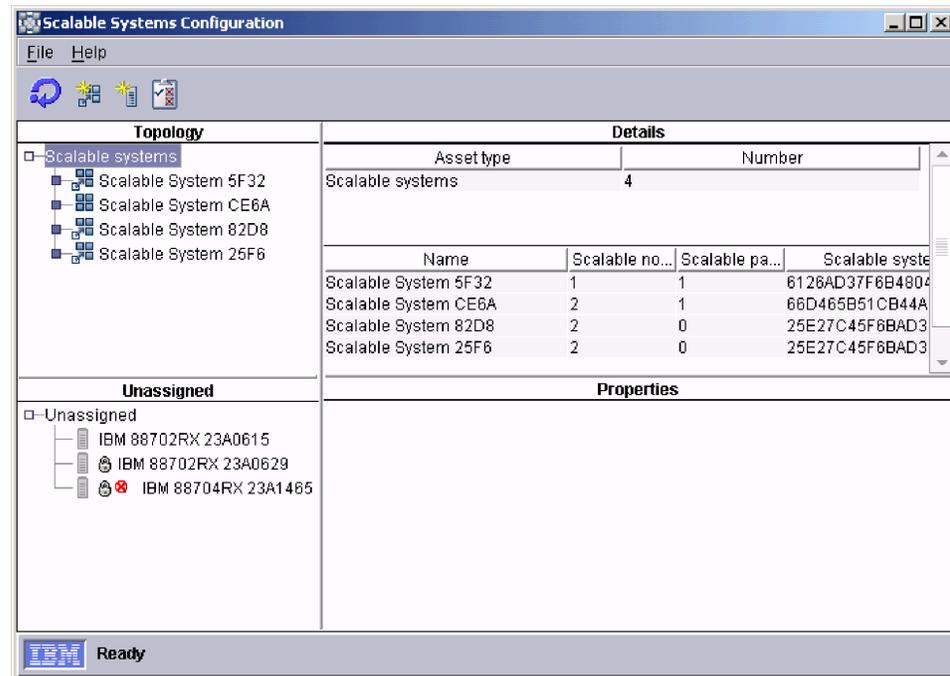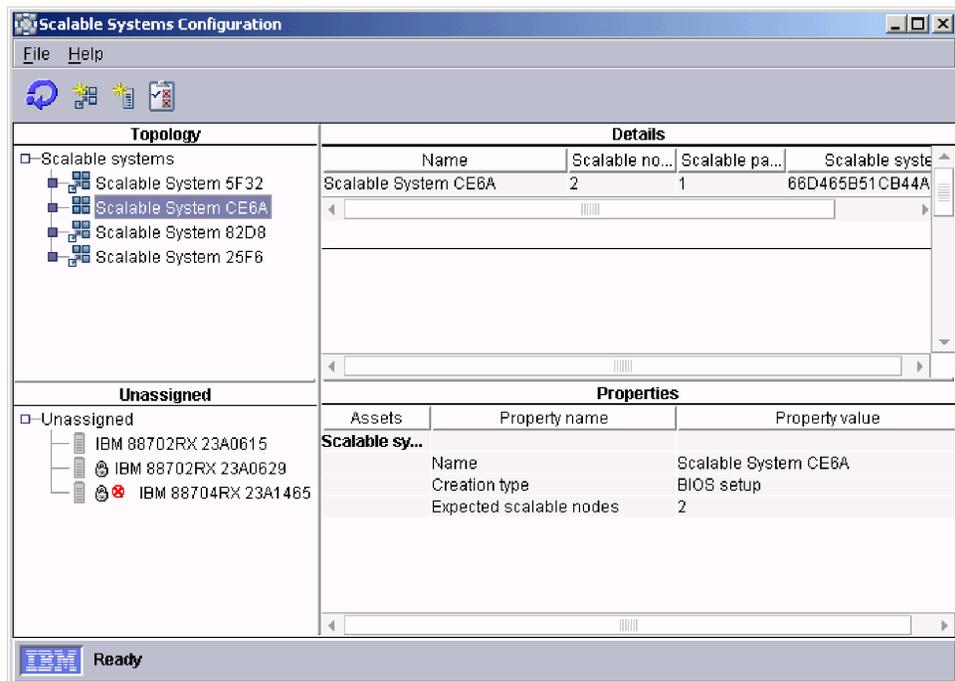

Figure 33. Displaying information about all scalable systems

Scalable Systems Configuration displays the following information:

- The top table in the Details pane identifies the asset type (Scalable systems in this example) and the number of scalable systems that are defined in IBM Director Server.

- The bottom table in the Details pane provides information about each scalable system that is defined on the server. For each scalable system, the table lists one row identifying its name, the number of scalable nodes in the scalable system, the number of scalable partitions in the scalable system, and the UUID of the scalable system.

Displaying information about one scalable system

Complete the following steps to display information about one scalable system:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.

2. In the Topology pane, expand the tree structure until you can see the scalable system that you want to display information about.

3. Click the scalable system.

Scalable Systems Configuration updates the Details and Properties panes to display information about the selected scalable system.

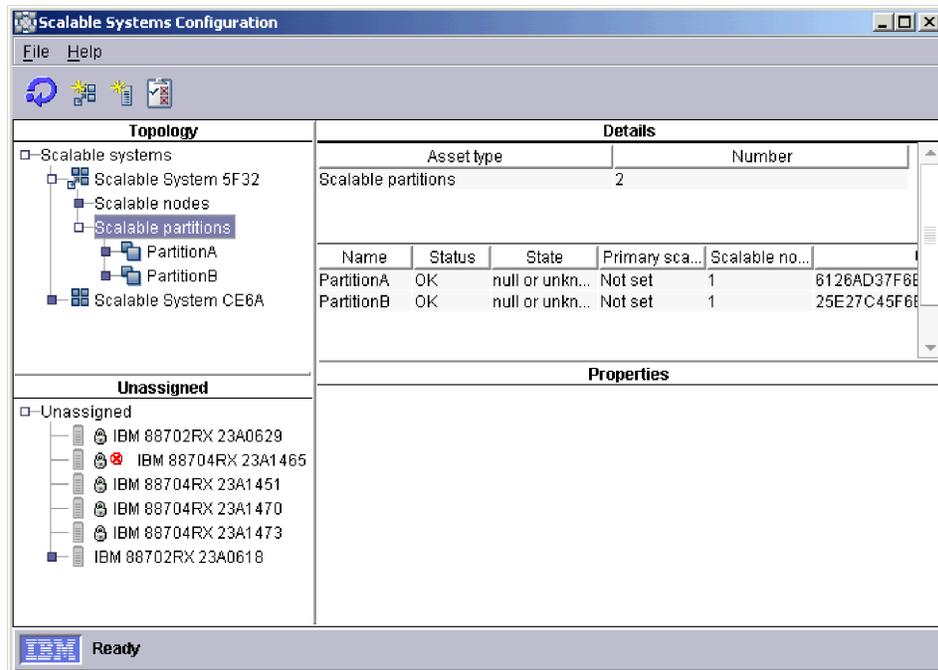*Figure 34. Displaying information about one scalable system*

Scalable Systems Configuration displays the following information:

- The Details pane identifies the name of the scalable system, the number of scalable nodes in the scalable system, the number of scalable partitions in the scalable system, and the UUID of the scalable system.

- The Properties pane displays the name of the scalable system, its creation type, and the number of expected scalable nodes in the scalable system.

  Note: The "Expected scalable nodes" property is used only for scalable systems that are created from BIOS-created partition descriptors. It is not applicable for scalable systems that were discovered by Scalable Systems Discovery or created in Scalable Systems Configuration.

Creation type of scalable systems

The creation-type property identifies the method that was used to create a scalable system. Scalable Systems Configuration displays this property and its value when you select one scalable system for which to display details. Possible values are shown in Table 24.

Table 24. Creation type values for scalable systems

Creation type value	Method used to create scalable system
Scalable systems configuration	The scalable system was created in Scalable Systems Configuration. This method creates manageable scalable systems.
Scalable systems discovery	The scalable system was discovered by the Scalable Systems Discovery task. This method creates manageable scalable systems.
BIOS setup	The scalable system was created from interrogation of the service processor NVRAM to locate a BIOS-created partition descriptor. This method can create manageable or view-only scalable systems.

Displaying details about scalable partitions

You can use Scalable Systems Configuration to display information about all scalable partitions that are defined for one scalable system or display information about one specific scalable partition in a scalable system.

Display information about all scalable partitions in a scalable system

Complete the following steps to display information about all scalable partitions in a scalable system:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Topology pane, expand the tree structure for the scalable system.
3. Click **Scalable partitions**.

Scalable Systems Configuration updates the Details pane to display information about all scalable partitions in that scalable system. The Properties pane is not used when the word **Scalable partitions** is selected in the Topology pane.
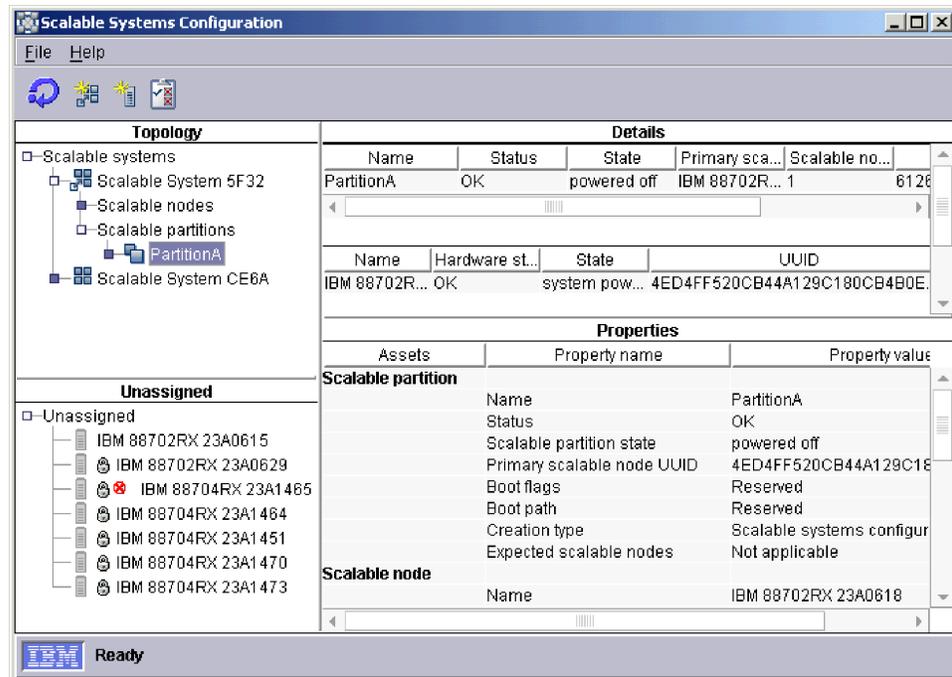
Figure 35. Displaying information about all scalable partitions in a scalable system

Scalable Systems Configuration displays the following information:

- The top table in the Details pane identifies the asset type (Scalable partitions in the example) and the number of scalable partitions in the scalable system.
- The bottom table in the Details pane provides information about each scalable partition in the scalable system. For each scalable partition, the table lists one row identifying the following information:
  - The name of the scalable partition
  - Its status
  - Its state
  - The name of the primary scalable node or "Not set" if none has been selected
  - The number of scalable nodes in the scalable partition
  - Its UUID

For more information, see "Hardware status of scalable nodes and remote I/O enclosures" on page 74 and "States of scalable partitions" on page 83.

Displaying information about one scalable partition

Complete the following steps to display details about one scalable partition:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.
2. In the Topology pane, expand the tree structure for the scalable system.
3. Click the scalable partition.

Scalable Systems Configuration updates the Details and Properties panes to display information about the selected scalable partition.

Figure 36. Displaying information about one scalable partition

Scalable Systems Configuration displays the following information:

Details pane

The Details pane displays information about the scalable partition and about each scalable node in the scalable partition:

- The top table in the Details pane provides information about each scalable partition in the scalable system. For each scalable partition, the table lists one row identifying the following information:
  - The name of the scalable partition
  - Its status
  - Its state
  - The name of the primary scalable node or "Not set" if none has been selected
  - The number of scalable nodes in the scalable partition
  - Its UUID
- The bottom table in the Details pane displays one row for each scalable node in the scalable partition. Each row identifies the following information:
  - The name of the scalable node
  - Its hardware status
  - Its state
  - Its UUID

Properties pane

The Properties pane displays information about the properties of the scalable partition and the properties of each scalable node in the scalable partition:

- Under the **Scalable partition** asset, information about the scalable partition itself is shown:
  - The name of the scalable partition
  - Its status
  - Its state
  - The UUID of the primary scalable node in the scalable partition
  - Its boot flags

– Its boot path
– Its creation type
– The number of expected scalable nodes

> **Note:** The "Expected scalable nodes" property is used only for scalable partitions that are created from BIOS-created partition descriptors. It is not applicable for scalable partitions that were created in Scalable Systems Configuration.

For more information, see "Status of scalable partitions" on page 82, "States of scalable partitions" on page 83 and "Creation type of scalable partitions" on page 83.

- Under each **Scalable node** asset, information about each scalable node in the scalable partition is shown:
  – The name of the scalable node
  – Its hardware status
  – Its state
  – Its number of SMP Expansion Modules
  – Its number of SMP Expansion Ports
  – Its number of RXE Expansion Ports
  – Its mode
  – The host name of its service processor
  – Its machine type and model
  – Its serial number

For more information, see "Hardware status of scalable nodes and remote I/O enclosures" on page 74, "States of scalable nodes" on page 75, and "Modes of scalable nodes" on page 76.

Status of scalable partitions

Scalable Systems Manager assigns status to a scalable partition when particular events occur for that scalable partition. The status value reflects the validity of the scalable partition and not the hardware status of the underlying servers that are represented by the scalable nodes in the scalable partition.

The status values for scalable partitions are shown in Table 25.

Table 25. Status values for scalable partitions

Status value	Description
OK	The scalable partition is valid and has not generated any error events.
Error	One of the following Error events occurred while initiating the power-on operation for the scalable partition: • A scalable node attempts to be the primary scalable node, but the scalable partition recognizes a different scalable node as primary. • A scalable node attempts to be a secondary scalable node, but the scalable partition recognizes a different scalable node as secondary. • The Escape key is pressed on the keyboard of the primary or secondary scalable node while it is starting so that BIOS code can be set up or flashed. • The primary scalable node could not start the operating system on the scalable partition, which means that the scalable node is now in stand-alone reset mode.
Unknown	The status of the scalable partition could not be determined.

States of scalable partitions

When you use Scalable Systems Configuration to display details about scalable partitions, one of the details that are displayed is the current state of the scalable partition. Similarly, IBM Director Console displays power icons beside the ▢ icon to indicate its state. Table 26 describes the various states that a scalable partition can have and the icons used by IBM Director to reflect these scalable partition states.

Table 26. State values for scalable partitions

State value	IBM Director Console icon	Description
Null or unknown	none	The state of the scalable partition could not be determined.
Powering-on	⬆	The scalable nodes in the scalable partition are powering-on, and after they are completed, the operating system will be started on the scalable partition.
Powered-on	⬆	The operating system is running on the scalable partition.
Powering-off	⬇	The operating system is being shut down on the scalable partition.
Powered-off	none	The operating system is not running on the scalable partition.
Resetting	⬍	The operating system is being reset on the scalable partition.

Creation type of scalable partitions

The creation-type property identifies the method that was used to create a scalable partition. Scalable Systems Configuration displays this property and its value when you select one scalable partition for which to display details. Possible values are shown in Table 27.

Table 27. Creation type values for scalable partitions

Creation type value	Method used to create scalable partition
Scalable systems configuration	The scalable partition was created in Scalable Systems Configuration. This method can create scalable partitions in manageable scalable systems.
BIOS setup	The scalable partition was created from interrogation of the service processor NVRAM to locate a BIOS-created partition descriptor. This method can create scalable partitions in manageable or view-only scalable systems.

Displaying details about remote I/O enclosures

Complete the following steps to display details about a remote I/O enclosure:

1. From IBM Director Console, in the Tasks pane, double-click the **Scalable Systems Configuration** task. The "Scalable Systems Configuration" window opens.

2. In the Topology pane or Unassigned pane, expand the tree structure for a scalable node.

3. Click the remote I/O enclosure.

Scalable Systems Configuration updates the Details and Properties panes to display information about the selected remote I/O enclosure and how it is cabled (attached) to the scalable node shown in the tree structure.
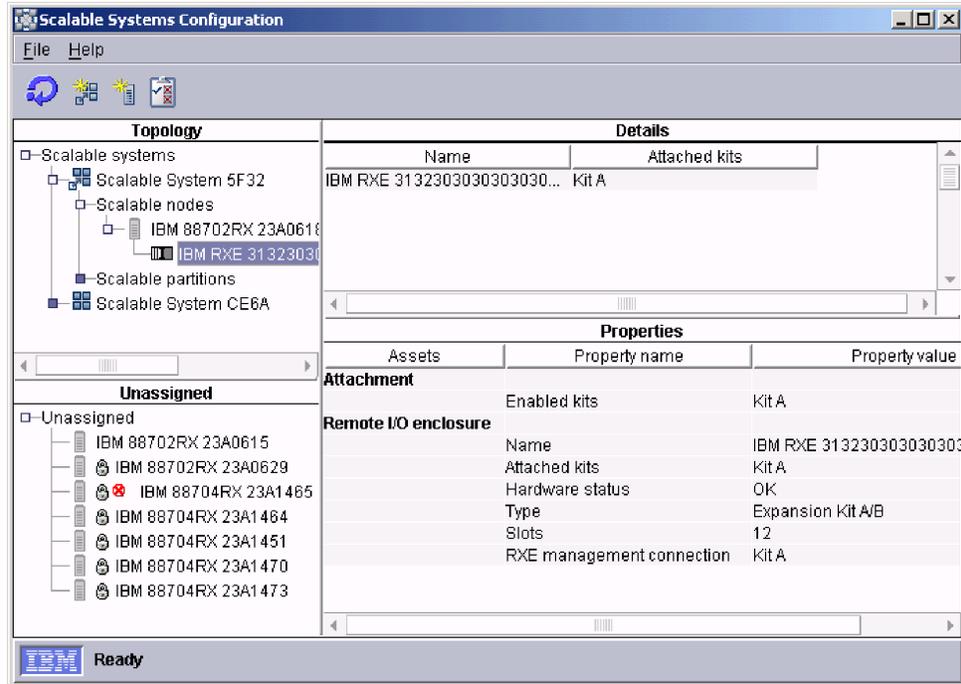

Figure 37. Displaying information about one remote I/O enclosure

The Details pane shows the name of the remote I/O enclosure and identifies the sides of the remote I/O enclosure that are cabled to the scalable node that is shown by the tree structure. Possible values are expansion kit A, B, or A/B.

The Properties pane shows information about how the remote I/O enclosure is attached to a scalable node and information about the remote I/O enclosure itself as follows:

- Under the **Attachment** asset, the value of the "Enabled kits" property identifies the expansion kits of the enclosure that are enabled for use by the scalable partition that contains the scalable node to which the enclosure is attached. Possible values are expansion kit A, B, or A/B. Use the RXE Configuration wizard to change the configuration of expansion kits for the remote I/O enclosure.
- Under the **Remote I/O enclosure** asset, details about the remote I/O enclosure itself are shown, such as the following properties:
  - Its name
  - Which expansion kits are physically cabled to its attached scalable nodes (expansion kit A, B, or A/B)
  - Its hardware status
  - The type, which means the type of expansion kits that are installed (expansion kit A, B, or A/B)
  - How many adapter slots it contains (6 or 12)
  - The RXE management connection being used by Scalable Systems Manager and IBM Director for communications (expansion kit A, B, or A/B)

For more information, see "Hardware status of scalable nodes and remote I/O enclosures" on page 74.

Chapter 9. Solving Scalable Systems Manager problems

The following table lists some of the problem symptoms and suggested solutions for Scalable Systems Manager.

Table 28. Solving Scalable Systems Manager problems

Symptom	Suggested action
Service processor communications	
Scalable Systems Manager cannot communicate with the service processors on the servers that are represented as scalable nodes in IBM Director.	Make sure that both of these conditions are met:
	• Service-processor networks either have static IP address assignments or have DHCP configured to maintain consistent IP addresses for the service processors. To do so, create reservations in DHCP that identify explicit IP addresses for the service processors.
	• The IP addresses that are assigned to the service processors do not change after the servers are discovered in IBM Director. This limitation is true whether you are using DHCP or statically assigning IP addresses.
	If the IP address of a service processor in a server that is used in a scalable partition changes, complete the following steps to reestablish out-of-band communications with the server:
	1. From IBM Director or Scalable Systems Configuration, remove the affected scalable objects:
	a. Delete the scalable system that contains the server.
	b. Delete any remote I/O enclosure that is attached to the scalable nodes that represent the server.
	c. Delete the scalable nodes that represent the server.
	2. From IBM Director Console, click **Tasks** → **Discover Systems** → **Physical Platforms** to rediscover the scalable nodes.
	3. From Scalable Systems Configuration, recreate the scalable system and scalable partition. For details, see "Creating scalable systems" on page 50 and "Creating scalable partitions" on page 52.
	4. If the scalable node is not already unlocked, from the IBM Director Group Contents pane, right-click the scalable node and click **Request Access**.
	From IBM Director Console, start one of the MPA subtasks against the physical-platform managed object. Verify that MPA is able to establish communications with the service processor. If it is not, use MPA to configure out-of-band communications with the service processor. For more information about communicating with service processors in an IBM Director environment, see the *IBM Director 4.11 Installation and Configuration Guide*.

Table 28. Solving Scalable Systems Manager problems (continued)

Symptom	Suggested action

A scalable partition cannot be powered-on or powered-off in Scalable Systems Configuration.	Complete the following steps:
	
	1. If you have not done so already, write down the IBM Director system names of the scalable nodes in the affected scalable partition. This information is shown in the Group Contents pane of IBM Director Console.
	
	2. From IBM Director or Scalable Systems Configuration, remove the affected scalable objects:
	  • Delete the scalable system that contains the scalable partition. (This will delete the scalable partition as well.)
	  • Delete any remote I/O enclosure that is attached to the scalable nodes that were in the scalable partition.
	  • Delete the scalable nodes that were part of the scalable partition.
	
	3. From the Web interface to the service processor on the primary server, use the Restore Defaults page to reset the service-processor configuration to the factory defaults.
	Attention: When you click Restore Defaults, you will lose all the modifications that you made to the service processor. You also will lose the remote control of the remote servers. This includes the IP address, logins, and passwords. If you had site-specific customizations, you must reconfigure them after the defaults have been restored.
	
	4. If the affected servers are part of a service-processor network with one or more DHCP servers, make sure that DHCP is configured to maintain consistent IP addresses for the service processors.
	
	5. From IBM Director Console, click **Tasks** → **Discover Systems** → **Physical Platforms** to rediscover the scalable nodes.
	
	6. From Scalable Systems Configuration, recreate the scalable system and scalable partition. For details, see "Creating scalable systems" on page 50 and "Creating scalable partitions" on page 52.
	
	7. If the scalable node is not already unlocked, from the IBM Director Group Contents pane, right-click the scalable node and click **Request Access**.

Table 28. Solving Scalable Systems Manager problems (continued)

Symptom	Suggested action

A scalable node that was previously unlocked becomes locked again, or a scalable node cannot be unlocked at all.	Make sure that only one management server is being used to manage scalable objects with Scalable Systems Manager. Using multiple management servers to manage scalable objects causes unpredictable results, such as being unable to power-on a scalable partition when the scalable node has suddenly become locked.
	
	Sometimes a scalable node cannot be unlocked because the IP address of the server that it represents has changed since it was discovered by IBM Director. Alternatively, the scalable node just needs to be rediscovered by IBM Director.
	
	If needed, complete the following steps to reestablish out-of-band communications with the server:
	1. From IBM Director or Scalable Systems Configuration, remove the affected scalable objects:
	   a. Delete the scalable system that contains the server.
	   b. Delete any remote I/O enclosure that is attached to the scalable nodes that represents the server.
	   c. Delete the scalable nodes that represent the server.
	2. Unplug all cables to the server and unplug it.
	3. Recable the server and plug in its power cable.
	4. From IBM Director Console, click **Tasks** → **Discover Systems** → **Physical Platforms** to rediscover the scalable nodes.
	5. From Scalable Systems Configuration, recreate the scalable system and scalable partition. For details, see "Creating scalable systems" on page 50 and "Creating scalable partitions" on page 52.
	6. If the scalable node is not already unlocked, from the IBM Director Group Contents pane, right-click the scalable node and click **Request Access**.

After changing the configuration of a remote I/O enclosure, Scalable Systems Configuration displays incorrect information about its expansion kit configuration.	Complete the following steps:
	1. From IBM Director or Scalable Systems Configuration, remove the following scalable objects:
	   a. Delete the scalable system that contains the remote I/O enclosure. (This will delete the scalable partition that contains the remote I/O enclosure.)
	   b. Delete the remote I/O enclosure that reflects the incorrect configuration.
	   c. Delete the scalable nodes that represent the server to which the remote I/O enclosure was attached.
	2. Unplug all cables between the remote I/O enclosure and the server and unplug their power cables.
	3. Recable the remote I/O enclosure to the server and plug in their power cables.
	4. From IBM Director Console, click **Tasks** → **Discover Systems** → **Physical Platforms** to rediscover the scalable node.
	5. From Scalable Systems Configuration, recreate the scalable system and scalable partition. For details, see "Creating scalable systems" on page 50 and "Creating scalable partitions" on page 52.
	6. If the scalable node is not already unlocked, from the IBM Director Group Contents pane, right-click the scalable node and click **Request Access**.

Ping responses from discovery and validation	

Table 28. Solving Scalable Systems Manager problems (continued)

Symptom	Suggested action
The results from Scalable Systems Discovery or Scalable Systems Validation are not what was expected.	Use the Management Processor Configuration subtask of MPA to make sure that an alert-forwarding profile is configured for the servers that are represented by the scalable nodes that you targeted. In the alert-forwarding profile, the connection type must be set to "IBM Director Comprehensive" and the IP address must be set to the IP address of the one management server that is being used to manage scalable objects with Scalable Systems Manager. Otherwise, the ping responses will not occur and neither Scalable Systems Discovery or Scalable Systems Validation will work properly.
Event action plans	
An event action in an event action plan does not occur when an event occurs.	Make sure that all of these conditions are met: • The event action plan is targeting the proper managed object. That is, if an event is for a physical-platform managed object, the event action plan must target a physical-platform managed object and not a scalable partition. • Only one management server is used to manage scalable objects in a multi-server IBM Director environment. • Each server must be configured with an alert-forwarding profile that has a connection type of "IBM Director Comprehensive" events. For more details, see "Event filters and actions for use with scalable objects" on page 8.
Automatic refreshes	
In the "Scalable Systems Configuration" window, the power state of a scalable partition is not automatically updated.	Make sure that both of these conditions are met: • Only one management server is used to manage scalable objects in a multi-server IBM Director environment. • Each server must be configured with an alert-forwarding profile that has a connection type of "IBM Director Comprehensive" events. For more details, see "The Scalable Systems Configuration window" on page 44.

Appendix A. Error codes

This appendix describes the error codes that Scalable Systems Manager can return through error messages displayed in Scalable Systems Configuration.

Additionally, you can see these error codes when you use the IBM Director Event Log task to view details on Scalable Systems Manager event types that have been received and logged by IBM Director Server.

Power-on failed events

Power-on failed events occur when Scalable Systems Manager is unable to power-on a scalable partition. When a power-on failed event occurs, Scalable Systems Manager returns the following message through Scalable Systems Configuration:

```
Scalable partition "name" power on failed with error code: number
```

where

- *name* is the name of the scalable partition that received the power-on failed event
- *number* is the error code returned with the event

Table 29 describes the error codes that can be returned in power-on failed messages.

Table 29. Error codes returned by power-on failed events

Error code	Description
8	The primary scalable node is locked. This error can also occur when Scalable Systems Manager cannot locate the primary scalable node or when the management server has an invalid primary node.
14	The scalable partition failed to power-on for an unknown reason.
16	None of the scalable nodes in the scalable partition have been set as the primary scalable node.
19	Scalable Systems Manager could not write the partition descriptor for the scalable partition to NVRAM on the service processor of the primary scalable node. This error can also occur when the service processor on the primary scalable node could not communicate with the secondary scalable nodes in the scalable partition.
22	The primary scalable node in the scalable partition is locked and cannot be accessed.
23	In a scalable partition generated from a BIOS-created partition descriptor, the number of expected nodes as read from the partition descriptor does not match the number of actual nodes in the scalable partition.
24	Scalable Systems Manager could not activate the partition descriptor for the scalable partition in NVRAM on the service processor of the primary scalable node.
30	Scalable Systems Manager located the primary scalable node, but it cannot communicate with its service processor.
34	One of the scalable nodes in the scalable partition is locked and cannot be accessed.

Table 29. Error codes returned by power-on failed events (continued)

Error code	Description
35	The scalable partition does not contain any scalable nodes.
36	A scalable node in the scalable partition cannot be used in a manageable scalable system.
1000	Scalable Systems Manager failed due to an unforeseen problem on the management server.

Assign failed events

Assign failed events occur when Scalable Systems Manager is unable to write a partition descriptor to a service processor. When an assign failed event occurs, Scalable Systems Manager returns the following message through Scalable Systems Configuration:

```
Scalable partition "name" assign failed with error code: number
```

where

- *name* is the name of the scalable partition that received the assign failed event
- *number* is the error code returned with the event

Table 30 describes the error codes that can be returned in assign failed messages.

Table 30. Error codes returned by assign failed events

Number	Event description
9	Scalable Systems Manager could not write the partition descriptor for the scalable partition to NVRAM on the service processor of the primary scalable node. This error can also occur when the service processor on the primary scalable node could not communicate with the secondary scalable nodes in the scalable partition.
16	None of the scalable nodes in the scalable partition have been set as the primary scalable node.
23	In a BIOS-created partition descriptor, the number of expected nodes in the partition descriptor does not match the number of expected nodes detected by Scalable Systems Manager.
34	One of the scalable nodes in the scalable partition is locked and cannot be accessed.
35	The scalable partition does not contain any scalable nodes.
36	A scalable node in the scalable partition cannot be used in a manageable scalable system.
1000	Scalable Systems Manager failed due to an unforeseen problem on the management server.

Shut down/power-off failed events

Shut down/power-off failed events occur when Scalable Systems Manager is unable to shut down and power-off a scalable partition. When a shut down/power-off failed event occurs, Scalable Systems Manager returns the following message through Scalable Systems Configuration:

```
Scalable partition "name" shut down/power off failed with error code: number
```

where
- *name* is the name of the scalable partition that received the shut down/power off failed event
- *number* is the error code returned with the event

Table 31 describes the error codes that can be returned in shut down/power-off failed events.

Table 31. Error codes returned by shut down/power-off failed events

Number	Event description
8	The primary scalable node is locked. This error can also occur when Scalable Systems Manager cannot locate the primary scalable node or when the management server has an invalid primary node.
15	A scalable partition failed to shut down and power-off.
30	Scalable Systems Manager cannot communicate with the service processor of the server that is represented by the primary scalable node.

Power-off failed events

Power-off failed events occur when Scalable Systems Manager is unable to power-off a scalable partition. When a power-off failed event occurs, Scalable Systems Manager returns the following message through Scalable Systems Configuration:

```
Scalable partition "name" power off failed with error code: number
```

where
- *name* is the name of the scalable partition that received the power-off failed event
- *number* is the error code returned with the event

Table 32 describes the error codes that can be returned in power-off failed events.

Table 32. Error codes returned by power-off failed events

Number	Event description
8	The primary scalable node is locked. This error can also occur when Scalable Systems Manager cannot locate the primary scalable node or when the management server has an invalid primary node.
17	A scalable partition failed to immediately power-off.
30	Scalable Systems Manager cannot communicate with the service processor of the server that is represented by the primary scalable node.

Scalable partition error events

Scalable partition error events occur when Scalable Systems Manager has errors with the primary or secondary scalable nodes in a scalable partition. When a scalable partition error event occurs, Scalable Systems Manager returns the following message through Scalable Systems Configuration:

```
Scalable partition "name" error occurred, error code: number
```

where
- *name* is the name of the scalable partition that received the error event

- *number* is the error code returned with the event

Table 33 describes the error codes that can be returned in scalable partition error events.

Table 33. Error codes returned by scalable partition error events

Number	Event description
39	A scalable node attempted to be the primary scalable node, but its scalable partition recognized a different scalable node as primary.
40	A scalable node attempted to be a secondary scalable node, but its scalable partition recognized a different scalable node as secondary.
41	The primary scalable node in a scalable partition started in stand-alone reset mode. This condition generally occurs when the Escape key is pressed on the keyboard of the primary scalable node while it is starting so that BIOS code can be set up or flashed. It can also happen when the primary scalable node could not start the operating system on the scalable partition because of an error. For more information, see "Modes of scalable nodes" on page 76.
42	The secondary scalable node in a scalable partition started in stand-alone reset mode. This condition generally occurs when the Escape key is pressed on the keyboard of the secondary scalable node while it is starting so that BIOS code can be set up or flashed. It can also happen when the primary scalable node could not start the operating system on the scalable partition because of an error. For more information, see "Modes of scalable nodes" on page 76.

Appendix B. Terminology summary and abbreviation list

This appendix provides a summary of Scalable Systems Manager terminology and a list of abbreviations that are used in this publication.

Scalable Systems Manager terminology summary

The hardware in an IBM Director environment is referred to in the following ways:

- A *system* is a server, workstation, desktop computer, or mobile computer.
- A *management server* is a server on which IBM Director Server is installed.
- A *managed system* is a system on which IBM Director Agent is installed.
- A *management console* is a system on which IBM Director Console and Scalable Systems Manager is installed.

In IBM Director, the following software concepts are applicable to Scalable Systems Manager:

- A *managed object* is an item that is managed by IBM Director such as a multi-node server, scalable partition, physical platform, scalable node, or remote I/O enclosure. 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).
- A *physical platform* is an IBM Director managed object that represents a remote system that is discovered out-of-band by IBM Director Server. The remote system is discovered through the use of the service location protocol (SLP) and the service processor on the remote system.
- A *scalable node* is a physical platform that has at least one SMP Expansion Module. At the time of publication, the xSeries 440 and xSeries 445 are the only server models that contain chassis that can be scalable nodes. 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.
- A *scalable system* is an IBM Director managed object that consists of scalable nodes and the scalable partitions that are made from the scalable nodes in the scalable system. When a scalable system contains two scalable nodes, the servers that they represent must be interconnected through their SMP Expansion Modules to make a 16-way configuration, for example, a 16-way xSeries 445 server. Scalable nodes that represent xSeries 445 servers in supported configurations are used in manageable scalable systems. Scalable nodes that represent xSeries 440 servers in a 16-way configuration are used in view-only scalable systems.
- A *scalable partition* is an IBM Director managed object that defines the scalable nodes that can run a single image of the operating system. A scalable partition is the logical equivalent of a physical platform: it can be powered-on and powered-off through IBM Director Console. IBM Director manages a scalable partition through the service processor on the primary scalable node of that scalable partition. Scalable partitions are associated with scalable systems and can contain only the scalable nodes from their associated scalable system.
- A *remote I/O enclosure* is an IBM Director managed object representing an expansion enclosure of PCI-X slots, for example, an RXE-100 Remote Expansion Enclosure. The enclosure consists of one or two expansion kits. Each expansion kit contains six hot-swap Active™ PCI-X adapter slots.

Abbreviation list

The following table lists abbreviations that are used in the Scalable Systems Manager publication.

Abbreviation	Definition
BIOS	basic input/output system
DHCP	Dynamic Host Configuration Protocol
GUI	graphical user interface
I/O	input/output
MPA	Management Processor Assistant
NVRAM	nonvolatile random-access memory
PDF	Portable Document Format
POST	power-on self-test
SLP	service location protocol
SMP	symmetrical multiprocessing
UUID	universal unique identifier

Appendix C. Getting help and technical assistance

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 appendix 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 (see Chapter 9, "Solving Scalable Systems Manager problems", on page 87), and use the diagnostic tools that come with your system.
- Go to the IBM Support Web site at http://www.ibm.com/pc/support/ to check for technical information, hints, tips, and new device drivers.
- Use an IBM discussion forum on the IBM Web site to ask questions.

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 comes 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 http://www.ibm.com/pc/support/ and follow the instructions. Also, you can order publications through the IBM Publications Ordering System at
http://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).

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services or through your IBM reseller, if your reseller is authorized by IBM to provide warranty service. Go to http://www.ibm.com/planetwide/ for support telephone numbers, or in the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

Appendix D. 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 give 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.

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.

Edition notice

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

e-business logo	Netfinity
@server	Redbooks
IBM	ServerProven
Intellistation	xSeries

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

A

association.   A way of displaying the members of a group in a logical ordering or a way to display additional information about the members of the group. For example, the Object Type association displays the managed objects in a group in folders based on their type. In the second case, 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 personal computer code that controls basic hardware operations, such as interactions with diskette drives, hard disk drives, and the keyboard. The Configuration/Setup Utility program is a menu-driven utility that is part of the BIOS code that comes with a server. You can start it with F1 during a specific point in the server startup (by watching the screen for a message about it).

BIOS.   See basic input/output system.

E

event.   A flag that identifies a change in a system process or a device, such that notification of that change can be generated and tracked, for example, notification that a managed system is offline.

event action.   The action that IBM Director takes in response to a specific event or events. In the Event Action Plan Builder, you customize an event action type by specifying certain parameters and saving the event action. You must assign the customized event action (and an event filter) to an event action plan before IBM Director can execute the event action.

event action plan.   A user-defined plan that determines how IBM Director will manage certain events. An event action plan consists of one or more event filters and one or more customized event actions. The event filters specify which events are managed, and the event actions specify what happens when the events occur.

Event Action Plan wizard.   An IBM Director Console wizard that can be used to create simple event action plans.

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 that the filter is assigned to.

extension.   See IBM Director extension.

G

group.   A logical set of managed objects. Groups can be dynamic, static, or task-based.

I

IBM Director Console.   A component of IBM Director software. When installed on a system, it provides a graphical user interface (GUI) and enables network administrators to access IBM Director Server. IBM Director Console transfers data to and from the management server using TCP/IP.

IBM Director extension.   A tool that extends the functionality of IBM Director. IBM Director extensions include the IBM Server Plus Pack, Remote Deployment Manager, Software Distribution, and others.

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.

in-band communication.   Communication that occurs through the same channels as data transmissions, for example, the interprocess communication that occurs between IBM Director Server, IBM Director Agent, and IBM Director Console.

M

managed group.   A group of systems or objects managed by IBM Director.

managed object.   An item managed by IBM Director. Managed objects include managed systems, Windows NT clusters, BladeCenter chassis, management processors, SNMP devices, multi-node servers (scalable systems), scalable partitions, physical platforms, scalable nodes, and remote I/O enclosures. 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 system.   A system (server, desktop computer, workstation, or mobile computer) on which IBM Director Agent is installed. Such a system is managed by IBM Director.

management console. A system (server, desktop computer, workstation, or mobile computer) on which IBM Director Console is installed.

Management Processor Assistant (MPA). An IBM Director task that can be used to configure, monitor, and manage service processors installed in Netfinity® and xSeries servers.

management server. The server on which IBM Director Server is installed.

MPA. See Management Processor Assistant.

N

nonvolatile random-access memory (NVRAM). Random access memory (storage) that retains its contents after the electrical power to the computer is shut off.

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 the LAN. In an IBM Director environment, such communication is independent of both the operating system and interprocess communication (IPC). In Scalable Systems Manager, out-of-band communication occurs exclusively between IBM Director Server and the service processor on supported servers.

P

partition. See scalable partition.

partition descriptor.

A small data structure located in NVRAM that defines a scalable system and the number of scalable nodes in it and defines any scalable partitions created from the scalable nodes. For multi-chassis servers, the partition descriptor defines one server as the primary server and the other as secondary.

The partition descriptor is read by BIOS code to start the scalable partition created from the scalable nodes. For some servers, the Configuration/Setup Utility program (in BIOS code) provides a default partition descriptor. However, it is recommended that you use Scalable Systems Manager to create and modify partition descriptors.

physical platform. (1) An IBM Director managed object that represents a remote system that is discovered out-of-band by IBM Director Server. The remote system is discovered through the use of the service location protocol (SLP) and the service

processor on the remote system. At the time of publication, the only server models whose chassis can be discovered as physical platforms in this manner are the xSeries 360, xSeries 440, and xSeries 445. A physical platform enables identification of some systems without communicating through the operating system or any IBM Director Agent that has been installed on that system. Because IBM Director Agent is not used to provide the support for physical platforms, only limited functionality exists. (2) An IBM Director managed object representing a system that has IBM Director Agent and the Management Processor Assistant (MPA) agent installed.

R

remote I/O enclosure. An IBM Director managed object representing an enclosure of PCI-X slots, for example, an RXE-100 Remote Expansion Enclosure. The enclosure consists of one or two expansion kits. Each expansion kit contains six hot-swap Active™ PCI-X adapter slots.

Remote Supervisor Adapter. A system adapter that provides advanced remote systems management for xSeries servers. This adapter also provides remote access to servers independent of the server status, remote control of the hardware and the operating system and Web-based management capabilities using standard Web browsers.

RXE Expansion Port. The dedicated high-speed port used to connect a remote I/O enclosure, for example, the RXE-100 Remote Expansion Enclosure, to a server such as the xSeries 445 server.

S

scalable node. A physical platform that has at least one SMP Expansion Module. At the time of publication, the xSeries 440 and xSeries 445 are the only server models whose chassis can be scalable nodes. 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: it can be powered-on and powered-off through IBM Director Console. IBM

Director manages a scalable partition through the service processor on the primary scalable node of that scalable partition. Scalable partitions are associated with scalable systems and comprise only the scalable nodes from their associated scalable systems. At the time of publication, only the xSeries 445 servers can be used to create scalable partitions in IBM Director.

scalable system. An IBM Director managed object that consists of scalable nodes and the scalable partitions that are made from the scalable nodes in the scalable system. When a scalable system contains two scalable nodes, the servers that they represent must be interconnected through their SMP Expansion Modules to make a 16-way configuration, for example, a 16-way xSeries 445 server. Scalable nodes that represent xSeries 445 servers in supported configurations are used in manageable scalable systems. Scalable nodes that represent xSeries 440 servers in a 16-way configuration are used in view-only scalable systems.

service location protocol (SLP). A protocol developed by the Internet Engineering Task Force (IETF) to discover the location of services on a network automatically. It is used by IBM Director Server to discover BladeCenter chassis and multi-node servers such as the xSeries 440 and 445.

service processor. A hardware-based management processor used in IBM Netfinity and xSeries servers. A generic term for Remote Supervisor Adapters, Advanced System Management processors, Advanced System Management PCI adapters, and integrated system management processors.

SLP. See service location protocol.

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. The IBM xSeries 440 is the first hardware platform that uses SMP Expansion Modules.

SMP Expansion Port. A dedicated high-speed port used to interconnect SMP Expansion Modules.

SMP Expansion Cable. The cable used to connect two SMP Expansion Ports.

static partition. A view-only scalable partition that consists of xSeries 440 servers.

U

universal unique identifier (UUID). A 128-bit character string guaranteed to be globally unique and used to identify components under management. The UUID enables inventory-level functionality and event tracking of scalable nodes, scalable partitions, scalable systems, and remote I/O enclosures.

UUID. See universal unique identifier.

V

vital product data (VPD). The key information about a server, its components, POST/BIOS, and service processor. This includes machine type, model and serial number, component FRU number, serial number, manufacturer ID, and slot number; POST/BIOS version number, build level, and build date; and service processor build ID, revision numbers, file name, and release date.

VPD. See vital product data.

Index

Numerics

16-way configurations
    supported servers 24, 27
4-way configurations
    supported servers 26
8-way configurations
    supported servers 27

A

administration, user 7
alert-forwarding profiles 90
Assign event 9
assign failed events 92
associations 7

B

BIOS
    interrogation 21
    updates by Scalable Systems Manager 29

C

configurations
    default remote I/O enclosures 59
        16-way servers 62
        4-way servers 60
        8-way servers 61
    supported servers 25
        16-way 27
        4-way 26
        8-way 27
configurations, supported servers
    16-way 24
    graphic notations 25
creation type property
    scalable partitions 83
    scalable systems 79
critical hardware status 75
customer support xii

D

Details pane 46
DHCP, configuration limitations 87
discovery
    physical platform 19
    Scalable Systems Discovery 35
disk space requirements 13

E

eFixes xii
Enabled kits property 84
error codes 91

error codes *(continued)*
    assign failed events 92
    power-off failed events 93
    power-on failed events 91
    scalable partition error events 93
    shut down/power-off failed events 92
Error event 9
error status 82
event actions 10
events
    MPA events for scalable objects 8
    Scalable Systems Manager 8
expansion enclosures
    See remote I/O enclosures

G

graphic notations
    default remote I/O enclosure configurations 59
    supported server configurations 25
groups 6

H

hardware requirements
    management servers and consoles 13
    out-of-band systems 13
hardware status 74
    icons 45
help
    IBM Director Web sites xii
    Scalable Systems Manager 46

I

IBM Director
    associations for scalable objects 7
    groups for scalable objects 6
    icons
        hardware status 45
        scalable objects 5
        scalable partition states 83
        toolbar 47
    interrogation of BIOS 21
    managed systems for scalable partitions 30
    physical platform access 20
    physical platform discovery 19
    publications xii
    scalable object information 32
        restoring 32
        saving 32
    user administration for scalable objects 7
    Web sites xii
icons
    hardware status 45
    scalable objects 5
    scalable partition states 83

IBM ®

Part Number: 01R0540

Printed in U.S.A.