Parallel System Support Programs for AIX

Installation and Migration Guide

Version 3 Release 4

GA22-7347-03

Parallel System Support Programs for AIX

Installation and Migration Guide

Version 3 Release 4

GA22-7347-03

Note!
Before using this information and the product it supports, be sure to read the general information under “Notices” on page 305.

Fourth Edition (December 2001)

This edition applies to version 3 release 4 of the IBM Parallel System Support Programs for AIX (PSSP) Licensed Program (product number 5765-D51) and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces GA22-7347-02. Significant changes or additions to the text and illustrations are indicated by a vertical line (|) to the left of the change.

IBM welcomes your comments. A form for readers' comments may be provided at the back of this publication, or you may address your comments to the following address:

International Business Machines Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, NY 12601-5400
United States of America

FAX (United States & Canada): 1+845+432-9405
FAX (Other Countries):
Your International Access Code +1+845+432-9405

IBMLink (United States customers only): IBMUSM10(MHVRCFS)
Internet e-mail: mhvrcfs@us.ibm.com

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:

- Title and order number of this book
- Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Tables

About this book

This book contains information to help you install, configure, maintain, and migrate the RS/6000 SP. It includes concepts and instructions pertaining to:

- Installing and configuring the IBM Parallel System Support Programs for AIX
- Installing on an existing RS/6000 SP system
- Reconfiguring the system
- Performing system maintenance
- Migrating the system

For a list of related books and information about accessing online information, see the Bibliography in the back of the book.

This book applies to PSSP Version 3 Release 4. To find out what version of PSSP is running on your control workstation (node 0), enter the following:

```
splst_versions -t -n0
```

In response, the system displays something similar to:

```
0 PSSP-3.4
```

If the response indicates **PSSP-3.4**, this book applies to the version of PSSP that is running on your system.

To find out what version of PSSP is running on the nodes of your system, enter the following from your control workstation:

```
splst_versions -t -G
```

In response, the system displays something similar to:

```
1 PSSP-3.4
2 PSSP-3.2
7 PSSP-3.1.1
8 PSSP-2.4
```

If the response for a particular node indicates **PSSP-3.4**, this book applies to the version of PSSP that is running on that node.

If you are running mixed levels of PSSP, be sure to maintain and refer to the appropriate documentation for whatever versions of PSSP you are running.

Who should use this book

This book is intended for system administrators responsible for installing, configuring, and maintaining the RS/6000 SP system. It assumes the administrators have a working knowledge of AIX or UNIX and experience with network systems.

The System Administrators Guild of USENIX (SAGE), has developed a classification for skills required for system administrators. Administrators of RS/6000 SP systems are expected to have level II skills (Junior System Administrator) or greater depending on the complexity of your site and system. See “Appendix A. SAGE job descriptions” on page 285 for more information on these job skills.

How this book is organized

- "Chapter 1. Overview of the installation and migration processes" on page 1 explains what happens when you perform the installation and customization procedure.
- "Chapter 2. Installing and configuring a new RS/6000 SP system" on page 11 provides the information you need to get the RS/6000 SP system set up and functioning in your environment.
- "Chapter 3. Installing and configuring the High Availability Control Workstation" on page 109 explains how to set up and configure your optional High Availability Control Workstation (HACWS).
- "Chapter 4. Migrating the software on your RS/6000 SP system" on page 127 describes how to migrate your system to the latest levels of AIX and PSSP.
- "Chapter 5. Reconfiguring security" on page 179 provides information on adding an authentication configuration to an existing SP system.
- "Chapter 6. Reconfiguring the RS/6000 SP system" on page 195 discusses adding and replacing hardware components.
- "Chapter 7. Performing software maintenance" on page 261 provides information on updating installation images, LPs, and support programs.
- "Chapter 8. Performing hardware maintenance" on page 271 provides information on replacing a frame supervisor, a fixed disk, a disk not in **rootvg**, and an I/O Planar card or Ethernet adapter.
- "Chapter 9. Installing the optional PSSP T/EC adapter" on page 279 provides information on installing a PSSP Tivoli Enterprise Console (T/EC) adapter.
- "Chapter 10. Installing extension nodes" on page 281 provides information on installing an extension node.
- "Appendix A. SAGE job descriptions" on page 285 contains system administrator job descriptions defined by SAGE.
- "Appendix B. Directory information" on page 289 contains a list of the directories created when you install the RS/6000 SP LP image.
- "Appendix C. SP Perspectives tasks" on page 291 provides instructions for completing common tasks using the SP Perspectives graphical user interface.
- "Appendix D. Boot/install server configuration commands" on page 295 contains a list of the boot/install server configuration commands.
- "Appendix E. User-supplied node customization scripts" on page 297 contains information on user-supplied customization scripts.
- "Appendix F. Overriding the PPSIZE in a mksysb image" on page 301 contains information on how to override the PPSIZE in a mksysb image.
- "Appendix G. Reserving ports" on page 303 contains information on reserving port numbers.

The back of the book includes a glossary, a bibliography, and an index.

Typographic conventions

This book uses the following typographic conventions:

Typographic	Usage
Bold	• **Bold** words or characters represent system elements that you must use literally, such as commands, flags, and path names.
Italic	• *Italic* words or characters represent variable values that you must supply. • *Italics* are also used for book titles and for general emphasis in text.
`Constant width`	Examples and information that the system displays appear in `constant width` typeface.
[ ]	Brackets enclose optional items in format and syntax descriptions.
{ }	Braces enclose a list from which you must choose an item in format and syntax descriptions.
|	A vertical bar separates items in a list of choices. (In other words, it means “or.”)
< >	Angle brackets (less-than and greater-than) enclose the name of a key on the keyboard. For example, **<Enter>** refers to the key on your terminal or workstation that is labeled with the word Enter.
...	An ellipsis indicates that you can repeat the preceding item one or more times.
<Ctrl-*x*>	The notation **<Ctrl-*x*>** indicates a control character sequence. For example, **<Ctrl-c>** means that you hold down the control key while pressing **<c>**.
****	The continuation character is used in coding examples in this book for formatting purposes.

Interface instructions

Some sections of this book give step-by-step instructions for performing tasks with a graphical user interface. The instructions use a format that distinguishes between the user action and the system response.

User actions appear in uppercase bold type.

PRESS Cancel

Selections from a menu bar are indicated with an →.

SELECT SP → Topology

Chapter 1. Overview of the installation and migration processes

This chapter provides an overview of the installation and migration processes. Before you install the SP system, you must perform several planning activities. These activities include completing worksheets or checklists that reflect your decisions regarding, but not limited to, the following:

- System partitioning (SP Switch only)
- Node and switch configuration
- System management options
- Boot/install servers and node relationships
- Location and type of authentication servers

It is essential that you plan your system carefully before attempting to install it. Refer to *RS/6000 SP: Planning, Volume 1, Hardware and Physical Environment* and *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for planning details.

Always review the *READ THIS FIRST* document that accompanies the PSSP installation media for the latest information. Assuming that the PSSP CD-ROM is in **/dev/cd0**, issue the following command to view that document:

```
installp -iq -d /dev/cd0 all
```

What's new in PSSP and AIX?

See the "Introduction to system planning" chapter in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* to find out what is new in this release of PSSP and AIX.

Using Perspectives

Perspectives is a graphical user interface you can use to manage and monitor the SP system. Complete instructions for using Perspectives appear in the online help information.

Note: As of PSSP 3.1, the System Monitor graphical user interface is no longer available.

You can use Perspectives during the installation and migration process to perform a variety of tasks including aid in installing and verifying software.

To start Perspectives, issue:

```
perspectives &
```

Note that many of the installation steps require you to use SMIT. You can use SMIT directly by issuing SMIT commands or you can access SMIT using Perspectives. Often used SMIT fast-paths, such as **smit_verify**, are available from the Perspectives Launch Pad.

When launching Perspectives using the **install_cw** command or from the command line, you may receive the following message:

```
Warning: locale not supported by C library, locale unchanged
```

This message indicates that you are attempting to run Perspectives in a locale that is not currently installed on your control workstation. Although Perspectives will run properly in this case, if you want to avoid seeing this message, check your LANG environment variable and ensure that it is set to a locale that is supported on your machine.

Installing your SP system

This section discusses some common terminology, as well as the major sections of the installation process.

Installation involves the following major steps:
1. Prepare (install or migrate) the control workstation with PSSP and AIX.
2. Enter configuration information for nodes that are new to the system.
3. Install and customize the nodes.

You complete all installation and configuration steps centrally from the control workstation.

Terminology

Before installing your system, become familiar with the following terminology used throughout this book. The definitions provide you with an overview only. For more detailed discussions, refer to the appropriate section in this book or the remainder of the PSSP library:

- Overwrite Install

  This type of installation writes a new version of AIX to the hard drive without saving any information on the system.
- Migration Install

  This type of installation preserves all file systems except **/tmp**.
- Coexistence

  Coexistence occurs when you mix different releases of PSSP software within an SP. There may also be a mixture of AIX software that supports the PSSP level.
- Network Installation Manager (NIM)

  NIM is an AIX component that assists in installing workstations (nodes) over communication networks.
- NIM Master

  In the NIM environment, a NIM master is an AIX system that can install one or more NIM clients. A system must be defined as a NIM master before any NIM clients are defined. A NIM master must be at the latest AIX level with the NIM master file sets installed. A NIM master manages the configuration database containing the information for the NIM clients.
- NIM Client

  A NIM client is a system installed and managed by a NIM master. SP supports the stand-alone type of NIM client.

Prepare and install the control workstation

The first step in installing the SP system is preparing the control workstation. This involves connecting your frame's RS-232 lines to the serial ports of your control workstation, configuring the Ethernet connections on your control workstation, and verifying name resolution. Other steps include installing AIX and PSSP software with required PTFs.

Enter node and configuration information

After preparing the control workstation, you can begin configuring the nodes. To configure them, you can use PSSP commands or the Software Management Interface Tool (SMIT) interface to enter information into the System Data Repository (SDR). The information you enter is based on the planning worksheets you completed. If you have not yet completed these worksheets, refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* and understand the information presented there before you begin installing your system. (Both configuration methods are described throughout the detailed installation information.)

Install PSSP on the nodes (externals)

PSSP node installation consists of these phases:

- Network installation
- Customization
- Boot processing

In the network installation phase, PSSP uses the AIX Network Installation Management (NIM) support to restore an AIX mksysb image created with the **mksysb** command. It can also perform an AIX migration install to the hard disk of the node. During the customize phase, PSSP personalizes the node with host names, default routes, and TCP/IP network adapter information from the SDR. AIX reboots the node. Final customization takes place after the node reboots, during boot processing.

Use the **spbootins** command to change the **bootp_response** attribute of the Node object in the SDR. You can set this attribute to **install**, **customize**, or **migrate**. For more information on these attributes, refer to *PSSP: Command and Technical Reference*.

Network installation

To install the nodes, network boot the nodes. Boot/install servers on the same network respond by installing the AIX image on the nodes.

Customization

Customization involves the following:

- Updating the node's host name and default route
- Defining the network adapters
- Installing the PSSP LP options on the node
- Running the **script.cust** you updated for your site-specific customization
- Running the **tuning.cust** you updated for your site-specific customization
- Running the **firstboot.cust** you updated for your site-specific customization

You customize your system after installing it. You can also customize a node that already has AIX and PSSP installed at any time. Situations that would call for a customize-only step include:

- If you have added or changed the networking or routing for one or more nodes.
- If you have added or changed a Kerberos Version 4 (V4) authentication server or reinitialized it, invalidating the Kerberos V4 service keys stored on the nodes.

You can do additional customization in one of three ways:

- By copying the **/usr/lpp/ssp/samples/firstboot.cust**, **/usr/lpp/ssp/samples/script.cust**, and **/usr/lpp/ssp/samples/tuning.cust** files on the control workstation to the **/tftpboot** directory and by tailoring these files to meet your needs. Commands in these files are run after the mksysb installation

and PSSP customization and before the reboot of each node. IBM suggests you use **firstboot.cust** and **script.cust** to complete your installation and **tuning.cust** to set initial network performance values.

- By selecting one of the IBM-supplied tuning files to set your performance parameters. One file is for the commercial environment, one is for the development environment, and one is for the scientific environment. For more information on these files, refer to “Chapter 2. Installing and configuring a new RS/6000 SP system” on page 11 or *PSSP: Command and Technical Reference*.
- By not selecting any existing files. If you choose this method, PSSP automatically uses the **tuning.default** file to set your initial performance parameters.

Before the initial boot of a node, the required RS/6000 SP options (such as **ssp.clients**, **ssp.basic**, and **ssp.sysctl**) are installed if they are not part of the AIX image already on the node.

Boot processing

After PSSP completes installation and customization, AIX reboots the node. PSSP can do additional configuration during the node reboot depending upon the choices you made in the SMIT Site Environment panel or with the **spsitenv** command.

Install PSSP on the nodes (internals)

The previous section “Install PSSP on the nodes (externals)” on page 3 provided a high-level discussion of the installation process. This section provides a more detailed discussion of the process and is intended to provide you with additional, in-depth information.

AIX provides the NIM environment to install AIX on the nodes. PSSP defines the control workstation as a NIM master. If a system has more than 40 nodes, boot/install servers are defined as NIM masters so they can install nodes. In this case, the first node in each frame, by default, is a NIM master for its frame. You can change this by using the **spchvgobj** command. You can define other nodes as boot/install servers and have PSSP make them NIM masters. This section discusses the parts of NIM exploited by the SP system. For more detailed information about the NIM environment, see the version of *IBM AIX Network Installation Management Guide and Reference* that is appropriate for your level of AIX.

The NIM environment includes a NIM master, a network, and its clients. The NIM master contains the information necessary to install its NIM clients.

Installing and migrating nodes (NIM clients)

After PSSP configures the NIM masters and the NIM clients, you can begin installing the nodes. If you have fewer than 40 nodes, the control workstation installs the nodes. Otherwise, the PSSP default is for the first node in each frame to install other nodes in that frame. You may select other SP nodes to be the NIM masters. This depends upon the physical layout of the installation network configuration.

When the node is network booted, the NIM master completes the installation and migration steps as follows:

1. Installs the AIX mksysb image for installation or performs an AIX migration install for migration.
2. Runs **pssp_script** on the node to do the following:
   a. Defines the host name and configures the primary adapter interface

b. Installs the necessary AIX and PSSP file sets if they are not already installed

c. Configures the switch if one is installed.

d. Copies and runs **tuning.cust** if it exists in the boot/install server's **/tftpboot** directory.

e. Copies and runs **script.cust** if it exists in the boot/install server's **/tftpboot** directory.

3. Reboots the node for more PSSP configuration. During this initial post-installation reboot, **psspfb_script** is run on the node from **/etc/inittab**. The **psspfb_script** performs the following tasks:

   a. Defines and configures the network adapters specified in the SDR for this NIM client.

   b. Calls **spauthconfig** to set up the security environment on the node.

   c. Copies and runs **firstboot.cust** if it exists in the boot/install server's **/tftpboot** directory.

Migrating your SP system

"Chapter 4. Migrating the software on your RS/6000 SP system" on page 127 of this book discusses how to migrate your system to the latest level of AIX and PSSP. You should also refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* to understand how to plan your migration. In general, you can migrate to PSSP 3.4 and AIX 4.3.3 or AIX 5L Version 5.1 (AIX 5L 5.1) from several starting points as shown in the list that follows. Note that you do not have to migrate your entire system at once. You can choose to migrate your system all at once or several nodes at a time.

- PSSP 3.2 and AIX 4.3.3
- PSSP 3.1.1 and AIX 4.3.3
- PSSP 2.4 and AIX 4.2.1 or 4.3.3

Important

If the starting point for a node migration is PSSP 2.4 and AIX 4.2.1, you must first migrate to PSSP 3.4 and AIX 4.3.3. A second migration is required to go from PSSP 3.4 and AIX 4.3.3 to PSSP 3.4 and AIX 5L 5.1. The two migrations cannot be combined.

To migrate a node, the node's boot/install server must be running at or above the PSSP and AIX levels that you want to migrate the node to. For example, if you want to migrate the node to PSSP 3.4 and AIX 4.3.3, the boot/install server must be running PSSP 3.4 and AIX 4.3.3 or PSSP 3.4 and AIX 5L 5.1. You can use the **spchvgobj** command to perform the following functions:

- Change the boot/install server
- Change the **lppsource_name**
- Change the **code_version**

You can use the **spbootins** command to change the node to migrate or customize.

Working with the basic AIX image

The **spimg** installp image contains one or more AIX mksysb images that you can install into the **/spdata/sys1/install/images** directory. You can use these mksysb images to install AIX on the nodes in your system after you have successfully installed and configured your control workstation. These mksysb images contain the supported AIX LPs and files to install the SP nodes. The minimal system is made up of the following:

- The basic operating system
- Extensions
- Networking
- Basic utilities

IBM provides the image as a base for getting your system up and running quickly. You can customize this mksysb image to meet the needs of your production environment.

See the *READ THIS FIRST* document for details on the supported AIX levels and APARs that are contained in the **spimg**.

Building a new AIX image

Prior to creating a mksysb, you must do the following:

Notes:

1. Make sure that no files named **/etc/niminfo** or **/etc/niminfo.prev** exist. If they do exist, rename them. These files should be saved for possible debugging later on.
2. Verify the host name resolution for the control workstation and any nodes where the mksysb may be installed, if they are listed in the **/etc/hosts** file.
3. If DCE is running on the host that the mksysb image is made from, you must first turn autostart off for the DCE daemons. To do this, issue:

   ```
   config.dce -autostart off
   ```

   then create the mksysb image.

When you are installing your system for the first time, you can generate a mksysb image and install that on your nodes during the initial installation. To generate this image, do either of the following:

- Generate a mksysb on a node after you have installed the node using the basic image and additional LPs required for your configuration. You first should install the SP node using the base spimg and the PSSP LPs. Then, you can use **firstboot.cust** and **script.cust** to install additional LPs and service your system requires. At this point, you can test the node with applications specific to your users' needs to make sure it meets your environment requirements. When you are satisfied this node has all of the LPs and service, you can make a mksysb image from it by using the SMIT mksysb menu.
- Use a standalone workstation (either the control workstation or a different standalone workstation).

  If you use the control workstation to create the mksysb image, the image should be created before defining the space for the **/spdata** directory. After you have completed "Step 21: Initialize RS/6000 SP Kerberos V4 (optional)" on page 38, you cannot use the control workstation to build a new image. Running the **setup_authent** command on any workstation makes the workstation unsuitable

for generating a mksysb image for a node. This is because **setup_authent** sets up authentication databases that are unique to the machine.

The control workstation (or a different standalone workstation) must be at the correct level of the AIX operating system and must have all required PTFs applied. See the *READ THIS FIRST* document for the list of required AIX operating system level and PTFs. One way to ensure this is to install one of the mksysb images that resides in the spimg installp image on its installation media.

After the image is created, you must copy it to the **/spdata/sys1/install/images** directory on the control workstation. This directory is the starting point for all the mksysb images used for the RS/6000 SP system. Make sure the directory has permissions `rwxr-sr-x` and the images have permissions `rw-r--r--`.

Changing the node's install image attributes

After creating and copying a new image to the control workstation, you are ready to update the SDR with the name of the new image for the nodes you want to install. To define a new image, you can use the **spchvgobj** command or the SMIT **changevg_dialog** interface and the **setup_server** command to do the following:

- Redefine the node's attributes in the SDR.
- Create new installation and configuration files for the boot/install server.
- Copy the new image to the boot/install server for the nodes to be reinstalled.

Detailed overview of the network installation of a node

The information in this section is optional. Read this section if you want to understand the details of the network installation described in "Install PSSP on the nodes (externals)" on page 3.

NIM-specific terminology and concepts

- NIM Master

  A site may have more than one NIM master, but these NIM masters and their clients are considered to be separate NIM environments. NIM does not provide any services to share information across these separate environments.

  As a default, the control workstation is the NIM master for a single frame system. If you have more than one frame, the default configuration is different: the control workstation installs the first node in each frame and defines it as a NIM master. Then, each NIM master installs the rest of the nodes in its frame. You can change the default by using SMIT **changevg_dialog** or the **spchvgobj** command.

  Whether you use the default configuration or define your own, you use the **spchvgobj** command to define the boot/install server. The **spchvgobj** command updates the SDR with the information specified with the command. The **setup_server** command can be invoked on the boot/install servers to define them as NIM masters.

  The **setup_server** command configures and defines the NIM master, resources, and clients that allow the nodes to be installed using the NIM environment.

- NIM Objects

  NIM stores information about the NIM environment as objects in the NIM database. The types of objects defined for the SP system are:
  - Network
  - Machine
  - Resource

Network objects represent information about the network interfaces required for installing a client. The SP system uses the SP Ethernet administrative local area network (LAN) adapter interface on the client for the NIM installation.

Although the SP system may support multiple Ethernet networks depending upon the physical and defined subnets for your system, only the client's SP Ethernet administrative LAN adapter interface is supported for the NIM installation.

Machine objects represent the machine configuration for a NIM client. Each machine object contains attributes which define it. Examples of machine attributes are:
- Type of machine (standalone)
- Hardware address of the client's network installation interface (SP Ethernet administrative LAN adapter)
- TCP/IP host name of the client
- Name of the network object defining the network to which the client is connected
- Type of processor - Uniprocessor (UP) or Multiprocessor (MP).

Resource objects represent available resources in the NIM environment. All operations on a client in a NIM environment require one or more resources to be allocated. Examples of resources available in a NIM environment are:
- SPOT - Shared Product Object Tree (SPOT)
- mksysb - AIX mksysb image
- script - Scripts run as part of the NIM customization on the client
- lppsource - Directory of installable LP images
- bosinst_data - Prompt, noprompt, and migrate installation information

See *IBM AIX Network Installation Management Guide and Reference* for a complete list and descriptions of the resource objects.

The SP system supports installation of mksysb images only through NIM mksysb resource objects.

Network installation

When a node boots over the network, it issues a **bootp** request on that network specifying its network device hardware address. (For the RS/6000 SP, this is the node's SP Ethernet administrative LAN adapter). The boot/install server has a list of nodes it boots and their associated hardware Ethernet addresses.

The **bootp** daemon on the boot/install server gets the request and looks in its table (**/etc/bootptab**) for this node's hardware Ethernet address. If it finds the address in its table, it responds by sending the node's IP address.

The node's IPL Read-Only Storage (ROS) requests a boot image transferred to the node using **tftp**. The boot image is run, the rootvg gets created, and NIM performs the installation of the mksysb image. After the mksysb installation is complete, NIM invokes the **pssp_script** script resource which transfers the install information files from the boot/install server to the node and performs the customization. These install files contain information necessary to successfully complete the netinstall.

Customization for network install

pssp_script transfers the following netinstall information files from the boot/install server to the node and then performs the customization:
- **/tftpboot/***node_name***.install_info**

- **/tftpboot/***node_name.***config_info**

Fields in the **install_info** file contain the following information:
- Client's SP Ethernet administrative LAN adapter IP address and host name
- SP system and node-related data from the SDR
- Netinstall server IP address and host name
- Control workstation IP information and host name
- lppsource
- PSSP code version
- Kerberos Version 4 server's information
- DCE Master Security and Cell Directory Services (CDS) servers' information
- Node's system partition host name and IP address

Fields in the **config_info** file contain the following information:
- Node number
- Switch node information
- Default route
- Initial host name
- Root volume group information
- An entry for each adapter defined in the SDR, listing the following:
  - Adapter name
  - Adapter IP address host name
  - Netmask
  - ring_speed (for token ring)
  - Cable type (bnc, dix, tp, fiber, NA)
  - Duplex (full, half, auto)
  - Ethernet speed (10, 100, 1000, auto)
  - Aggregate adapter information
  - Physical location code
  - SP Ethernet administrative LAN

After a node is installed, **pssp_script** gets a copy of the **config_info** file from its master. The information in the **config_info** file is used to define the adapters, set the hostname, define the default route, and to customize the adapters with information such as host name and default route. This file is transferred from the server to the node during customization. The file is read at this time and the node is configured.

The PSSP file sets **ssp.basic**, **ssp.perlpkg**, **ssp.sysctl**, **ssp.ha**, **ssp.clients**, and **ssp.sysman** are installed if they are not already installed. The **ssp.css** file set is always installed on systems with a switch. The **ssp.st** option is installed if the node's boot/install server has it installed.

As part of completing the network install, NIM modifies the boot/install server's **/etc/bootptab** file so that the node does not perform a netinstall again on the next boot. PSSP changes the SDR node object **bootp_response** attribute to **disk**.

Boot processing

After customization, a bootable image is written to the node's disk. The node is then shut down and rebooted. When it is rebooted, PSSP checks the SDR to determine if any of the PSSP Site Environment options were chosen. If an option is selected, the node configuration for that system management option is performed now. When the node reboots, if it is a boot/install server, **setup_server** is run.

Configuring boot/install servers as NIM masters

When you run **/usr/lpp/ssp/bin/setup_server** to configure a boot/install server, it does the following:

- Creates **/tftpboot/***node_name***.install_info** and **/tftpboot/***node_name***.config_info** for all its client nodes.
- Configures the server as a NIM master
- Defines among other items, the spot, mksysb, script, lppsource, bosinst NIM resources
- Defines nodes to be installed or booted in diagnostics or maintenance as NIM clients
- Allocates NIM resources to the NIM clients as needed
- Copies the mksysb images (if any of its nodes need to be installed) to the boot/install server that services the nodes.
- Creates required authentication files, Kerberos V4 keyfiles
- Executes the NIM **bos_inst** command which allocates the NIM boot and NIM script resources, updating **/etc/bootptab**

Chapter 2. Installing and configuring a new RS/6000 SP system

This chapter describes how to install the IBM Parallel System Support Programs for AIX (PSSP), configure your control workstations, servers and processor nodes, and initialize the system.

Notes:

1. Perform the steps in this chapter only if you are installing a new system or if you are doing a complete reinstall. Do not perform the steps if you are migrating from one level of PSSP or AIX to another level. For migration steps, refer to "Chapter 4. Migrating the software on your RS/6000 SP system" on page 127.
2. Make sure you have your configuration worksheets complete and available before you proceed with installation. During the procedure, you are prompted for network, host name, and configuration data. It is easier to install the system if you plan your configuration beforehand. Refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for planning information.
3. You can install your system by using either the System Management Interface Tool (SMIT) or PSSP commands. SMIT is a graphical menu-driven interface to enter information into the SDR. Experienced users may prefer to issue the PSSP commands directly. Throughout this chapter, instructions for both methods are given. Use default options unless instructed to do otherwise.
4. Systems with Default Options Installed or Preloaded Systems

   One of the decisions you must make prior to installing your system is how you want to receive the SP. You have the option of purchasing your SP with the default software installed or you can purchase one that IBM has preloaded with software to meet your organization's specific needs. For more detailed information on these preloaded options, refer to the "Introduction to system planning" chapter in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*.
5. Reserving Port Numbers

   Some of the subsystems managed by Syspar Controller (**syspar_ctrl**) allocate port numbers from the range 10000 to 10100, inclusive. Therefore, if any customer subsystem, such as DB2, uses port numbers in this range, such port numbers must be reserved in **/etc/services**. For details, refer to "Appendix G. Reserving ports" on page 303.
6. Restricted Root Access

   As of PSSP 3.2, you have the option of running your SP system with an enhanced level of security. With the restricted root access (RRA) option enabled, PSSP does not internally issue **rsh** and **rcp** commands as a root user from a node. Also, PSSP does not automatically grant authorization for a root user to issue **rsh** and **rcp** commands from a node. If you enable this option, some procedures might not work as documented. For example, to run HACMP, an administrator must grant the authorizations for a root user to issue **rsh** and **rcp** commands that PSSP otherwise grants automatically. See the ″Planning for security″ chapter in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for a description of this function and a complete list of limitations.
7. Secure remote commands

   Refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more information.

8. If you plan on installing a firewall in your SP, review *Implementing a Firewalled RS/6000 SP System* before starting to install your SP. With proper planning, the network and security configuration required for a firewall can be implemented during the installation. For example, you may want to:
   a. Determine your trusted and untrusted nodes and configure your SP administrative LAN to physically separate the trusted and untrusted sides on to two Ethernet subnets before starting the installation. The logical configuration of the two Ethernets would be done in "Step 37: Enter the required node information" on page 62.
   b. Enable restricted root access at "Step 30: Enter site environment information" on page 49.
   c. Enable a secure remote command method at "Step 30: Enter site environment information" on page 49.

Finding related installation information

To find out more information, refer to the following manuals:
- *IBM AIX General Concepts and Procedures for RS/6000*
- *IBM AIX Network Installation Management Guide and Reference*
- *IBM AIX Commands Reference*

Task A. Prepare the control workstation

This section describes the steps you take to prepare the control workstation. Note that it is a prerequisite for the control workstation to already have the correct level of AIX installed.

Step 1: Verify the control workstation requirements

Ensure that the following requirements are met:

Step 1.1: Verify the inst_root directories

If your control workstation has had the **inst_root** directories removed, the system cannot be used as a SPOT server. To determine if the **inst_root** directories have been removed, issue:

```
/usr/lib/instl/inurid -q; echo $?
```

If the result is 1, you will need to reinstall.

In addition to running the **inurid -r** command manually, the **inst_root** directories may have been removed on your system:

- If the control workstation was installed by a non-mksysb NIM installation using a bosinst resource that had the **RM_INST_ROOT** variable set to **yes**.
- If the control workstation was installed from a mksysb that had the **inst_root** directories removed before it was created.

Step 1.2: Review the READ THIS FIRST document

Refer to the *READ THIS FIRST* document that accompanies the PSSP installation media for the latest information for supported AIX release and software update (PTF) levels.

Step 1.3: Verify the control workstation required software

The RS/6000 machine you use as the control workstation must have AIX Version 4.3.3, AIX 5L 5.1, or later Base Operating System software installed.

Step 1.4: Verify the control workstation serial ports

The control workstation requires one serial port available per SP frame for the RS-232 cable that supports hardware monitoring and control.

The following tables lists the number of tty ports you must allocate for the specific servers.

SP model numbers	tty port values
IBM @server pSeries 690	Does not require a tty port value, but the Hardware Management Console (HMC) must be connected by the SP Ethernet administrative LAN
RS/6000 H80, M80, and IBM @server pSeries 660 (6H0, 6H1, 6M1)	1
RS/6000 S70, S7A, and S80 or IBM @server pSeries 680	2

Refer to “Step 33: Enter non-SP frame information and reinitialize the SDR (optional)” on page 56 for more information.

Step 1.5: Verify the control workstation disk space

The control workstation must have ample disk space for the boot/install server files. See *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for detailed information about space requirements.

Step 1.6: Install the Distributed Computing Environment (DCE) on the control workstation (optional)

If you plan to have PSSP use DCE, you must install DCE either as a client or server on the control workstation.

Note: If using DCE, you should use a separate file system for the DCE configuration files. DCE-related file systems should be created under **/var**. Refer to *IBM Distributed Computing Environment for AIX: Administration Guide—Introduction* for more information.

Step 1.7: Verify the AIX error log size

You must verify that the AIX Error Log is at least 4096000 bytes. For example, issue:

```
/usr/lib/errdemon -l
```

To change the size of the AIX Error Log, issue:

```
/usr/lib/errdemon -s 4096000
```

Step 1.8: Update the root user path

You must run all installation tasks from root on the control workstation. All of the SP commands in this chapter are located in the following directories:

- **/usr/lpp/ssp/bin**
- **/usr/lib/instl**
- **/usr/sbin**
- **/usr/bin**
- **/usr/lpp/ssp/kerberos/bin**

Note: Refer to the “Naming conflict if using both Kerberos Version 4 and DCE” section in the “Security features on the SP system” chapter of *PSSP: Administration Guide* for additional information on using the **k4destroy**, **k4init**, and **k4list** commands.

To avoid entering the complete path name each time you want to invoke a command, you should add the directories in the previous list to the root user's path. For example, if using the Korn shell, add the directories to the path statement in root's **.profile** file. The following is an example of the root's user path statement:

```
PATH=$PATH:/usr/lpp/ssp/bin:/usr/lib/instl:/usr/sbin: \
/usr/lpp/ssp/kerberos/bin
```

If you would like to use the man pages during your install, you must set the MANPATH as follows:

```
MANPATH=$MANPATH:/usr/lpp/ssp/man
```

In order for this path statement to take effect, you first have to execute **.profile**. For example, issue:

```
. /.profile
```

To verify that this step completed successfully, issue:

```
echo $PATH
```

Step 1.9: Verify name resolution on the control workstation

All names and addresses of all IP interfaces on SP nodes must be resolvable on the control workstation before you install and configure the SP.

Ensure that the control workstation host name is set. After you have set the host names and IP addresses on the control workstation, you should not change them for the duration of the installation.

Step 1.10: Install the secure remote command software

PSSP 3.4 provides the ability to remove the dependency PSSP has on the **rsh** and **rcp** commands issued as root by enabling the use of a secure remote command method. It is the administrator's responsibility to choose the secure remote command software and install it on the control workstation. When installing the software, you should use a separate file system for the secure remote shell configuration files and keys to ensure that they are not picked up in the mksysb to be installed on the nodes.

All nodes must be at PSSP 3.2 or later before the secure remote command method can be enabled.

The secure remote command software will be installed on the nodes by editing the **script.cust** file to install the secure remote command software after the PSSP code is installed.

There is usually a requirement that root be the owner of it's home directory to run secure remote commands. By default, this is not true in a general AIX installation (earlier than AIX 5.1) with the root home directory defaulting to (***/***) owned by bin. Either **/** must be changed to be owned by root, or the administrator must create a new home directory for root. In this environment, it is also beneficial to create root's home directory in a separate file system to make sure that any keys generated by the root user are not contained in the mksysb and become installed on the nodes.

Public keys must be generated and known_hosts file set up such that the PSSP code can run the secure remote commands as root without being prompted for passwords or passphrases. The known_hosts files must contain both long and short names for the nodes. Applications that have changed to support the secure remote command environment such as **dsh** will appear to hang if the secure remote command code issues a prompt when the command is run. See *SSH The Secure*

Shell The Definitive Guide for additional information. See *PSSP: Diagnosis Guide* for diagnosing secure remote setup problems.

After the secure remote command software is installed and the daemon is started on the control workstation, the administrator should ensure that the root ID can issue a secure remote command and secure remote copy command from the control workstation to the control workstation without being prompted for input (for example, passwords, passphrases). This ability to run secure remote commands and copies is required to run the PSSP code using the secure remote command method.

Some of the secure remote copy products default the location of the secure remote command to **/usr/local/bin/ssh**. If your product does this, be sure that the secure remote command executable is located in this directory or linked to this directory.

When the PSSP code issues secure remote commands and copies, it uses no options on the call to the secure remote product. In this way, we have built no dependency on a particular implementation of a secure remote command product.

Step 2: Verify the network requirements

Using your node worksheets, check all SP Ethernet administrative LAN and additional adapter information to ensure that you have assembled the correct host name, IP address, and netmask data for installation.

Step 3: Connect frames to your control workstation

Connect RS-232 and Ethernet cables from the SP system frames and from the SP-attached servers or clustered enterprise servers to the control workstation according to your SP Control Workstation Network Worksheet. Your IBM Customer Engineer (CE) performs this step. The CE will need access to the control workstation to run diagnostics. See *RS/6000 SP: Installation and Relocation* for instructions.

Step 4: Configure RS-232 control lines

You must allocate the appropriate number of tty ports in this step. Each SP frame in your system requires a serial port on the control workstation configured to accommodate the RS-232 line. SP-attached servers require different numbers of tty ports depending on the server type. For example, an S80 requires two tty ports, so if you have two SP frames and one S80 SP-attached server, configure four tty terminals. Refer to “Step 1.4: Verify the control workstation serial ports” on page 13 for valid tty port values.

If using:	Do this:	
SMIT	**TYPE**	**smit tty** • The TTY menu appears
	SELECT	Add a TTY
	SELECT	tty rs232 Asynchronous Terminal
	SELECT	An appropriate serial port (parent adapter) connection. For example, **sa0**. • A data entry window appears
	PRESS	**List** to show the available port numbers. Select a port number. For example, **0, 1, 2, 3**.
	PRESS	**Ok** to configure the line.

mkdev	Enter a command similar to this example, which defines and configures an RS-232 line for parent adapter sa0 on port serial port 1:
	`mkdev -c tty -t tty -s rs232 -p sa0 -w 1`
	This example configures a second port of a two-frame system:
	`mkdev -c tty -t tty -s rs232 -p sa1 -w 2`	

You can accept the default Baud rate of 9600. The internal system daemon called **hardmon** changes the rate to 19200.		

Step 5: Tune all control workstation network adapters

Various models of network adapters can have different values for transmit and receive queue sizes. The queue setting for Micro Channel adapters is 512. For PCI adapters, the queue setting is 256 or greater.

You can set these values using SMIT or the **chdev** command. If the adapter you are changing is also the adapter for the network you are logged in through, you will have to make the changes to the database only. Then reboot the control workstation for the changes to become effective.

If using:	Do this:	
SMIT	**TYPE**	**smit devices**
		• The Devices menu appears
	SELECT	Communication
		• The Communication menu appears
	SELECT	The adapter you want to reset (for example, `Ethernet Adapter`)
		• The Ethernet Adapter menu appears
	SELECT	Adapter
		• The Adapter menu appears
	SELECT	Change / Show Characteristics of an Ethernet Adapter
		• An Ethernet Adapter window appears
	SELECT	An adapter from the list shown
		• The Change / Show Characteristics of an Ethernet Adapter window appears
	CHANGE	The HARDWARE TRANSMIT queue size. (See note 2 below.)
	CHANGE	`Apply change to DATABASE only` to **yes**. (Press List and select yes.)
	PRESS	**Ok** to apply the changes to the database.
chdev	Enter `chdev -P -l ent0 -a xmt_que_size=512`	

1. You must reboot the control workstation in order for the changes to take effect.
2. To determine the name of the TRANSMIT queue size, issue:

   ```
   lsattr -l adapter_name -E
   ```

   In response, a list of all of the values for the different variables, what they are, and the name of the variable will be displayed.

   In the previous example, an MCA adapter with AIX 4.2.1 or later was used.

Step 6: Configure the control workstation Ethernet adapters

Use SMIT or the **chdev** command to configure each Ethernet adapter connecting the nodes on your frames to the control workstation. For details on the correct use of **chdev**, see *IBM AIX Commands Reference*, the man pages, or the online information database.

Refer to your SP Control Workstation Network Worksheet in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*.

If the adapter is not yet defined or configured, use **smit mkinet** or the **mkdev** command instead of **smit chinet** or **chdev** to specify a new IP host name and netmask values. If you are adding an extension node to your system, you may want to configure the adapters now. For more information, refer to “Chapter 10. Installing extension nodes” on page 281.

If using:	Do this:	
SMIT	**TYPE**	**smit chinet** • The Available Network Interfaces menu appears
	SELECT	The Ethernet interface to be configured
	ENTER	IP interface and netmask information from worksheet.
	SELECT	**up** as the `Current STATE` option
	PRESS	**Ok** to complete operation
chdev	Enter a command similar to this example, which configures an SP Ethernet administrative LAN adapter. `chdev -l en0 -a netaddr=129.33.41.1 -a netmask=255.255.255.0 -a state=up`	

mkdev	Enter a command similar to this example, which defines and configures an SP Ethernet administrative LAN adapter.
	`mkdev -c if -s EN -t en -a netaddr=129.33.34.1 \`
	`-a netmask=255.255.255.0 -a state=up -q -w en0`	

Step 7: Verify the control workstation interfaces

Verify the configuration for each Ethernet adapter in the control workstation. You can verify that the adapter is installed even if it is not cabled to the SP system yet.

Verify each Ethernet adapter by *pinging* its IP address and seeing if you get a proper response. If you do not receive a response, debug the network problem, and reconfigure the adapter.

For example:

```
ping -c 1 129.33.34.1
```

Step 8: Ensure that the necessary daemons are running on the control workstation

1. Check to see if the System Resource Controller (SRC) is running by issuing the following:

   ```
   lssrc -a
   ```

   If you get an error message similar to the following, then SRC is not running:

   ```
   The System Resource Controller daemon is not active
   ```

   To start SRC, uncomment or add the srcmstr entry in the **/etc/inittab** file and reboot the control workstation using the **shutdown -Fr** command.

   In **/etc/inittab**, the srcmstr entry should look like:

   ```
   srcmstr:2:respawn:/usr/sbin/srcmstr # System Resource Controller
   ```

2. Check to see if the bootps and tftp entries are uncommented in the **/etc/inetd.conf** file. If they are commented (a leading pound sign designates a comment), uncomment the two entries in the file and issue the following:

   ```
   refresh -s inetd
   ```

Step 9: Change the control workstation maximum default processes

When you first install your system, the number of processes is set to an AIX default. You will not be able to continue installing your system with this default value—you must increase the number. IBM suggests changing the maximum number to 256.

If using:	Do this:	
SMIT	**TYPE**	smit system • The System Environments menu appears
	SELECT	Change / Show Characteristics of Operating System
	ENTER	256 for the maximum number of processes allowed per user.
	PRESS	Ok
chdev	Enter: `chdev -l sys0 -a maxuproc='256'`	

Step 10: Change the control workstation tunables and tunable values

When you first install your system, the network tunable values are set to AIX defaults. (A **tunable** is a performance parameter you can set to a value that makes your system run its workload most efficiently.) Your system may not run efficiently with the default values. Use the **no** command to display these values. This command and all its parameters are described in *IBM AIX Commands Reference*.

When you install PSSP on your control workstation, change the network tunables on the control workstation to the suggested values in the following table.

Table 1. Control Workstation Tunables

Tunable	Recommended Initial Value	Description
thewall	16384	The upper bound on the amount of real memory that can be used by the communications subsystem. The units are in 1 KB increments. **Note:** As of AIX 4.3.3, the recommended initial value should not be set because it is automatically sized by the system at boot.
sb_max	163840	Upper limit on the size of the TCP and UDP buffers in mbufs of allocated space to a connection.
ipforwarding	1	Specifies whether the kernel should forward packets. A value of **1** forwards packets when they are not for the local system; a value of **0** prevents forwarding.
tcp_sendspace	65536	The default size of the TCP send window.
tcp_recvspace	65536	The default size of the TCP receive window.
udp_sendspace	32768	The default size of the UDP send buffer. The effective maximum is 65536 (64K).
udp_recvspace	65536	The default size of the UDP receive buffer.
tcp_mssdflt	1448	Maximum package size for remote network.
tcp_pmtu_discover	0	TCP MTU path discovery (AIX 4.3.1 or later).
udp_pmtu_discover	0	UDP MTU path discovery (AIX 4.3.1 or later).

Using the no Command

To display network tunable values, enter:

```
no -a
```

To change the value of **tcp_mssdflt**, enter:

```
no -o tcp_mssdflt=1448
```

When you change the network tunables, they take effect immediately. However, they are not preserved across a boot. To make the changes to the tunables effective across boots, add the **no -o** commands you used to change the network tunables to the last section of the **/etc/rc.net** file. Using the same syntax, place the commands under the line:

```
/usr/sbin/no -o extendednetstats=0 >>/dev/null 2>&1
```

For example:

```
/usr/sbin/no -o tcp_mssdflt=1448 >>/dev/null 2>&1
```

Step 11: Define space for the /spdata directory

The **/spdata** directory contains, among other items, mksysb, and installp file sets. IBM suggests you create a separate volume group for the **/spdata** file system. These file sets require a minimum of 2 GB of disk space. You will require additional disk space if you need to support multiple AIX and PSSP release levels, and multiple mksysb images. If you have not done so already, use *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* to help you estimate how much space you need to define.

To define space for the **/spdata** directory, refer to the following table to determine which steps to follow. Do either Procedure A or B. Do not perform more than one procedure.

Procedure	To:	Follow these steps:.
A	Define a new volume group	1. Do “Step 11.1: Define a volume group” 2. Do “Step 11.2: Create the logical volume for /spdata” on page 21 3. Do “Step 11.3: Create the file system for /spdata” on page 21 4. Do “Step 11.4: Mount the /spdata file system” on page 22 5. Do “Step 12: Create the required /spdata directories” on page 22
B	Use an existing volume group (such as **rootvg**) and define a new file system.	1. Do “Step 11.2: Create the logical volume for /spdata” on page 21 2. Do “Step 11.3: Create the file system for /spdata” on page 21 3. Do “Step 11.4: Mount the /spdata file system” on page 22 4. Do “Step 12: Create the required /spdata directories” on page 22

Step 11.1: Define a volume group

The logical volume manager (LVM) configuration you define can also include placing the **/spdata** file system in its own volume group rather than using the default volume group **rootvg**.

You cannot create a separate volume group if all control workstation disks are in use or the control workstation has only one physical volume. When working with larger disks that are greater than 4 GB, you must specify the physical partition size to be 8 MB or more for the spdata volume group.

If using:	Do this:
SMIT	**TYPE** **smit vg** • The Volume Groups menu appears **SELECT** Add a Volume Group • The Add a Volume Group window appears **TYPE** The volume group name (for example, **spdatavg**). **TYPE** The physical volume name (press F4 and select a volume). **PRESS** **Ok** to create the volume group. SMIT automatically varies on the volume group.

mkvg	The following command will create a new logical volume named **spdatavg**, using **hdisk1** as the physical volume.
	`mkvg -f -y spdatavg hdisk1`
	After you create the new volume group **spdatavg**, vary on the volume group on your control workstation using the **varyonvg** command.
	`varyonvg spdatavg`

Step 11.2: Create the logical volume for /spdata

You now need to create the logical volume called spdata_lv with 2 GB, assuming a 4 MB physical partition size.

If using:	Do this:
SMIT	**TYPE** **smit mklv** • The Add a Logical Volume menu appears **TYPE** The volume group name (for example, **spdatavg** or press F4 and select a volume). **TYPE** The logical volume name (for example, **spdata_lv**). **TYPE** The number of logical partitions (for example, 500). **TYPE** The maximum number of logical partitions (for example, 512). **PRESS** **Ok** to create the logical volume.
mklv	To create a logical volume spdata_lv, use the **mklv** command: `mklv -y spdata_lv -x 512 spdatavg 500`

Step 11.3: Create the file system for /spdata

You now need to create the **/spdata** file system on the control workstation. The file system must be mounted as **/spdata**.

If using:	Do this:	
SMIT	**TYPE**	**smit crfs**
	SELECT	Add a Journaled File System (AIX 5L 5.1 systems only)
	SELECT	Add a Journaled File System on a Previously Defined Logical Volume
	SELECT	Add a Standard Journaled File System
	SELECT	Your logical volume (press F4 and select a volume).
	SPECIFY	**/spdata** as mount point
	SELECT	**yes** for `Mount AUTOMATICALLY at system restart?`
	PRESS	**Ok** to create the file system.
crfs	To create a file system **/spdata** within the newly-created logical volume **spdata_lv**, use the **crfs** command. The following example creates the **/spdata** file system, using logical volume **spdata_lv** and mount point **/spdata**. It also adds the **/spdata** into the **/etc/filesystems** to be automatically mounted at system restart. `crfs -v jfs -d spdata_lv -m /spdata -A yes -p rw -t no -a bf=true`	

Step 11.4: Mount the /spdata file system

After you create the new **/spdata** file system, you need to mount it on the control workstation. Use the AIX command **mount**.

The following example mounts the **/spdata** file system:

```
mount /spdata
```

Step 12: Create the required /spdata directories

Make sure you mount the new **/spdata** file system before you create the **/spdata** directories. The SP requires that you create subdirectories on the **/spdata** file system for storing critical PSSP data. Make sure the directories have the permissions `rwxr-sr-x`. Table 2 lists the required directories.

Table 2. Required /spdata Directories

Directory	Description	mkdir Command
/spdata/sys1/install/*name*/lppsource	Location of required AIX file sets	mkdir -p /spdata/sys1/install/*name*/lppsource
/spdata/sys1/install/images	Location of all required AIX mksysb images	mkdir /spdata/sys1/install/images
/spdata/sys1/install/pssplpp/*code_version*	Location of all SP installp file sets	mkdir -p \ /spdata/sys1/install/pssplpp/*code_version*
/spdata/sys1/install/pssp	Location of NIM configuration data files	mkdir /spdata/sys1/install/pssp
name is the new lpp_source name for the nodes (such as aix433 if that is what you called the subdirectory with the AIX 4.3.3 lppsource). Keep in mind that the **setup_server** program looks for this name later on during the installation process. By default, it is set to the string “default,” so that if you use that as your subdirectory name, you do not have to change the name here. *code_version* is the version of code of the form PSSP-x.y. (such as, PSSP-3.4).		

Step 13: Define space for the NIM boot images

Before creating an AIX boot/install server, ensure that there is sufficient space in the root (/) file system or create a separate file system for **/tftpboot** to manage the space required for the boot images (approximately 25 MB per lppsource level supported) created by NIM. For example, to increase the size of the root file system by 25 MB, issue:

```
chfs -a size=+51200 /
```

If you want to create a logical volume and file system in either the **rootvg** or **/spdatavg** volume group, follow the instructions in “Step 11: Define space for the /spdata directory” on page 20.

Step 14: Copy the AIX LP images and other required AIX LPs and PTFs

You must copy the AIX file sets into an lppsource directory on the hard drive of your control workstation. For AIX 4.3.3, the files go into **/spdata/sys1/install/***name***/lppsource**.

In AIX 5L 5.1, changes were made to allow installation with installers other than **installp** to allow new installation media formats. With AIX 5L 5.1, two new commands (**geninstall** and **gencopy**) were introduced, which call **installp** or **bffcreate**, or other commands as appropriate. Additional subdirectories have also been added into the NIM LPP_SOURCE with AIX 5L 5.1. For NIM, instead of just putting everything in the LPP_SOURCE directory, appropriate subdirectories are created by the **gencopy** and **bffcreate** commands and the images are copied to those subdirectories based on the format of the install package.

You must copy the AIX file sets under the **/spdata/sys1/install/***name***/lppsource** directory on your hard disk on the control workstation. The **bffcreate** command places the files into the appropriate subdirectory:

AIX 4.3.3 **/spdata/sys1/install/***name***/lppsource**

AIX 5L 5.1 **/spdata/sys1/install/***name***/lppsource/installp/ppc**
/spdata/sys1/install/*name***/lppsource/rpm/ppc**

You can download all of the AIX file sets (a very large number) or only the minimal required AIX file sets (approximately 500 MB for AIX 4.3.3 and approximately 1 GB for AIX 5L 5.1). Download the AIX file sets and the required AIX LPs into **/spdata/sys1/install/***name***/lppsource**. The AIX file sets and required AIX LPs must exist in this directory. Links to file sets in other directories are not allowed. If you change the path name in any way, the installation fails.

The following is the minimal list of AIX file sets required to perform mksysb installations. The *prefix.** syntax in the list refers to everything that starts with the *prefix*. For example, **devices.*** refers to all of the file sets starting with **devices**.

Minimal list of AIX 4.3.3 file sets:

```
Java.rte.*                  bos.diag.*
X11.apps.*                  bos.html.en_US.topnav.*
X11.base.*                  bos.mp.*
X11.compat.*                bos.net.*
X11.Dt.*                    bos.powermgt.*
X11.fnt.*                   bos.rte.*
X11.loc.*                   bos.sysmgt.*
X11.motif.*                 bos.terminfo.*
X11.vsm.*                   bos.up.*
```

```
X11.msg.*                   devices.*
bos                         perfagent
bos.64bit.*                 perl.*
bos.adt.*
```

Minimal list of AIX 5L 5.1 file sets:

```
IMNSearch.bld.*             bos.net.*
IMMSearch.rte.*             bos.perf.*
Java130.rte.*               bos.powermgt.*
Tivoli_Management_Agent.*   bos.rte.*
X11.Dt.*                    bos.svprint.*
X11.adt.*                   bos.sysmgt.*
X11.apps.*                  bos.terminfo.*
X11.base.*                  bos.txt.*
X11.compat.*                bos.up.*
X11.fnt.*                   devices.*
X11.loc.*                   ifor_ls.base.*
X11.motif.*                 invscout.ldb.*
X11.msg.*                   invscout.rte.*
X11.vsm.*                   perl.*
bos                         perfagent.tools.*
bos.64bit.*                 printers.rte.*
bos.adt.*                   rpm.rte.*
bos.diag.*                  rsct.*
bos.doc*                    sysmgmt.help.msg.en_US.*
bos.help.msg.en_US.*        sysmgmt.msg.en_US.websm.*
bos.html.en_US.topnav.*     sysmgmt.sguide.*
bos.iconv.*                 sysmgmt.websm.*
bos.loc.iso.*               xlC.aix50.*
bos.man.en_US.*             xlC.cpp.*
bos.mp.*                    xlC.msg.en_US.cpp.*
bos.mp64.*                  xlC.rte.*
box.msg.en_US.*
```

Additional files you may want to add to your lppsource:

bos.acct.* Required if you plan to use PSSP accounting

bos.cpr.* Required to install LoadLeveler 3.1 or Parallel Environment 3.2

dce.* The DCE file sets are required only if DCE will be configured by PSSP anywhere on the system. You will need the client portion of the DCE file sets because the installation code installs the DCE client code.

Java130.xml4j
Required for pSeries 690 servers

CIMOM Copy from the AIX toolbox for Linux applications CD. It is labeled in the Contents as openCIMOM, with the file name of Rpms/noarch/openCIMOM-0.61-1.aix5.1.noarch.rpm

Notes:

1. Refer to your disk usage planning in the “Combining the space requirements” section of *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* to determine if you have allocated enough space to accomplish this task.
2. Allow at least 1-3 hours for moving all the file sets from media to disk.

To copy the AIX LP images, login to the control workstation as root and run **bffcreate** using SMIT or the command line. The following example shows the product media on **cd0** and the selection of all LPs. Using **all** may load unnecessary file sets into the directory.

```
bffcreate -qvX -t/spdata/sys1/install/name/lppsource -d /dev/cd0 all
```

The following warning message is issued—ignore it:

```
bffcreate:
Warning: important size information is missing from
the table of contents file.  Consequently, there may not
be enough free file system space to successfully create
the bff image(s).  Continuing anyway...
```

If you choose not to use **bffcreate**, you need to run the **inutoc** script as follows:

```
cd /spdata/sys1/install/name/lppsource

inutoc .
```

Step 15: Copy the Correct level of PAIDE

The perfagent.server file set is part of the Performance Aide for AIX (PAIDE) feature of the Performance Toolbox for AIX (PTX), a separate product. This product provides the capability to monitor your SP system's performance, collects and displays statistical data for SP hardware and software, and simplifies runtime performance monitoring of a large number of nodes.

The Performance Toolbox for AIX, Agent Component (PAIDE) is required. The correct level of AIX PAIDE (perfagent.tools) needs to be installed on the control workstation and copied to all of the lppsource directories. The perfagent.tools file set is part of AIX 4.3.3 and AIX 5L 5.1.

The required level of perfagent is dependent upon the level of AIX and PSSP as shown in the following table:

Table 3. perfagent File Sets

AIX Level	PSSP Level	Required File Sets
AIX 4.2.1	PSSP 2.4	perfagent.server 2.2.1.x, where *x* is greater than or equal to 2
AIX 4.3.3	PSSP 2.4	perfagent.server 2.2.33.*
AIX 4.3.3	PSSP 3.1.1	perfagent.tools 2.2.33.*
AIX 4.3.3	PSSP 3.2	perfagent.tools 2.2.33.*
AIX 4.3.3	PSSP 3.4	perfagent.tools 2.2.33.*
AIX 5L 5.1	PSSP 3.4	perfagent.tools 5.1.0.*

Verify that your lppsource contains the correct level of PAIDE.

Refer to the *READ THIS FIRST* document for the latest information on PAIDE levels.

Note that important PTFs for perfagent.server are distributed on the AIX Update CD-ROM. The level of PAIDE copied to each lppsource directory must match the level of AIX in that directory.

Task B. Install PSSP on the control workstation

This section describes the steps you take to install PSSP on the control workstation. After you prepare the control workstation, you are ready to install the PSSP software.

Step 16: Copy the PSSP images

The RS/6000 SP package is comprised of these install images and file sets:

Image	Description
pssp.installp	Contains the PSSP install file sets
rsct.basic	Contains the RS/6000 Cluster Technology Availability Subsystems (for use with AIX 4.3.3)
rsct.clients	Contains the RS/6000 Cluster Technology Availability Subsystems (for use with AIX 4.3.3)
rsct.core	Contains the RS/6000 Cluster Technology Availability Subsystems (for use with AIX 4.3.3)
ssp.resctr	Contains the SP Resource Center
ssp.vsdgui	Contains the IBM Virtual Shared Disk Perspective
vsd.cmi	Contains the IBM Virtual Shared Disk Centralized Management Interface
vsd.hsd	Contains the IBM Virtual Shared Disk Hashed Shared Disk
vsd.rvsd.hc	Contains the IBM Recoverable Virtual Shared Disk Connection Manager
vsd.rvsd.rvsdd	Contains the IBM Recoverable Virtual Shared Disk Connection Daemon
vsd.rvsd.scripts	Contains the IBM Recoverable Virtual Shared Disk Recovery Scripts
vsd.sysctl	Contains the IBM Virtual Shared Disk sysctl commands
vsd.vsdd	Contains the IBM Virtual Shared Disk device driver

The RS/6000 SP package also contains several PSSP prerequisites files. They are:

ipfx	IBM Information Presentation Facility/6000
vacpp.cmp	Contains the VisualAge C++ Compiler
vacpp.ioc	Contains the VisualAge C++ IBM Open Class library runtime file set
xlC.aix43	Contains the VisualAge C++ Runtime file sets specific to AIX 4.3
xlC.aix50	Contains the VisualAge C++ Runtime file sets specific to AIX 5L 5.1
xlC.rte	Contains the VisualAge C++ Runtime file sets

Before you install the PSSP images on the control workstation, you first need to copy the images from the installation media to **/spdata/sys1/install/pssplpp/PSSP-3.4** directory on your hard disk.

Step 16.1: Copy PSSP images from media

Login to the control workstation as root and run **bffcreate** using SMIT or the command line.

If using:	Do this:	
SMIT	**TYPE**	**smit bffcreate** • The Copy Software to Hard Disk for Future Installation window appears.
	PRESS	**List** (F4) to show the available devices. Select the device containing the product installation media.
	PRESS	**Ok** to display the target parameters. • The Copy Software to Hard Disk for Future Installation window appears.
	TYPE	**/spdata/sys1/install/pssplpp/PSSP-3.4** in the `DIRECTORY for storing software` field.
	PRESS	**Ok** to begin the install process.
bffcreate	This example shows the product media on **cd0**. Enter: `bffcreate -d /dev/cd0 -t /spdata/sys1/install/pssplpp/PSSP-3.4 -X all`	

The following warning message is issued—ignore it:
`bffcreate:`
`Warning: important size information is missing from`
`the table of contents file. Consequently, there may not`
`be enough free file system space to successfully create`
`the bff image(s). Continuing anyway...`		

Step 16.2: Update the image table of contents (.toc)

When **bffcreate** completes, rename **ssp.3.4.0.0.I**, **rsct.clients.1.2.1.0.I**, **rsct.basic.1.2.1.0.I**, **rsct.core.1.2.1.0.I**, in **/spdata/sys1/install/pssplpp/PSSP-3.4**.

Enter the following:

```
cd /spdata/sys1/install/pssplpp/PSSP-3.4
mv ssp.3.4.0.0.I pssp.installp
mv rsct.basic.1.2.1.0.I rsct.basic
mv rsct.clients.1.2.1.0.I rsct.clients
mv rsct.core.1.2.1.0.I rsct.core
inutoc .
```

Step 16.3: Move prerequisite files

Several PSSP prerequisite files that are shipped on the PSSP media must be moved to your AIX lppsource.

Enter the following:

```
cd /spdata/sys1/install/pssplpp/PSSP-3.4
```

If your lppsource is for **AIX 4.3.3**, the following prerequisite files must be copied:

```
cp xlC.rte.* /spdata/sys1/install/name/lppsource
cp xlC.aix43.* /spdata/sys1/install/name/lppsource
cp ipfx.* /spdata/sys1/install/name/lppsource
cp vacpp.ioc.* /spdata/sys1/install/name/lppsource
cp vacpp.cmp.* /spdata/sys1/install/name/lppsource
cd /spdata/sys1/install/name/lppsource
inutoc .
```

If your lppsource is for **AIX 5L 5.1**, the following prerequisite files must be copied:

```
cp xlC.rte.* /spdata/sys1/install/name/lppsource/installp/ppc
cp xlC.aix5.* /spdata/sys1/install/name/lppsource/installp/ppc
cp ipfx.* /spdata/sys1/install/name/lppsource/installp/ppc
```

```
cp vacpp.ioc.* /spdata/sys1/install/name/lppsource/installp/ppc
cp vacpp.cmp.* /spdata/sys1/install/name/lppsource/installp/ppc
cd /spdata/sys1/install/name/lppsource/installp/ppc
inutoc .
```

Remove the prerequisite files from the PSSP lppsource directory since they have been moved to the AIX lppsource directories.

```
cd /spdata/sys1/install/pssplpp/PSSP-3.4
rm xlC*
rm ipfx*
rm vacpp*
inutoc.
```

If you never intend to install any nodes with AIX 4.3.3, you can also remove the RSCT files from the PSSP lppsource directory and then rerun inutoc.

Step 17: Copy a basic AIX (mksysb) image

Note that there is no root password in the basic (minimal) AIX/6000 SP mksysb image. If you choose to use this image (it is the default) to install your nodes, you should take appropriate steps to make the system more secure. If your site uses NIS, you can use the **firstboot.cust** file to define the NIS client. If you are not using NIS, you can use the **script.cust** file to copy the **/etc/passwd** and **/etc/security/passwd** files from the boot/install server. Refer to the example in the **/usr/lpp/ssp/samples/firstboot.cust** file to determine how to copy a file.

Note: In order to reinstall your nodes, the mksysb image and the lppsource that you use must both contain the same version, release, modification, and fix levels of AIX. If you do not have a mksysb image at the same level as your lppsource, you may do one of the following:

1. Make your own updated mksysb image. In order to do this, you will need to:
   a. Update an existing lppsource to the most recent maintenance level of AIX.
   b. Perform a BOS node upgrade on a single node as described in "BOS node upgrade" on page 153 **or** follow the steps in "Installing updates on a per node basis" on page 264.
   c. Make a mksysb image of that node as described in "Installing updates through reinstallation" on page 266.
   d. Use the mksysb created in Step 1c along with your updated lppsource to install your remaining nodes.
2. Contact IBM Level 1 service to obtain an updated mksysb image.

The media shipped with the SP hardware contains the **spimg** installp image. This image contains one or more AIX mksysb images. You may install any of these images for use on your nodes or use mksysb images of your own. You need to only install the AIX images that you intend to use.

If you intend to use your own mksysb image, copy it to **/spdata/sys1/install/images** and continue with "Step 18: Install PSSP prerequisites" on page 29.

Note: If DCE is running on the host that the mksysb image is made from, you must first turn autostart off for the DCE daemons. To do this, issue:

```
config.dce -autostart off
```

then create the mksysb image.

If using:	Do this:
SMIT	**TYPE** **smit install_latest** • The Install Software window appears. **TYPE** The input device (press F4 and select a device). **PRESS** **Ok** to begin the install.
installp	Enter: `installp -a -d /dev/cd0 -X spimg`

Step 18: Install PSSP prerequisites

PSSP has prerequisites for certain file sets.

Step 18.1: Install bos.net files

Make sure that the bos.net (TCP/IP and NFS) and bos.net.uucp (for Kerberos V4 systems only) files are installed on your control workstation.

Step 18.2: Install the perfagent.tools file set

Make sure that the perfagent.tools file set, which is part of AIX 4.3.3 or later, is installed on your control workstation. This file should have been placed in the lppsource directory in “Step 15: Copy the Correct level of PAIDE” on page 25. If it is not already installed on the control workstation, it should be installed now.

If using:	Do this:
SMIT	**TYPE** **smit install_latest** • The Install Software window appears. **ENTER** For AIX 4.3.3, **/spdata/sys1/install/***name***/lppsource** for Input Device **ENTER** For AIX 5L 5.1, **/spdata/sys1/install/pssplpp/***name***/lppsource/installp/ppc** for Input Device **PRESS** **Ok** to display the default install parameters. **PRESS** **List** for SOFTWARE to install to show options. **SELECT** One or more program options based on your AIX level as shown in the **installp** example that follows. **SELECT** Select program options as shown in the **installp** section of this table that follows. **PRESS** **Ok** to complete option selection and to begin installation. When the installation is complete, check the SMIT log file for the installation status. If errors occur, see *IBM AIX Problem Solving Guide and Reference*.

installp	For AIX 4.3.3, enter:
	`installp -agXd /spdata/sys1/install/name/lppsource \`
	`perfagent.tools`
	For AIX 5L 5.1, enter:
	`installp -agXd /spdata/sys1/install/name/lppsource/installp/ppc \`
	`perfagent.tools`

Step 18.3: Install the runtime files

PSSP has prerequisites for runtime libraries from the VisualAge C++ product.

For AIX 4.3.3, they are:

vacpp.ioc.aix43.rte 5.0.2.0

xlC.aix43.rte 5.0.2.0

For AIX 5L 5.1, they are:

vacpp.ioc.aix50.rte 5.0.2.0

xlC.aix50.rte 5.0.2.0

These file sets may not be part of the AIX installation package. These files and their associated prerequisites were placed in your AIX lppsource during “Step 16.3: Move prerequisite files” on page 27. They must be installed now.

There may be more recent levels of these files available. Please check the AIX Fix Distribution Service Web site at:

http://techsupport.services.ibm.com/rs6k/fixdb.html

If using:	Do this:
SMIT	**TYPE** **smit install_latest** • The Install Software window appears. **ENTER** For AIX 4.3.3, **/spdata/sys1/install/***name***/lppsource** for Input Device **ENTER** For AIX 5L 5.1, **/spdata/sys1/install/pssplpp/***name***/lppsource/installp/ppc** for Input Device **PRESS** **Ok** to display the default install parameters. **PRESS** **List** for SOFTWARE to install to show options. **SELECT** One or more program options based on your AIX level as shown in the **installp** example that follows. **SELECT** Select program options as shown in the **installp** section of this table that follows. **PRESS** **Ok** to complete option selection and to begin installation. When the installation is complete, check the SMIT log file for the installation status. If errors occur, see *IBM AIX Problem Solving Guide and Reference*.

installp	For AIX 4.3.3, enter:
	`installp -agXd /spdata/sys1/install/name/lppsource xlC.rte \`
	`xlC.aix43.rte vacpp.ioc.aix43.rte`
	For AIX 5L 5.1, enter:
	`installp -agXd /spdata/sys1/install/name/lppsource/installp/ppc \`
	`xlC.rte xlC.aix50.rte vacpp.ioc.aix50.rte`

Step 18.4: Install the RSCT files

If you are installing with AIX 5L 5.1, you must install the RSCT shipped with AIX 5L 5.1. You can skip this step if you are installing PSSP 3.4 on AIX 4.3.3.

If using:	Do this:
SMIT	**TYPE** **smit install_latest** • The Install Software window appears. **ENTER** **/spdata/sys1/install/***name***/lppsource/installp/ppc** for Input Device **PRESS** **Ok** to display the default install parameters. **PRESS** **List** for SOFTWARE to install to show options. **SELECT** One or more RSCT program options, or select the header files (called **rsct.basic**, **rsct.compat.basic**, **rsct.compat.clients**, and **rsct.core** with **ALL** on the far right side) to do the full installation. **PRESS** **Ok** to complete option selection and to begin installation. When the installation is complete, check the SMIT log file for the installation status. If errors occur, see *IBM AIX Problem Solving Guide and Reference*.
installp	Enter: `installp -agXd /spdata/sys1/install/name/lppsource/installp/ppc rsct`

Step 18.5: Install pSeries 690 files (optional)

Install the pSeries 690 files using the following command:

```
/bin/rpm -i openCIMOM-0.61-1.aix5.1.noarch.rpm
```

Step 19: Install PSSP on the control workstation

The PSSP images are made up of one or more file sets. Some of these file sets must be installed on the control workstation while others are optional. A subset of the file sets is installed on the individual nodes later in the installation process. Refer to the following table for more information.

Note: Do not create the **/usr/lpp/ssp** directory as a separate file system. It must be part of the **/usr** file system.

File sets installed on the control workstation

File set	Required on CWS	Description
rsct.basic.hacmp	Yes	**RS/6000 Cluster Technology** basic function (HACMP/ES for both AIX 4.3.3 and AIX 5L 5.1)
rsct.basic.rte	Yes	**RS/6000 Cluster Technology** basic function (HACMP/ES for both AIX 4.3.3 and AIX 5L 5.1)
rsct.basic.sp	Yes	**RS/6000 Cluster Technology** basic function (HACMP/ES for both AIX 4.3.3 and AIX 5L 5.1)
rsct.clients.hacmp	Yes	**RS/6000 Cluster Technology** client function (HACMP realm for AIX 4.3.3 only)
rsct.clients.rte	Yes	**RS/6000 Cluster Technology** client function (all realms for AIX 4.3.3 only)
rsct.clients.sp	Yes	**RS/6000 Cluster Technology** client function (SP realm for AIX 4.3.3 only)
rsct.compat.basic.hacmp	Yes	**RS/6000 Cluster Technology** Event Management basic function (AIX 5L 5.1 only)
rsct.compat.basic.rte	Yes	**RS/6000 Cluster Technology** Event Management basic function (AIX 5L 5.1 only)

File set	Required on CWS	Description
rsct.compat.basic.sp	Yes	**RS/6000 Cluster Technology** Event Management basic function (AIX 5L 5.1 only)
rsct.compat.clients.hacmp	Yes	**RS/6000 Cluster Technology** Event Management client function (AIX 5L 5.1 only)
rsct.compat.clients.rte	Yes	**RS/6000 Cluster Technology** Event Management client function (AIX 5L 5.1 only)
rsct.compat.clients.sp	Yes	**RS/6000 Cluster Technology** Event Management client function (AIX 5L 5.1 only)
rsct.core.auditrm	Yes	**RS/6000 Cluster Technology** Audit Log Resource Manager (AIX 5L 5.1 only)
rsct.core.errm	Yes	**RS/6000 Cluster Technology** Event Response Resource (AIX 5L 5.1 only)
rsct.core.fsrm	Yes	**RS/6000 Cluster Technology** File System Resource (AIX 5L 5.1 only)
rsct.core.gui	Yes	**RS/6000 Cluster Technology** Graphical User Interface (AIX 5L 5.1 only)
rsct.core.hostrm	Yes	**RS/6000 Cluster Technology** Host Resource Manager (AIX 5L 5.1 only)
rsct.core.rmc	Yes	**RS/6000 Cluster Technology** Resource Monitoring and Control Host Resource Manager (AIX 5L 5.1 only)
rsct.core.sec	Yes	**RS/6000 Cluster Technology** Security (AIX 5L 5.1 only)
rsct.core.sr	Yes	**RS/6000 Cluster Technology** Registry (AIX 5L 5.1 only)
rsct.core.utils	Yes	**RS/6000 Cluster Technology** Utilities (for both AIX 4.3.3 and AIX 5L 5.1)
rsct.msg.EN_US.*	Yes	**RS/6000 Cluster Technology** Message files associated with the other rsct.* file sets (AIX 5L 5.1 only)
rsct.msg.en_US.*	Yes	**RS/6000 Cluster Technology** Message files associated with the other rsct.* file sets (AIX 5L 5.1 only)
ssp.authent	Yes, if CWS is Kerberos V4 authentication server	**SP Authentication Server** Contains the server code that provides Kerberos V4 ticket-granting services and utility commands
ssp.basic	Yes	**SP System Support Package** Code for installing and monitoring the SP, including: • SP System Monitor • SDR • Centralized Management Interface (CMI—the SMIT panels) • Installation and Configuration Commands • Distributed shell • Login control
ssp.cediag		SP CE Diagnostics

File set	Required on CWS	Description
ssp.clients	Yes	**SP Authenticated Client Commands** User authentication commands, sysctl, monitor command line interfaces, logging daemon, Resource Manager client library, **jm_status** command. **Note:** If you want to run PSSP DCE administration commands on a cell administrator workstation remote from the SP system, you need to install the **ssp.clients** file set and its prerequisites on that cell administrator workstation.
ssp.css	Yes, if switch	**SP Communication Subsystem Package** Device drivers and switch support including: • Switch initialization and reconfiguration • Software error detection • Switch clock API
ssp.docs		SP man pages, PDF files, and HTML files
ssp.gui	Yes	**SP Perspectives** GUI (Launch Pad, Hardware Perspective, Event Management Perspective)
ssp.ha_topsvcs.compat	Yes	Compatibility for ssp.ha and ssp.topsvcs clients
ssp.hacws		**SP High Availability Control Workstation** Includes scripts to create a backup control workstation, error notification object samples, error log templates, and verification programs
ssp.jm		**Resource Manager** If no nodes are running PSSP 2.4, do not install the **ssp.jm** file set. It should be installed only on the control workstation if there are nodes running PSSP 2.4, which use the Resource Manager functionality that was merged into LoadLeveler 2.1 and for PSSP 3.1.1.
ssp.msg.En_US.*		US English IBM-850 message file sets associated with the other ssp.* file sets
ssp.msg.en_US.*		US English ISO 8859-15 message file sets associated with the other ssp.* file sets
ssp.perlpkg	Yes	**SP PERL Distribution Package** Includes Perl4, and Perl5 links
ssp.pman		**SP Problem Management**
ssp.public		**Public Code Compressed Tar files** Including tar files for public domain code Perl, SUP, Tcl, TclX, Tk, and Expect
ssp.resctr.rte		**SP Resource Center** Front end interface to online documentation and resources
ssp.spmgr		**SP Extension Node SNMP Manager** Required for extension node support
ssp.st		**Job Switch Resource Table Services Package** Low-level application programming interface for loading, unloading, and querying the job switch resource table

File set	Required on CWS	Description
ssp.sysctl	Yes	**SP Sysctl Package** The Sysctl remote execution facility server, daemon, commands, and configuration files
ssp.sysman	Yes	**Optional System Management programs** SP Management Tools including: • User Management Support • File Collections • Accounting Support • NTP • Parallel management commands • Error log management
ssp.tecad		**SP HA TEC Event Adapter Package**
ssp.top	Yes, if switch	**SP Communication Subsystem Topology Package** The system partitioning configuration directory and files including the System Partitioning Aid.
ssp.top.gui		**SP System Partitioning Aid Perspective** GUI
ssp.ucode	Yes	**SP Supervisor Microcode Package**

Notes:

1. **ssp.authent** contains the parts required on a system that are used by the Kerberos V4 authentication server.

   You must install **ssp.authent** if the control workstation will be configured as a Kerberos V4 authentication server. If you are using MIT Kerberos V4 or Andrew File System (AFS) authentication services, **ssp.authent** is not required. You can install **ssp.authent** on any other RS/6000 SP system that is used as a Kerberos V4 authentication server. You cannot install **ssp.authent** if the system already has an MIT Kerberos V4 or AFS authentication server installed. If you want to use the SP authentication facilities, you must first remove the other authentication service.
2. If you are planning to set up system partitions, you must install **ssp.top**.
3. **ssp.spmgr** contains an SNMP manager. If you already have an SNMP manager running on the control workstation, UDP port 162 usage by the SP manager included as part of this file set must change. You will need to update the **spmgrd-trap** entry in the **/etc/services** file on the control workstation to specify an unused port. Any new port configured on the control workstation also has to be configured on the SNMP agents supporting dependent nodes.

For a complete list of file sets, refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment.*

File sets installed on the control workstation in later steps

File set	Required on CWS	Description
ssp.vsdgui		**IBM Virtual Shared Disk Perspective** GUI
vsd.cmi		**IBM Virtual Shared Disk Centralized Management Interface**
vsd.rvsd.hc		**IBM Recoverable Virtual Shared Disk Connection Manager**

File set	Required on CWS	Description
vsd.rvsd.rvsdd		**IBM Recoverable Virtual Shared Disk Connection Daemon**
vsd.rvsd.scripts		**IBM Recoverable Virtual Shared Disk Recovery Scripts**
vsd.sysctl		**IBM Virtual Shared Disk sysctl commands**
vsd.vsdd		**IBM Virtual Shared Disk device driver**

Note: You must install the IBM Virtual Shared Disk file sets in "Step 26: Complete IBM Virtual Shared Disk installation (optional)" on page 47.

Installation without AIX preinstalled

Because your system may not have AIX preinstalled on the nodes, you should add an install image to your list of installation options. You can install one of the mksysb images shipped with the PSSP package.

Or if you prefer, you can provide your own AIX image for installation on the nodes.

PSSP installation instructions

Login to the control workstation as root, install the file sets selected for the control workstation, and follow one of the procedures described in the following table.

Notes:

1. If you are installing PSSP on AIX 4.3.3, install RSCT in **/spdata/sys1/install/pssplpp/PSSP-3.4**.
2. Do not install RSCT from this directory if you are installing PSSP 3.4 on AIX 5L 5.1.

If using:	Do this:	
SMIT	**TYPE**	**smit install_latest** • The Install Software window appears.
	ENTER	**/spdata/sys1/install/pssplpp/PSSP-3.4** for Input Device
	PRESS	**Ok** to display the default install parameters.
	PRESS	**List** for SOFTWARE to install to show options.
	SELECT	One or more program options, or select the header file (called **ssp** with **ALL** on the far right side) to do the full installation. If installing AIX 4.3.3 file sets, you must also select RSCT. If installing AIX 5L 5.1 file sets, RSCT is is already installed on AIX 5L 5.1.
	PRESS	**Ok** to complete option selection and to begin installation.
	When the installation is complete, check the SMIT log file for the installation status. If errors occur, see *IBM AIX Problem Solving Guide and Reference*.	

If using:	Do this:

installp	You can use **installp** to install multiple file sets. For example, if installing PSSP on AIX 4.3.3:
	`installp -a -g -d /spdata/sys1/install/pssplpp/PSSP-3.4 -X ssp rsct`
	For example, if installing PSSP on AIX 5L 5.1:
	`installp -a -g -d /spdata/sys1/install/pssplpp/PSSP-3.4 -X ssp`
	Note: For AIX 4.3.3 or later, **installp** automatically commits the packaging file set when you specify the **-a** option.
	To list all of the options for **ssp**, enter:
	`installp -l -d /spdata/sys1/install/pssplpp/PSSP-3.4/pssp.installp`

SP administrative locale

When PSSP is installed on the control workstation, an SP administrative language is created. This locale is used on the SP to determine:

- The language that the data can be written in to the SDR
- The default AIX locale to set when installing a new node

This locale can also be used by some SP subsystems for locale-specific operations. It is not necessary for every node on the SP to operate in the same locale. Nodes can operate in a locale that is different from the SP administrative locale.

The SP administrative locale is initially set to the base AIX locale installed on the control workstation. This value can be changed at anytime using standard PSSP procedures for modifying site environment variables (see “Step 30: Enter site environment information” on page 49).

A related site environment variable is used to control the type of information that can be written to the SDR. This variable indicates whether only ASCII data can be written to the SDR (that is, data in the '00'x to '7F'x code range), or whether non-ASCII data is allowed.

Be careful when setting the SP system to allow non-ASCII data in the SDR. This should be done only if all nodes on the SP will be operating in the same locale and you have no future requirements to change the SP administrative locale. The base ASCII code range is available in all currently AIX-supported locales. Non-ASCII data written in one locale cannot be properly processed when operating in a different locale. Therefore, switching from one SP administrative locale to another is prohibited if the SDR contains non-ASCII data.

System language environment

PSSP runs in the base AIX locale for the machine. PSSP ships message catalogs only for en_US and En_US. Running in a locale for which a message catalog does not exist (including the C and POSIX locales) can result in text similar to the following embedded in messages:

```
Message not found
```

Refer to the “Considering AIX and PSSP in another language” section in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for additional information.

Installing the ssp.docs HTML files

The **ssp.docs** file set includes HTML files that contain online versions of the PSSP publications. Once you have installed the **ssp.docs** file set, the PSSP HTML publications will be located at **/usr/lpp/ssp/html**. Since other parts of PSSP link to the HTML publications, these files should not be moved from the **/usr/lpp/ssp/html** directory.

A sample index file, **/usr/lpp/ssp/html/psspbooks.html**, has also been provided. It shows you how to set up a single launching point from which users can access all of the online books.

Installing the RS/6000 SP Resource Center

The RS/6000 SP Resource Center (**ssp.resctr**) provides a single interface to all of the online SP documentation and information resources. It contains links to SP publications, READMEs, product information, performance information, Redbooks, white papers, education, and up-to-date service information.

When the SP Resource Center is run, it detects which documentation file sets are installed (**ssp.docs**, **LoadL.html.en_US**, **ppe.docs**, and **mmfs.gpfs**). The SP Resource Center contains links to documents that are locally installed, or if a document is not installed, the link points to the document on the IBM World Wide Web site. If you are unsure that you have access to the World Wide Web, the documentation file sets should be installed to allow you to view them from the SP Resource Center.

The SP Resource Center consists of HTML, Java, and JavaScript. The files are installed in **/usr/lpp/ssp/resctr**.

The SP Resource Center does not have any requisites to other PSSP file sets, so it may be installed on any machine that is running AIX Version 4.2.1 or later. You must have the Netscape Navigator Version 4 or later to run the SP Resource Center. The SP Resource Center can also be run from a CD-ROM that can be used on AIX, or on the Microsoft Windows 95, 98, or NT platforms.

Once the SP Resource Center is installed, you can invoke it by issuing:

```
/usr/lpp/ssp/bin/resource_center
```

You can also invoke the SP Resource Center by selecting its icon from the CDE Desktop or by selecting its icon from the Perspectives Launch Pad. The first time you invoke the SP Resource Center, you will be prompted to enter the path name to the Netscape Navigator. This path name is stored on a per-user basis in **$HOME/.resctr**.

Step 20: Set authentication methods for AIX remote commands on the control workstation

When filling out your worksheet in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*, you decided which types of authentication methods you wanted to use on your SP system. You must select one or more authentication method for the control workstation. Your choices are **k5**, **k4**, or **standard**. This setting is used to determine initial security settings for PSSP in “Step 25: Complete system support installation on the control workstation” on page 46 when the **install_cw** script is run.

Valid authentication settings for AIX remote commands are:

If using:	Do this:
DCE	Enter: `chauthent -k5`
Kerberos V4	Enter: `chauthent -k4`
Standard AIX	Enter: `chauthent -std`
DCE and Kerberos V4	Enter: `chauthent -k5 -k4`
DCE and Standard AIX	Enter: `chauthent -k5 -std`
Kerberos V4 and Standard AIX	Enter: `chauthent -k4 -std`
DCE, Kerberos V4, and Standard AIX	Enter: `chauthent -k5 -k4 -std`

Notes:

1. If you are using Kerberos V4 and the primary Kerberos V4 server is on an external system, issue the **chauthent** command to ensure **rsh** can be issued by the external Kerberos V4 server to the control workstation. For example, issue:
   ```
   chauthent -k4 -std
   ```
2. After issuing the **chauthent** command, you can verify that your authentication setting is accurate by issuing the **lsauthent** command.

Step 21: Initialize RS/6000 SP Kerberos V4 (optional)

Prior to performing this step, you must have decided what type of Kerberos V4 authentication server to use: RS/6000 SP, AFS, or another MIT Kerberos V4 implementation. In preparation, you should have completed the checklist in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*. See that book for more information.

RS/6000 SP authentication provides a program, **/usr/lpp/ssp/bin/setup_authent**, to initialize RS/6000 SP authentication services on RS/6000 SP workstations (including the control workstation) for Kerberos V4 authentication servers and authentication client systems. This program defines instances of the **hardmon** and **rcmd** authenticated services, and does one of the following:

1. Creates a primary Kerberos V4 authentication server and database
2. Creates a secondary Kerberos V4 authentication server and database
3. Configures the control workstation as a Kerberos V4 client
4. Initializes the control workstation or other RS/6000 SP workstations to use AFS authentication

Note the following when running **setup_authent**:

- <screenclear> shows where **setup_authent** clears the screen before displaying an explanation of the next part of the initialization procedure.
- *root* refers to the administration ID. It can be **root** or any name you choose. It must be a user in the system.

The procedure for completing this step varies, depending on the authentication configuration you select. Optionally, you can set up other workstations as secondary servers or client systems. Each configuration includes an example where the **setup_authent** command is invoked. Review the examples. Substitute the principal names and passwords on your system for the *DescriptiveTerms* shown in the examples, and use them to initialize the authentication services on your system.

If initializing as:	Refer to:
Primary Kerberos V4 Authentication Server	"Step 21.2: Initializing as the primary Kerberos V4 authentication server"
Secondary Kerberos V4 Authentication Server	"Step 21.1: Setting up an external primary server" and "Step 21.3: Initializing as a secondary Kerberos V4 authentication server" on page 41
Authentication Client System	"Step 21.1: Setting up an external primary server" and "Step 21.4: Initializing as an authentication client system" on page 42
Use AFS Authentication	"Step 21.5: Initializing to use AFS authentication" on page 43
Select only one authentication step to follow. Do not perform the steps for the Kerberos V4 authentication server you did not choose.	

Step 21.1: Setting up an external primary server

Perform the following tasks to set up a primary Kerberos V4 server as an external workstation (not the control workstation).

1. Install the **ssp.authent** file set, if you have not already done so.
2. Set up the configuration file on the workstation. See the section on "Creating the Kerberos V4 configuration files" in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more information.
3. Run the **setup_authent** command on the external server.

After performing these tasks, follow the instructions in either "Step 21.3: Initializing as a secondary Kerberos V4 authentication server" on page 41 or "Step 21.4: Initializing as an authentication client system" on page 42.

Step 21.2: Initializing as the primary Kerberos V4 authentication server

Follow this procedure to initialize your primary Kerberos V4 authentication server on the RS/6000 SP control workstation or another RS/6000 SP system:

1. Create a **/etc/krb.conf** file, unless you want **setup_authent** to create a default configuration file with one Kerberos V4 authentication server for the default local realm name.
2. Create a **/etc/krb.realms** file, if you need to map any domains to the local realm name.
3. Run the **/usr/lpp/ssp/bin/setup_authent** program.

For more information, see "Installing and configuring Kerberos V4" on page 181.

The following example shows the interaction you can expect when you run **setup_authent** when initializing the primary Kerberos V4 authentication server.

```
#setup_authent
<screenclear>
********************************************************************
Creating the Kerberos Database

Invoking the kdb_init and kstash utilities to create the database.
```

```
You must decide on a master password for the database. You will be
prompted to enter it twice.  Save this password in a very secure
place, since it is used to encrypt all keys in the database and you
will need it for other administrative tasks.

After you complete this task, the Kerberos daemons will be started:
kerberos for ticket-granting services, kadmind for administration.

For more information see the kdb_init and kstash man pages.
********************************************************************
You will be prompted for the database Master Password.
It is important that you NOT FORGET this password.

Enter Kerberos master key:  YourDatabasePassword

Enter Kerberos master key:  YourDatabasePassword

0513-004 The Subsystem or Group, kerberos, is currently inoperative
0513-083 Subsystem has been Deleted
0513-071 The kerberos Subsystem has been added
0513-059 The kerberos Subsystem has been started.  Subsystem PID is 18394
<screenclear>
********************************************************************
Defining an Administrative Principal to Kerberos

The kdb_edit utility is used to define the initial Kerberos users.
You must define a user whose UID is 0 as a Kerberos database
administrator. This user will have to login to Kerberos with this name
prior to performing installation tasks that result in execution of the
setup_server command, during installation or whenever network
interfaces have been added or renamed in the SP system configuration.

kdb_edit prompts you separately for the name and the instance. First
enter the user name, specifying the login name of the user who will be
the primary Kerberos administrator for the local realm. When you are
prompted for the instance, you must enter admin.  You must assign a
Kerberos password for this user and enter it twice (you may use the
AIX login password). To take default values on other options, hit
<Enter>.

You may create any number of other Kerberos principals at this time.
To exit kdb_edit, hit <Enter> when prompted for another
principal name.

For more information see the kdb_edit man page.
********************************************************************

Opening database...
Previous or default values are in [brackets];
hit <enter> to leave the same, or new value.

Principal name: root
Instance: admin

<not found>, Create [yes]?  <Enter>

Principal: root, Instance: admin, kdc_key_ver: 1
New Password: <password>
Verifying, please re-enter
New Password: <password>

Principal's new key version = 1
Expiration date (enter yyyy-mm-dd) [2037-12-31] ?<Enter>
Max ticket lifetime [ 255 ] ? <Enter>
Attributes [ 0 ] ? <Enter>
Edit O.K.
Principal name: <Enter>
```

```
********************************************************************************

Logging into Kerberos as an admin user
You must assume the role of a Kerberos administrator <user>.admin
to complete the initialization of kerberos on the local system. The
k4init command is invoked and will prompt you for the password. If you
are setting up your primary server here, you just defined it. If you
have defined multiple administrative principals, or if your primary
authentication server is on another system, you must first enter the
name of an administrative principal who has root privilege (UID 0).
You need to be authenticated as this administrator so that this program
can create the principals and service keyfiles for the authenticated
services that run on the SP system. For more information, see the
k4init man page.
********************************************************************************
Kerberos Initialization for "root.admin"
Password: rootPassword
```

Step 21.3: Initializing as a secondary Kerberos V4 authentication server

- To do this step, the primary Kerberos V4 authentication server must already be initialized.
- Ensure clock synchronization between the primary and secondary Kerberos V4 authentication servers.

For more information, see “Installing and configuring Kerberos V4” on page 181.

Follow this procedure to initialize a secondary Kerberos V4 authentication server on the control workstation or another RS/6000 SP workstation.

1. Copy the **/etc/krb.conf** file from the primary Kerberos V4 authentication server to this secondary Kerberos V4 server system.
2. Add a line to the **/etc/krb.conf** file, listing this system (by its full host name) as a secondary Kerberos V4 server for the local authentication realm.

   For example, to add **sp2cw.xyz.com** as a secondary Kerberos V4 server for the authentication realm **XYZ.COM**, add this line to **/etc/krb.conf**:

   ```
   XYZ.COM sp2cw.xyz.com
   ```
3. Copy the **/etc/krb.realms** file from the primary Kerberos V4 server to this secondary Kerberos V4 server system.
4. Run the **setup_authent** program.

   setup_authent requires you to login to the authentication service using the same administrative principal name that was defined for the primary Kerberos V4 server. The remainder of the initialization of authentication services on this secondary local Kerberos V4 system takes place automatically.
5. After **setup_authent** completes, add an entry for the new server to the **/etc/krb.conf** file on all SP systems on which you have already initialized authentication.
6. On the primary Kerberos V4 server, if this is the first secondary Kerberos V4 server, you should create a root **crontab** entry that invokes the script **/usr/kerberos/etc/push-kprop** to periodically propagate database changes.

The following example shows the interaction you can expect when you run **setup_authent** when initializing as a secondary Kerberos V4 authentication server:

```
#setup_authent
<screenclear>
********************************************************************************
Logging into Kerberos as an admin user
```

```
You must assume the role of a Kerberos administrator <user>.admin
to complete the initialization of kerberos on the local system.
The k4init command is invoked and will prompt you for the password.
If you are setting up your primary server here, you just defined it.
If your primary server is on another system, you must first enter
the user name of an administrative principal defined on that server.

You need to be authenticated as an administrator so that this
program can create the service principals required by the
authenticated services that are included in the ssp package.

hardmon - for the System Monitor facilities
rcmd    - for sysctl and Kerberos-authenticated rsh and rcp

For more information, see the k4init man page.
********************************************************************
setup_authent: Enter name of admin user: root
Kerberos Initialization for "root.admin"
Password: rootPassword
backup.abc.com: success.backup.abc.com:         Succeeded
#
```

The last two messages shown in the previous example are issued by the programs that transfer the database from primary to secondary Kerberos V4 servers, to indicate that the backup database has been installed.

Step 21.4: Initializing as an authentication client system

To do this step, the primary Kerberos V4 authentication server must already be initialized.

For more information, see “Installing and configuring Kerberos V4” on page 181.

Follow this procedure to initialize the control workstation or another RS/6000 SP system as an authentication client system.

1. Copy the **/etc/krb.conf** file from the primary Kerberos V4 authentication server to this system.
2. Copy the **/etc/krb.realms** file from the primary Kerberos V4 server.

   If the new workstation is outside the realm of the primary server, you must add this new workstation to the **/etc/krb.realms** file on the primary Kerberos V4 server before you copy the **/etc/krb.realms** file from the primary Kerberos V4 server to the new workstation.
3. Run the **setup_authent** program.

   setup_authent requires you to login to the authentication service using the same administrative principal name that was defined when the primary Kerberos V4 server was set up.
4. The **.klogin** file on a client workstation contains just the administrative principal name you used to install authentication. You may want to edit the **/spdata/sys1/spsec/.klogin** file to add other principals in your configuration.

The following example shows the interaction you can expect when you run **setup_authent** when initializing as an authentication client system. The initial warning message shown in the example is issued if you have installed the **ssp.authent** option on a system configured as a client rather than a server.

```
#setup_authent
setup_authent: This system is not listed as a Kerberos server in
/etc/krb.conf. Continuing setup as a Kerberos client system only.
<screenclear>
********************************************************************
Logging into Kerberos as an admin user
```

```
You must assume the role of a Kerberos administrator <user>.admin
to complete the initialization of kerberos on the local system.
The k4init command is invoked and will prompt you for the password.
If you are setting up your primary server here, you just defined it.
If your primary server is on another system, you must first enter
the user name of an administrative principal defined on that server.

You need to be authenticated as an administrator so that this
program can create the service principals required by the
authenticated services that are included in the ssp package.

hardmon - for the System Monitor facilities
rcmd    - for sysctl and Kerberos-authenticated rsh and rcp

For more information, see the k4init man page.
********************************************************************************
```

setup_authent: Enter name of admin user: *root*

```
Kerberos Initialization for "root.admin"
```

Password: *rootPassword*

Step 21.5: Initializing to use AFS authentication

To do this step, the AFS primary authentication server must already be initialized.

For more information, see “Installing and configuring Kerberos V4” on page 181.

Follow this procedure to initialize using AFS authentication servers.

1. If the AFS configuration files (**ThisCell, CellServDb**) are not in **/usr/vice/etc**, you must create a symbolic link from **/usr/vice/etc** to the directory containing those files:
2. If the **kas** command is not installed in **/usr/afsws/etc**, create a symbolic link from **/usr/afsws/etc** to the directory containing the **kas** command.
3. If you are using AFS Version 3.4, you must reconcile the conflicting port assignments used by the kaserver and RS/6000 SP authentication commands, as described in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* and *PSSP: Administration Guide*.
4. Run the **setup_authent** program.

   setup_authent requires you to enter the name and password of the AFS administrator.

The following example shows the interaction you can expect when you run **setup_authent** when initializing to use AFS authentication. The message always appears when the workstation has AFS installed, either as a client or server:

```
#setup_authent
<screenclear>
********************************************************************************
Option to Use AFS

Because this system is configured for use of AFS, you may choose to
use the AFS authentication servers instead of installing RS/6000 SP
authentication servers or using other Kerberos V4 servers.

The choice of AFS indicates that you will be using AFS authentication
servers exclusively in your RS/6000 SP system's local realm.
```

```
Do you want to set up authentication services to use AFS servers?
**********************************************************************

Enter y or n: y

afs_add_principal: Enter afs admin principal name [login-name] user-name
Password: UserNamePassword
```

Step 22: Configure DCE for the control workstation (required for DCE)

Restrictions

1. You cannot use both DCE authentication and HACWS.
2. You cannot use IPv6 aliasing with DCE, HACMP, and HACWS.

If you want PSSP to use DCE authenticated services, you must:

1. Install DCE on the control workstation.

   If you plan to install DCE on the control workstation, become familiar with "Tips for installing DCE on the SP." If you modify the **/etc/environment** file, you will need to reboot the control workstation in order for the DCE processes to use those changes.

   Tips for installing DCE on the SP

   DCE will use all configured network interfaces available for any DCE runtime traffic. There may be circumstances where certain network interfaces or addresses should not be used. DCE provides a mechanism to exclude these interfaces or adapters. Excluding these interfaces does not preclude their use for remote command traffic.

   DCE accomplishes this through the use of environment variables. These are: **RPC_UNSUPPORTED_NETADDRS** and **RPC_UNSUPPORTED_NETIFS**. The two variables accomplish the same task, so only use one of these variables. The recommended value to use is **RPC_UNSUPPORTED_NETIFS**.

   Within the SP, there are specific adapters or interfaces, like the switch (css*#*) adapters, which do not communicate between the control workstation and the nodes. These adapters are prime candidates for exclusion from DCE traffic.

   For example, to exclude the switch adapter, css0, do one of the following:

   - Edit the **/etc/environment** file on all nodes and add **RPC_UNSUPPORTED_NETIFS=css0**.
   - On the command line, enter **export RPC_UNSUPPORTED_NETIFS=css0**

   Start DCE within the same session the previous command was entered. If there are adapters on the control workstation, through which no DCE communication is expected, exclude these adapters as well using the same method described previously.

2. Be configured either as a client or server in the cell.

Step 22.1: Update the spsec_overrides file (optional)

The **config_spsec** command reads from two files. The defaults file is **/usr/lpp/ssp/config/spsec_defaults**. If the defaults need to be modified, for example, if any of the names in **spsec_defaults** conflict with items already in the DCE database, the **/spdata/sys1/spsec/spsec_overrides** file should be modified.

Note: If the **spsec_overrides** file has been modified on the control workstation, it must be copied to the remote workstation in order to run the **config_spsec** command off of the SP.

For more information, refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* and *PSSP: Command and Technical Reference*.

Step 22.2: Create DCE groups, organizations, principals, and accounts

As the cell administrator on the control workstation, issue the following command to create SP Trusted Services groups, organizations, and principals for the control workstation:

```
config_spsec -c -v
```

Note: Refer to the **config_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag to run this command remotely off of the SP.

Step 22.3: Create SP administrative principals

There must be a DCE principal that is a member of **hm-admin**, **sdr-admin**, **sdr-system-class-admin**, **sdr-restricted**, **spsec-admin**, and the **hm-control** groups to continue the install.

Use the appropriate DCE commands to define an administrative principal. The principal can be added to the SP access groups by a cell administrator using **dcecp**:

```
dcecp -c group add sdr-admin -member your_principal
dcecp -c group add hm-admin -member your_principal
```

Note: The administrative principals may need access to additional SP groups. Refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for a complete list of groups defined by PSSP to DCE. The access groups (**ACC-GRP**) whose name does not end in "**-services**" are intended for end users. For example, these control facilities could be sysctl, problem management, event management, the switch commands, LoadLeveler, Parallel Environment, and so on.

Step 22.4: Create control workstation-specific keyfiles

As root on the control workstation with default credentials, issue the following command to create control workstation-specific keyfiles:

```
create_keyfiles -c -v
```

Step 23: Set the authentication method for SP Trusted Services on the control workstation

Depending on the authentication method you selected in either "Step 21: Initialize RS/6000 SP Kerberos V4 (optional)" on page 38 or "Step 22: Configure DCE for the control workstation (required for DCE)" on page 44, determine the appropriate authentication method to use for SP Trusted Services during installation.

Notes:

1. If the authentication methods enabled for use by SP Trusted Services includes DCE, the authentication methods enabled for use by the AIX remote commands must include Kerberos V5.
2. If the authentication methods enabled for use by SP Trusted Services includes compatibility, the authentication methods enabled for use by the AIX remote commands must include Kerberos V4.

If using:	Do this:
DCE	Enter: `chauthts dce`
Kerberos V4	Enter: `chauthts compat`
Both DCE and Kerberos V4	Enter: `chauthts dce compat`
None	Enter: `chauthts`

To verify your settings, issue the **lsauthts** command. If your setting was DCE, **DCE** will be returned. If your setting was Kerberos V4, **Compatibility** will be returned.

Step 24: Obtain credentials

If DCE or Kerberos V4 was enabled in “Step 23: Set the authentication method for SP Trusted Services on the control workstation” on page 45, you must obtain credentials using **dce_login** or **k4init**. If DCE was selected, you should **dce_login** to the SP administrative principal created in “Step 22.3: Create SP administrative principals” on page 45. If Kerberos V4 was selected, you should use the appropriate administrative principal created in “Step 21: Initialize RS/6000 SP Kerberos V4 (optional)” on page 38.

Step 25: Complete system support installation on the control workstation

This step does the following:

- Configures the control workstation
- Installs PSSP SMIT Panels
- Starts SP daemons
- Configures the SDR
- Updates **/etc/inittab** and **/etc/services**
- Sets *node_number* for the control workstation to 0 in the Object Data Management (ODM) databases
- Creates the **hardmon hmacls** file
- Configures the system as one system partition. There is one system partition corresponding to the name of the control workstation in the SP object. It is called the **default** or **persistent** system partition because it always exists.
- Sets security attributes for the default system partition.

Using install_cw

Use the **install_cw** command to finish installing PSSP on the control workstation.

If using:	Do this:
install_cw	Enter: `install_cw`

There are certain conditions that can cause the **install_cw** command to fail. This will be shown by a message such as:

```
The SDR_init script completed unsuccessfully with a return code of 1.
Exiting...
```

Additional messages in **/var/adm/SPlogs/sdr/SDR_config.log** will provide more detailed information about the failure. Typical conditions that can cause a failure are:

1. The **lsauthts** command indicates that the **chauthts** command was never run. (The **chauthts** command should have been run in “Step 23: Set the authentication method for SP Trusted Services on the control workstation” on page 45.) To recover from this failure, run the **chauthts** command and rerun the **install_cw** command.
2. In a system where the **lsauthts** command indicates DCE, **install_cw** was invoked by a user lacking **sdr-system-class-admin** and **sdr-admin** authority. To recover from this failure, **dce_login** to a correctly-authorized principal and rerun the **install_cw** command.

Note: At this point, you can bring up the Perspectives GUI for use in the rest of the installation. The Launch Pad will let you bring up SMIT menus and issue commands, but the **sphardware** Perspective may not function at this stage of the installation. See “Appendix C. SP Perspectives tasks” on page 291 for more information on using Perspectives. If you do not wish to use Perspectives for the install steps, you can use SMIT or the command line.

To bring up the Perspectives Launch Pad, make sure your DISPLAY environment variable is set correctly and enter the following command:

```
perspectives &
```

You may receive the following message which you can ignore:

```
Warning: locale not supported by C library, locale unchanged.
```

Use the **splstdata** command to check the initial system partition security settings.

If using:	Do this:
splstdata	Enter: `splstdata -p`

Step 26: Complete IBM Virtual Shared Disk installation (optional)

Note: Perform this step only if you are installing IBM Virtual Shared Disk.

Use SMIT or the **installp** command to install the IBM Virtual Shared Disk file sets.

If using:	Do this:
SMIT	**TYPE** **smit install_latest** • The Install Software window appears. **ENTER** **/spdata/sys1/install/pssplpp/PSSP-3.4** for Input Device **PRESS** **Ok** to display the default install parameters. **PRESS** **List** for SOFTWARE to install to show options. **SELECT** Select **vsd** with **ALL** on the far right side to do the IBM Virtual Shared Disk installation. **PRESS** **Ok** to complete option selection and to begin installation. When the installation is complete, check the SMIT log file for the installation status. If errors occur, see *IBM AIX Problem Solving Guide and Reference*.

installp	You can use **installp** to install multiple file sets. For example:
	`installp -a -g -d /spdata/sys1/install/pssplpp/PSSP-3.4 -X vsd`
	Note: For AIX 4.3.3 or later, **installp** automatically commits the packaging file set when you specify the **-a** option.
	To list all of the options for IBM Virtual Shared Disk, enter:
	`installp -l -d /spdata/sys1/install/pssplpp/PSSP-3.4/vsd`

Step 27: Apply PSSP PTFs (optional)

Software maintenance (PTFs) may now be applied to the **ssp** and **rsct** file sets installed on the control workstation. Refer to “Installing program updates” on page 262 for planning considerations. Follow the instructions in “Preparing the control workstation” on page 263 to install the PTFs.

Step 28: Add the PSSP T/EC adapter (optional)

At this point, you can optionally add the PSSP T/EC adapter to your system. Refer to “Chapter 9. Installing the optional PSSP T/EC adapter” on page 279 for more information.

Step 29: Run SDR and System Monitor verification tests

If using:	Do this:	
Perspectives	**SELECT**	**smit SP_verify on CWS** from Launch Pad • The RS/6000 SP Installation/Configuration Verification menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE**	**smit SP_verify** • The RS/6000 SP Installation/Configuration Verification menu appears.
	SELECT	System Data Repository
	PRESS	**Done** to return to the previous screen
	SELECT	System Monitor Installation

SDR_test
spmon_itest
	Enter:
	`SDR_test`
	`spmon_itest`	

After the tests are run, the system creates the **spmon_itest.log** in **/var/adm/SPlogs/spmon** and the **SDR_test.log** in **/var/adm/SPlogs**.

See *PSSP: Command and Technical Reference* for more information about **SDR_test** and **spmon_itest** and on what these tests do.

Task C. Enter site environment, frame, node, switch, and security information

This section describes the steps you take to enter the information defining the SP configuration. After you prepare the control workstation and install the PSSP software, you are ready to enter the data to define your SP configuration.

Step 30: Enter site environment information

If you changed the lppsource *name* to something other than the default (you would have changed the name in “Step 12: Create the required /spdata directories” on page 22), you must perform this step.

If you do not want to change any of the default site environment variables, skip this step and continue with “Step 32: Enter SP or multiple NSB frame information and reinitialize the SDR” on page 54. To check the default site environment variables settings, issue the following command:

```
splstdata -e
```

Use Perspectives, SMIT, or the **spsitenv** command to enter information about your site environment. Use the worksheet in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* to enter site environment values.

The site environment data is written to the SDR. Before you run any of the installation scripts, you must enter the following data on the control workstation. Refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for additional explanation.

Site environment data includes:

- The name of the default network install image
- Your method of time service, the name of your time servers, and the version of NTP in use
- Whether you want to have the SP services configure and manage the Automounter
- User Admin information
- Whether you want to use RS/6000 SP User Management
- Whether you want RS/6000 SP File Collection Management installed and where the daemon will run

  Note: If the user administration interface is set to **false** and file collections is set to **true**, the following user and group files will not be updated:
  - **/etc/passwd**
  - **/etc/group**
  - **/etc/security/passwd**
  - **/etc/security/group**

  Refer to the "user.admin collection" section in the "Managing file collections" chapter of *PSSP: Administration Guide* for more information.
- Whether you want to use RS/6000 SP Accounting
- Whether you use the default lppsource directory as the location of the AIX file sets. If you changed the lppsource name from the default, you also need to change it in the Site Environment Data label.

  You must ensure that the AIX level of the LP source (indicated by the *cw_lppsource_name*) matches the AIX level installed on your control workstation.
- Whether you use the base AIX locale installed on the control workstation
- Whether ASCII-only data can be written to the SDR or whether non-ASCII data is allowed

  Note: If any of your nodes are running a version of PSSP earlier than PSSP 3.2, only ASCII data may be written to the SDR.
- SP model number and SP serial number

  The SP model number and SP serial number are used in gathering vital product data for the SP system. Entering these values is not required. If not entered at this time, the data will need to be provided the first time the **get_vpd** command is run to collect vital product data for the SP system.
- Whether you want to run with restricted root access (RRA) enabled, which precludes the root user from being able to issue the **rsh** and **rcp** commands from the nodes to the control workstation or to other nodes.

  Note: Some applications (such as GPFS, Problem Management, and HACMP) cannot run in this environment.
- Whether you want to run with a secure remote command method instead of the default **rsh** and **rcp** commands from the PSSP code. It also gives you the ability to enter the location of your remote command and remote copy executables.

Restricted root access must be enabled to use the secure remote command method. Also, the secure remote command software should be installed and running on the control workstation before enabling on the site environment menu. Once set in the SDR, the PSSP installation and configuration code, and CSS commands will automatically begin to use the secure remote commands instead of the **rsh** and **rcp** commands. Refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more information.

Three environment variables are defined so that the user can override the SDR settings for the Remote command method, the Remote command executable, and the Remote copy executable on any given PSSP command or script invocation. These are also used to drive the remote command method used by user commands such as **pcp** and **dsh** which only use these environment variables to determine the remote command method to use.

The environment variables are:

RCMD_PGM=secrshell or RCMD_PGM=rsh

DSH_REMOTE_CMD=*remote_command_executable*

REMOTE_COPY_CMD=*remote_copy_command_executable*

To set up your system to run with the secure remote command method:

From SMIT:

– Enter **smit enter_data**
– Select **Site Environment Information**

The following options that appear on the Site Environment Information menu for remote command access must be modified. The following are the defaults.

– Root remote command access restricted: false
– Remote command method: rsh
– Remote command executable:

  The default depends on the setting of the remote command method. It is not displayed when using SMIT.
– Remote copy executable:

  The default depends on the setting of the remote command method. It is not displayed when using SMIT.

For the root remote command access restricted option, select **true**. This enables the restricted root rcmd method on the system which is required to enable a secure remote command method.

For the remote command method, select **secrshell** which will enable the PSSP code to use the secure remote command method. If the Remote command executable and the Remote copy executable are left blank, the executable will default to **/bin/ssh** and **/bin/scp**. If this is not the location of your secure remote command program, enter the full name (including path) of the location of your secure remote command and secure remote copy program.

The following command can be used in place of the SMIT panels to perform the same function:

```
spsitenv restrict_root_rcmd=true,rcmd_pgm=secrshell, \
         dsh_remote_cmd=/bin/ssh,remote_copy_cmd=/bin/scp
```

You can enter the site environment information using Perspectives, SMIT, or the **spsitenv** command. Whichever method you chose, keep in mind that you can easily change these options at any time after installation is complete.

If using:	Do this:
Perspectives	**SELECT** **smit config_data on CWS** from the Launch Pad • The SP Configuration Database Management menu appears. **SELECT** Enter Database Information • The Enter Database Information menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit enter_data** • The Enter Database Information menu appears. **SELECT** Site Environment Information • The Site Environment Information window appears. **TYPE** Your environment choices. Refer to the Site Environment Worksheet. **PRESS** **Ok** to complete operation • If you do not make any changes to the information already in the window, SP displays a usage information message. This is normal—continue on.

spsitenv	This example configures NTP service as consensus and specifies that file collection be installed.
	`spsitenv ntp_config=consensus filecoll_config=true`
	This example specifies that the control workstation lppsource directory be **/spdata/sys1/install/aix433/lppsource** for you installation configuration.
	`spsitenv cw_lppsource_name=aix433`

Step 31: Enter Hardware Management Console (HMC) information (optional)

You must perform this step if your SP system or clustered enterprise server system will contain IBM @server pSeries 690 servers.

The hardware control and monitor functions for the pSeries 690 server are managed through a network connection from the control workstation to the hardware management console (HMC) that is controlling the pSeries 690 server. Install the HMC for your pSeries 690 server following the instructions in *pSeries 690 Installation Guide*. The following operations must be performed locally on the HMC before the pSeries 690 server is defined to PSSP. Refer to *Hardware Management Console for pSeries Operations Guide* for details on performing these operations. Review the *READ THIS FIRST* document that accompanies the PSSP installation media for information on required HMC product and PTF levels and the corresponding pSeries 690 hardware and software product and PTF levels.

Perform the following steps for each HMC and each pSeries 690 server:

1. Ensure that the HMC is installed and configured to operate on the SP Ethernet administrative LAN network. Use the HMC *System Configuration* interface to customize the network settings. Ensure the netmask is properly assigned and that all IP addresses have been registered with your name server and can be resolved. Note the IP address assigned to this HMC SP Ethernet administrative LAN connection. This information will be required when defining the HMC to the control workstation in the next section.

2. Use the HMC *User Management* interface to define a user ID with the role of *System Administrator* and assign a password. This information will be required when defining the HMC to the control workstation in the next section.
3. Ensure that the pSeries 690 server is recognized by the HMC. Use the HMC *Partition Management* interface to determine if the server is present. Follow problem resolution procedures in *Hardware Management Console for pSeries Operations Guide* if an entry for the server is not displayed on the interface.
4. Use the HMC *Partition Management* interface to view the properties for the managed system object. If desired, change the system name from the default name set by the HMC. Note the defined system name. This information will be required when entering the non-SP frame information for this server in "Step 33: Enter non-SP frame information and reinitialize the SDR (optional)" on page 56. If the system name is changed in the future, the new name will then also need to be changed in the non-SP frame information stored in the SDR on the control workstation.
5. Use the HMC *Partition Management* interface to select the desired power-on mode for your system: full system partition (SMP) mode, logical partition standby mode, or physical partition mode.
6. If you selected logical partition standby mode or physical partition mode, use the HMC *Partition Management* interface to create partitions and profiles as necessary. Partition objects must be created at this time, but the partitions do not need to be activated or installed before proceeding with control workstation operations. When a partition is activated from the control workstation, the default profile for the partition will be used. If you want to use a different profile, change the default profile setting for the partition using the HMC *Partition Management* interface.
7. View the properties for each partition object and note the partition ID. Each partition is represented in PSSP as a node. The partition ID will be used by PSSP to assign a corresponding SP slot number and node number to the SP node for that partition.

If you need to change your partition settings at a later time, refer to "Reconfiguring IBM @server pSeries 690 logical partitions (LPARs)" on page 245 for instructions on how to do this.

The following operations must be performed on the control workstation before entering non-SP frame information for your pSeries 690 servers.

1. Use the AIX **ping** command to verify that the control workstation has network connectivity to each HMC:

   ```
   ping hmc_hostname_or_ipaddr
   ```

   where *hmc_hostname_or_ipaddr* is the host name or IP address of the HMC as configured on the HMC *System Configuration* interface in Step 1 on page 52 of the previous section. If the command fails, review your network and name server configurations on both the control workstation and the HMC.
2. Define the previously-created HMC user ID to PSSP for **hardmon**. Running the following command once for each HMC:

   ```
   /usr/lpp/ssp/bin/sphmcid hmc_hostname_or_ipaddr hmc_sysadmin_userid
   ```

   where *hmc_hostname_or_ipaddr* is the host name or IP address of the HMC as configured on the HMC *System Configuration* interface in Step 1 on page 52 of the previous section and *hmc_sysadmin_userid* is the system administrator user ID created on the HMC *User Management* interface in Step 2 of the previous section. You will be prompted to enter the password for the

hmc_sysadmin_userid. You must run this command again anytime the password is changed for the *hmc_sysadmin_userid*.

3. Define the switch node numbers for your nodes. If this system has an SP Switch or is a switchless SP system, you must define an SP switch node number for each logical partition in the pSeries 690 server. Manually edit the **/etc/switch.info** file to include one entry for each logical partition in the attached server. See *PSSP: Command and Technical Reference* for details on editing this file. You can skip this step if the system has an SP Switch2 or is a switchless clustered enterprise server system.

   An example **/etc/switch.info** file might contain the following entries for a pSeries 690 server that will be defined as frame 5, with four LPARs, attached to switch 2 in the system:

   ```
   # Node_number     Switch_node_number
   65                16
   66                17
   67                18
   68                19
   ```

   If you are running your pSeries 690 server in SMP mode or will only be defining one LPAR and it is assigned partition ID 1, you can skip this operation and simply enter the switch node number when you enter the the other non-SP frame information in “Step 33: Enter non-SP frame information and reinitialize the SDR (optional)” on page 56.

Step 32: Enter SP or multiple NSB frame information and reinitialize the SDR

You must perform this step at least once for each set of frames or multiple node switch board (NSB) frames that you are adding to the system. You do not need to reinitialize the SDR until you are entering the last set of frames.

SP frames containing nodes must be numbered between 1 and 128 inclusive. This is to ensure that nodes will be numbered between 1 and 2047. Larger frame numbers, up to 250, can be used for frames that will contain only switches or SP expansion I/O units.

SP frames

This step creates frame objects in the SDR for each SP frame in your system. At the end of this step, the SDR is reinitialized, resulting in the creation of node objects for each node attached to your frames.

You can enter information about your frames using Perspectives, SMIT, or the **spframe** command. You must be an authenticated administrative user to issue this command.

If frames are not contiguously numbered, repeat this step for each series of contiguous frames. To save time, do not specify reinitialization of the SDR until you are entering the final series of contiguous frames.

If using:	Do this:	
Perspectives	**SELECT**	**smit config_data on CWS** from the Launch Pad • The SP Configuration Database menu appears.
	SELECT	Enter Database Information • The Enter Database Information menu appears.
	From this point, you can follow the rest of the SMIT steps described in the next row of this table.	
SMIT	**TYPE**	**smit enter_data** • The Enter Database Information menu appears.
	SELECT	SP Frame Information • The SP Frame Information window appears.
	TYPE	The start frame number and the number of frames in the Frame Count field. (Start frame defaults to **1**.) The starting frame tty port defaults to dev/tty0. You may need to change this depending upon your configuration.
	SELECT	**yes** or **no** next to `Re-initialize the System Data Repository`, as follows: **no** if you have more (noncontiguous) frame entries to make in this panel. **yes** if you are entering only one series of contiguous frames, or entering the last series of noncontiguous frames.
	SELECT	**no** to `Multiple NSB Frame (SP Switch2 only)`.
	SELECT	**no** to `Allow Frame Numbers greater than 128`.
	PRESS	**Ok** to enter frame data to the SDR.
spframe	Specify **spframe** command with **-r yes** to reinitialize the SDR (when running the command for final series of frames), a starting frame number, a frame count, and the starting frame's tty port. The following example enters information for four frames (frame 1 to frame 4) and indicates that frame 1 is connected to **/dev/tty0**, frame 2 to **/dev/tty1**, and so on, and reinitializes the SDR. `spframe -r yes 1 4 /dev/tty0`	

Multiple node switch board (NSB) frames (SP Switch2 only)

In PSSP 3.4, you can install multiple NSBs in an SP frame. A multiple NSB frame can only contain switches in slots 1 through 16. You cannot install SP nodes in a multiple NSB frame.

You can enter information for non-SP frames using Perspectives, SMIT, or the **spframe** command. If frames or tty ports are not all contiguously numbered, repeat this step for each series of contiguous information. To save time, do not specify the reinitialization of the SDR until you are entering the final series of contiguous frames.

If using:	Do this:
Perspectives	**SELECT** **smit config_data on CWS** from the Launch Pad • The SP Configuration Database menu appears. **SELECT** Enter Database Information • The Enter Database Information menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit enter_data** • The Enter Database Information menu appears. **SELECT** SP Frame Information • The SP Frame Information window appears. **TYPE** The start frame number and the number of frames in the Frame Count field. (Start frame defaults to **1**.) The starting frame tty port defaults to dev/tty0. You may need to change this depending upon your configuration. **SELECT** **yes** or **no** next to `Re-initialize the System Data Repository`, as follows: **no** if you have more (noncontiguous) frame entries to make in this panel. **yes** if you are entering only one series of contiguous frames, or entering the last series of noncontiguous frames. **SELECT** **yes** to `Multiple NSB Frame (SP Switch2 only)`. **SELECT** **yes** or **no** next to `Allow Frame Numbers greater than 128`, as follows: **yes** if you want to use a frame number greater than 128. **no** if you want to use a frame number that is less than or equal to 128. **PRESS** **Ok** to enter frame data to the SDR.
spframe	Specify **spframe** command with **-r yes** to reinitialize the SDR (when running the command for final series of frames), a starting frame number, a frame count, and the starting frame's tty port. The following example enters information for two frames (frame 1 to frame 2) and indicates that frame 1 is connected to **/dev/tty0**, frame 2 to **/dev/tty1**, and reinitializes the SDR. `spframe -r yes -m 1 2 /dev/tty0`

Step 33: Enter non-SP frame information and reinitialize the SDR (optional)

If you entered SP or multiple NSB frame information in "Step 32: Enter SP or multiple NSB frame information and reinitialize the SDR" on page 54, you must reinitialize the SDR before continuing to enter frame information for non-SP frames. You must perform this step at least once for each frame protocol of non-SP frames that you are adding to the system.

If you want to add an SP-attached server or clustered enterprise server (for example, the RS/6000 Enterprise Server Model M80 or IBM @server pSeries 690)

without reinstalling its software, install the rest of the new SP system. Once you have completed the steps, follow the steps in “Chapter 6. Reconfiguring the RS/6000 SP system” on page 195 to integrate the new SP-attached server.

SP-attached servers and clustered enterprise servers also require frame objects in the SDR. These frames are referred to as non-SP frames and one object is required for each server attached to your SP. These objects have a non-SP hardware protocol associated with them which instructs PSSP as to which method of hardware communications is to be used for controlling and monitoring the node associated with this frame object. Valid hardware protocol values of the nodes within the frame are:

HMC IBM @server pSeries 690 servers

CSP RS/6000 H80, M80, and IBM @server pSeries 660 servers (6H0, 6H1, 6M1)

SAMI RS/6000 S70, S7A, and S80 or IBM @server pSeries 680 servers

The number of tty port values you must define depends on the hardware protocol type you selected.

HMC Does not require a tty port value, but the HMC must be connected by the SP Ethernet administrative LAN

CSP Requires one tty port value

SAMI Requires two tty port values

The servers that use the SAMI hardware protocol require two tty port values to define the tty ports on the control workstation to which the serial cables connected to the server are attached. The tty port value defines the serial connection to the operator panel on these servers for hardware controls. The s1 tty port value defines the connection to the serial port on the servers for serial terminal (s1term) support.

Switch port numbers are required on SP Switch or switchless systems for each SP-attached server in your system. This information is available from your Switch Configuration Worksheet. Although switch ports are not required for switchless or SP Switch2 clustered enterprise servers, you may want to specify a switch port if you plan to add an SP frame sometime in the future. *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* explains how to fill out your worksheet and provides details on assigning switch port numbers.

For pSeries 690 servers in an SP Switch or switchless system, a switch node number is required for each logical partition (LPAR). These switch node numbers must be specified to PSSP through the **/etc/switch.info** file. Manually edit the **/etc/switch.info** file to include one entry for each LPAR in the attached server. See the **switch.info** file in *PSSP: Command and Technical Reference* for details on editing this file.

An example of the **/etc/switch.info** file might contain the following entries for a pSeries 690 server that will be defined as frame 5 with four LPARs attached to switch 2 in the system:

```
# Node_number        Switch_node_number
65                          16
66                          17
67                          18
68                          19
```

If you are running your pSeries 690 server in full system partition (SMP) mode or will only be defining one LPAR and it is assigned partition ID 1, you can skip this operation. Instead just simply enter the switch node number when you enter the other non-SP frame information later in this step through the SMIT menu or the **spframe** command.

You can enter information for non-SP frames using Perspectives, SMIT, or the **spframe** command. If using SMIT, a different procedure is used for each hardware protocol. If frames, tty ports, or switch port values are not all contiguously numbered, repeat this step for each series of contiguous information. To save time, do not specify the reinitialization of the SDR until you are entering the final series of contiguous frames.

If using:	Do this:	
Perspectives	**SELECT**	**smit config_data on CWS** from the Launch Pad • The SP Configuration Database menu appears.
	SELECT	Enter Database Information • The Enter Database Information menu appears. From this point, you can follow the rest of the SMIT steps described in one of the next three rows of this table. Choose the row that matches the hardware protocol of the server for which you are entering non-SP frame data.
SMIT with **HMC** protocol	**TYPE**	**smit enter_data** • The Enter Database Information menu appears.
	SELECT	Non-SP Frame Information • The Non-SP Frame Information window appears.
	SELECT	HMC - pSeries 690 • The HMC - pSeries 690 Information window appears.
	TYPE	• The frame number in the Frame Number field. • The starting switch port number in the Starting Switch Port Number (SP Switch or switchless systems only) field (optional for the clustered enterprise server). • The IBM @server pSeries 690 server identifier as it is known to the Hardware Management Console in the domain name field. • The IP addresses or host names of the Hardware Management Console that controls the IBM @server pSeries 690 server in the HMC IP address field.
	SELECT	**yes** or **no** next to `Re-initialize the System Data Repository`, as follows: **no** if you need to make more non-SP frame entries. **yes** if you are entering the last non-SP frame.
	PRESS	**Ok** to enter frame data to the SDR.

If using:	Do this:	
SMIT with **CSP** protocol	**TYPE**	**smit enter_data** • The Enter Database Information menu appears.
	SELECT	Non-SP Frame Information • The Non-SP Frame Information window appears.
	SELECT	CSP - RS/6000 H80, M80, and pSeries 660 (models 6H0, 6H1, 6M1) • The CSP - RS/6000 H80, M80, and pSeries 660 (models 6H0, 6H1, 6M1) window appears.
	TYPE	• The starting frame number in the Start Frame field. • The number of frames in the Frame Count field. • The starting tty port in the Starting Frame tty port field. • The starting switch port number in the Starting Switch Port Number (SP Switch or switchless systems only) field (optional for the clustered enterprise server).
	SELECT	**yes** or **no** next to `Re-initialize the System Data Repository`, as follows: **no** if you need to make more non-SP frame entries. **yes** if you are entering the last non-SP frame.
	PRESS	**Ok** to enter frame data to the SDR.
SMIT with **SAMI** protocol	**TYPE**	**smit enter_data** • The Enter Database Information menu appears.
	SELECT	Non-SP Frame Information • The Non-SP Frame Information window appears.
	SELECT	SAMI - RS/6000 S70, S7A, and S80 and pSeries 680 • The SAMI - RS/6000 S70, S7A, and S80 and pSeries 680 Information window appears.
	TYPE	• The starting frame number in the Start Frame field. • The number of frames in the Frame Count field. • The starting tty port (**the tty for the operator panel**) in the Starting Frame tty port field. • The starting switch port number in the Starting Switch Port Number (for SP Switch or switchless systems only) field (optional for the clustered enterprise server). • The s1 tty port (**the tty for the serial terminal**) in the s1 tty port field only if the s1 tty port value is not incrementally one more than the tty port field.
	SELECT	**yes** or **no** next to `Re-initialize the System Data Repository`, as follows: **no** if you need to make more non-SP frame entries. **yes** if you are entering the last non-SP frame.
	PRESS	**Ok** to enter frame data to the SDR.

If using:	Do this:

spframe	Specify the **spframe** command with the **-n** option for each series of contiguous non-SP frames. The **-n** option is not required for switchless clustered enterprise servers or SP Switch2 systems. Specify the **-r yes** option when running the command for the final series of frames. Include the starting frame number, the number of frames, the starting tty port value, and the starting switch port number for each invocation of the command.
	
	The following example enters non-SP information for one S80 server (frame 5), one H80 server (frame 6), and one pSeries 690 server with four LPARs (frame 7).
	
	The first server has the following characteristics:
	
	`Frame Number: 5`
	`tty port for operator panel connection: /dev/tty4`
	`tty port for serial terminal connection: /dev/tty5`
	`switch port number: 10`
	
	The second server has the following characteristics:
	
	`Frame Number: 6`
	`tty port for operator panel connection: /dev/tty6`
	`switch port number: 11`
	
	The third server has the following characteristics:
	
	`Frame Number: 7`
	`switch port number: 12, 13, 14, 15`
	
	To define the first two servers to PSSP, enter:
	
	`spframe -r no -p SAMI -n 10 -s /dev/tty5 5 1 /dev/tty4`
	
	`spframe -r no -p CSP -n 11 6 1 /dev/tty6`
	
	Append the following to the **/etc/switch.info** file:
	
	`7,1 12`
	`7,2 13`
	`7,3 14`
	`7,4 15`
	
	To define the third server to PSSP and reinitialize the SDR, enter:
	
	`spframe -r yes -p HSC -d huntley -i 129.33.32.121 7`

Note: The SP-attached server and clustered enterprise server in your system will be represented with the node number corresponding to the frame defined in this step. For pSeries 690, each logical partition in the server will be represented as a node on that frame. PSSP will assign an SP slot number and node number that corresponds to the partition ID set by the HMC for that partition. Continue with the remaining installation steps to install the SP-attached server, clustered enterprise server, or logical partition as an SP node.

Step 34: Update the state of the supervisor microcode

This step ensures that you have the latest level of microcode required by the SP frames, nodes, and switches on your SP system.

Note: You must have the latest version of **ssp.ucode** installed that is appropriate for your PSSP level before proceeding.

If using:	Do this:
Perspectives	**SELECT** **smit supervisor on CWS** from the Launch Pad. • The RS/6000 SP Supervisor Manager menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit supervisor** • The RS/6000 SP Supervisor Manager menu appears. The first five selections on the menu allow you to query the state of the microcode on the supervisor cards. Once you have determined whether a supervisor requires an action, based on the displayed state, continue with these steps. Move the cursor to the RS/6000 SP Supervisor Manager heading and select the Help Key (F1). A list of hardware that supports supervisor microcode is displayed. To update all of the supervisors, SELECT Update ″All″ Supervisors That Require Action. To update a subset of supervisors, SELECT Update Selectable Supervisors That Require Action.

spsvrmgr	The following command gives the status in report form of all of your frames, nodes, and switches:
	`spsvrmgr -G -r status all`
	
	The following command updates the microcode on the frame supervisor of frame 3:
	`spsvrmgr -G -u 3:0`

Step 35: Verify System Monitor installation

Perform this step to verify that the System Monitor and Perspectives have been correctly installed.

If using:	Do this:
Perspectives	**SELECT** **smit SP_verify on CWS** from the Launch Pad • The RS/6000 SP Installation/Configuration Verification menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit SP_verify** • The RS/6000 SP Installation/Configuration Verification menu appears. **SELECT** System Monitor Configuration
spmon	Enter: `spmon_ctest`

After the tests are run, the system creates a log in **/var/adm/SPlogs/spmon** called **spmon_ctest.log**.

See the section on “Diagnosing System Monitor problems” in *PSSP: Diagnosis Guide* if the verification test fails.

Step 36: Verify frame information

All frames must be powered up and connected to the control workstation so that the nodes are automatically detected and added to the SDR.

If using:	Do this:
Perspectives	**SELECT** The **Hardware Perspective** icon by double clicking • The Hardware Perspective appears with the Nodes Pane showing by default. If you had the Hardware Perspective up before you added the frame information, you should delete and re-add the Nodes pane. Next open the Frames pane to verify that all of your hardware is displayed. The number of frames and assignment of nodes within the frames should match your configuration.
spmon	Type: `spmon -d -G`

You should see the SP frames represented with thin, wide, or high nodes, depending on your configuration. If using Perspectives, SP-attached servers are shown as a unique SP-attached server icon. If using **spmon -d**, SP-attached servers are represented as a one node frame. The pSeries 690 servers will be represented as one frame with one node per LPAR within that frame. For multiple node switch board frames and for intermediate switch board frames, you should see the SP frames represented with switches listed in the appropriate slot locations. If your frames are not correctly represented, you may have a hardware problem, such as a misplugged RS-232 cable. See the “Diagnosing hardware and software problems” chapter in *PSSP: Diagnosis Guide* for help in correcting the error. If an error occurred, the frame must be deleted, using the **spdelfram** command, prior to reissuing the **spframe** command. After updating the RS-232 connection to the frame, you should reissue the **spframe** command.

Step 37: Enter the required node information

- Verify that the node objects have been created by issuing **splstdata -n** and verify that there is an entry for each node in your system.
- Be sure to have your node configuration worksheet on hand with all the node information completed before attempting to perform this step. *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* explains how to fill out your worksheet.
- If multiple IP interfaces map to the same host name on the starting node, you must enter the Ethernet IP address for the starting node. Do not enter its host name.

  If multiple IP interfaces do not map to the same host name on the starting node and you decide to enter its host name, it must be identical to the default host name returned by the **host** command for the starting node SP Ethernet IP address. For example, if the SP Ethernet administrative LAN adapter IP address of a node is 123.45.678.90 and host 123.45.678.90 gives v64n90.xen.kry.arg.com, then this host name must be used.
- The host name of a node is case sensitive. If you choose to enter the host name for a node, it must match the format of the host name returned when you issue **/usr/bin/host** against the node’s IP address.
- Enter a correct value for Ethernet speed (10, 100, or auto), Duplex (full, half, or auto), and Type (bnc, dix, tp, or NA) for the SP Ethernet adapter on each node.

- When adding nodes with connected SP expansion I/O units, verify that:
  - The node expansion objects were created by issuing **splstdata -x**
  - An entry exists for each I/O unit in your system
  - The node connection information is correct (see “Step 72: Verify node expansion configuration information (optional)” on page 102
- If a node is not directly connected to an SP Ethernet adapter on the control workstation or the host name of the control workstation is not set to the name of that SP Ethernet adapter, the default route for the node must be an adapter that is automatically configured. See the **spadaptrs** command in *PSSP: Command and Technical Reference* for a list of adapter types that can be automatically configured.
- For pSeries 690 servers, specifying the SP Ethernet adapter by its physical location code is suggested, especially if there is more than one Ethernet adapter present in the node. The physical location code for an adapter can be determined in one of the following ways:
  - Run the **spadapter_loc** command for the node. The command will return a list of all the SP supported adapters installed on the node. Save the results of the command in a file. You can use the returned hardware Ethernet address for your SP Ethernet adapter in “Step 38: Acquire the hardware Ethernet addresses” on page 66 to save processing time later.
  - Visually locate the adapter on your server and look up the physical location code in the hardware publications distributed with the server.
  - If AIX is installed and running on the server, run the AIX **lscfg -p** command to list the physical location codes of all the hardware devices installed on your system.

This step adds IP address-related information to the node objects in the SDR. It also creates adapter objects in the SDR for the SP Ethernet administrative LAN adapters on your nodes. This information is used during node customization and configuration.

Note: The default route that you enter in this step is not the same as the default route on the node. The route that you enter here goes in the SDR Node Class. It is the route over which the node communicates with its boot/install server (for example, install, customize, and so on). The default route must be a valid path from the SP Ethernet administrative LAN adapter to the node’s boot/install server and the control workstation.

The default route on the node is the route it will use for its network communications if there is no specific route to the destination. During the boot process, this is set to the default route in the SDR. It can be changed later on in the boot process or after the node is running, but should not be changed permanently in the SDR. For FDDI, token ring, or other Ethernet adapters, create the route in **firstboot.cust**. The following example defines a route for an Ethernet adapter. This example also saves the route into the node’s ODM.

```
old_route_info=$($lsattr -E -l inet0 | $grep 'route *net,.*,
0, 0-9 . *' | $awk ' print $2; ' | $tail -n 1) #-
if -n "$old_route_info" ; then #-
$chdev -l inet0 -a delroute="$old_route_info" > /dev/null
2>&1 #-
fi #-
$chdev -l inet0 -a route="0,<route>"
```

In order for the route to remain set after customization, also set the route up in **/etc/inittab** after the line that runs **rc.sp**. For the switch, set the route up in **/etc/inittab** after the line that runs **rc.switch**.

Enter information about your nodes attached to each Ethernet adapter using Perspectives, SMIT, or the **spadaptrs** command.

If using:	**Do this:**
Perspectives	**SELECT** **smit config_data on CWS** from the Launch Pad. • The SP Configuration Database Management menu appears. **SELECT** Enter Database Information • The Enter Database Information menu appears. **SELECT** Node Database Information • The Node Database Information menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.

If using:	Do this:

SMIT	**TYPE** **smit node_data**
	• The Node Database Information menu appears.
	SELECT SP Ethernet Information
	• The SP Ethernet Information window appears:
	
	If you have wide nodes or high nodes which each occupy multiple slots in your system, entering **yes** next to `Skip IP Address for Unused Slots?` can be useful in assigning IP addresses that correspond to the slots in the frame, with each wide node address incrementing by 2, each thin node address incrementing by 1, and each high node node incrementing by 4.
	
	You can avoid skipping IP addresses for the empty slots by entering **no** next to `Skip IP Address for Unused Slots?` In this way IP addresses are assigned consecutively for both thin and wide nodes.
	
	The distribution of your IP addresses determines how many times you perform this step. You may have to do it more than once if:
	• There are gaps in your IP addresses that are not caused by wide or high nodes
	• You want to set up alternate default routes or netmasks for certain IP address ranges
	
	Enter the following information for each consecutive block of nodes:
	1. Start Frame, Start Slot, and Node Count **OR** Node Group (see *PSSP: Administration Guide* for more information on node groups) **OR** Node List
	2. Either adapter name, or for the pSeries 690, a physical location code with the adapter type. For models other than the pSeries 690 server, you must specify the en0 adapter name.
	3. Starting Node's IP Address or Hostname
	4. Netmask
	5. Default Route Hostname or IP Address
	6. Ethernet Adapter Type (bnc, dix, tp, or NA)
	7. Duplex (full, half, or auto)
	8. Ethernet Speed (10, 100, or auto)
	9. Skip IP Addresses for Unused Slots?
	
	If you specify nodes with a node list, you cannot specify **yes** for `Skip IP Addresses for Unused Slots?`
	
	TYPE Data in the fields as required. Refer to your worksheet.
	
	PRESS **Ok** to store the data.
	
	Starting Slot is always relative to the frame and not to the system. This means that the first slot in the second, third, and fourth frames is still slot 1 rather than slots 17, 33, and 49. For example, for the first frame you might enter:
	
	`Start Frame 1`
	`Start Slot  1`
	
	and for a second frame, you might enter:
	
	`Start Frame 2`
	`Start Slot  1`
	
	Node List is used to specify a group of node numbers separated by commas. Node numbers can be referenced for systems with more than one frame. (Node number 17 would be used for frame 2 slot 1.) For example:
	
	`1,5,7,9,15,17,19`
	
	You can also specify a file that contains a single line of data containing the node list, separated by commas. Enter the full-path name, unless the file is in your current directory. For example, if you have a list of nodes in **/tmp/***node_list*, enter the following in the *node_list* field:
	
	`/tmp/`*`node_list`*

If using:	Do this:

spadaptrs	This following example configures an SP Ethernet administrative LAN adapter network of 16 nodes with IP addresses ranging from 129.33.32.1 to 129.33.32.16, a netmask of 255.255.255.192, and a default route of 129.33.32.200 for a twisted-pair Ethernet using auto-negotiate for the communication transfer and rate:
	`spadaptrs -e 129.33.32.200 -t tp -d auto -f auto 1 1 16 en0 \`
	`129.33.32.1 255.255.255.192`
	
	The following example configures the adapter on the SP Ethernet administrative LAN adapter for the first logical partition of a pSeries 690 server. The adapter is a twisted-pair Ethernet adapter with communication transfer and rate set to auto-negotiate. The IP address is 129.33.32.65 with a netmask of 255.255.255.192. The pSeries 690 server is represented as frame 5, the node is assigned slot 1, and the adapter is located at physical location U1.9-P2-I2/E1.
	`spadaptrs -P U1.9-P2-I2/E1 -t tp -d auto -f auto 5 1 1 en 129.33.32.65 \`
	`255.255.255.192`

If you are adding an extension node to your system, you may want to enter required node information now. For more information, refer to “Chapter 10. Installing extension nodes” on page 281.

Step 38: Acquire the hardware Ethernet addresses

- Do not do this step on a production running system because it shuts down the nodes.
- If you are adding a node, select only the new node. All the nodes you select are powered off and back on.
- The nodes for which you are obtaining Ethernet addresses must be physically powered on when you perform this step. No ttys can be open in write mode.

This step gets hardware Ethernet addresses for SP Ethernet administrative LAN adapters for your nodes, either from a file or from the nodes themselves, and puts them into the Node Objects in the SDR. That information is used to set up the **/etc/bootptab** files for your boot/install servers. This step will also **ping** the default route set for this node.

If you know the hardware Ethernet addresses, you can speed this process by putting the addresses in the **/etc/bootptab.info** file. If you are performing this step for a pSeries 690 server, you may already have the hardware Ethernet addresses available to you from “Step 37: Enter the required node information” on page 62. Create the **/etc/bootptab.info** file as follows:

- Create a file named **/etc/bootptab.info** (if it does not already exist), listing your RS/6000 SP nodes by node number, followed by a blank and the hardware Ethernet address. For example, a file containing addresses for a frame might look like this:

```
1 08005ABAB177
3 08005ABAAEAB
5 08005ABAB161
7 08005ABAB17A
9 02608CF53067
13 02608CF527F2
17 08005ABAB1A0
19 08005ABAB062
21 002035D34F7A
22 002035D34FE2
23 002035D34F3C
24 002035D34F70
```

```
25 002035D34E65
26 002035D34E5F
27 002035D34FE5
28 002035D34F68
29 02608CF55E6D
```

The **/etc/bootptab.info** file is not required. If you do not know your hardware Ethernet addresses, and the **/etc/bootptab.info** file does not exist, use **sphrdwrad** to access the SP node and retrieve the hardware Ethernet address for you. (This makes **sphrdwrad** take longer to run.)

If using:	Do this:
Perspectives	**SELECT** **smit config_data on CWS** from the Launch Pad. • The SP Configuration Database Management menu appears. **SELECT** Enter Database Information • The Enter Database Information menu appears. **SELECT** Node Database Information • The Node Database Information menu appears. **SELECT** Get Hardware Ethernet Addresses From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit node_data** • The Node Database Information menu appears. **SELECT** Get Hardware Ethernet Addresses • The Get Hardware Ethernet Addresses window appears. **TYPE** The starting frame, slot, and node count (the number of nodes for each consecutive series of nodes), or node group. **or** The node list. Type the node numbers (separated by commas) or the file containing the list (for example, **/tmp/***node_list*) **PRESS** **Ok** to get the data.

sphrdwrad	This example gets all hardware Ethernet addresses for an RS/6000 SP system.
	`sphrdwrad 1 1 rest`
	This example gets all hardware Ethernet addresses for the nodes specified in the node list (the **-l** flag):
	`sphrdwrad -l 10,12,17`

Step 39: Verify that the Ethernet addresses were acquired

This step verifies that Ethernet addresses were placed in the SDR node object.

If using:	Do this:
Perspectives	**SELECT** **smit config_data on CWS** from the Launch Pad • The SP Configuration Database Management menu appears. **SELECT** List Database Information • The List Database Information menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit list_data** • The List Database Information menu appears. **SELECT** List Node Database Information • The List Node Database Information menu appears. **SELECT** List Node Boot/Install Information • A window appears listing the node Ethernet information. **TYPE** The Start Frame, Start Slot, and Node Count **OR** Node Group **OR** Node List **PRESS** **Ok**
splstdata	**Attention:** If your system is large, **splstdata** returns great quantities of data. You may want to pipe the command output through a filter to reduce the amount of data you see. To display SDR boot/install data, enter: `splstdata -b`

Step 40: Configure additional adapters for nodes

Perform this step if you have a switch or if you require any additional adapters.

If you are configuring more than eight of one particular adapter type, you must change the ifsize parameter in the **tuning.cust** file.

Be sure to have your switch configuration worksheet on hand with all the switch information completed before attempting to perform this step. *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* explains how to fill out your worksheet.

This step creates adapter objects in the SDR for each node. The data in the adapter objects is used during the customization or installation steps to configure the adapters on the nodes. You can configure the following adapter types with this procedure:

- Ethernet (en)
- FDDI (fi)
- Token ring (tr)
- css

To configure adapters such as ESCON and PCA, you must configure the adapter manually on each node, using **dsh**, or modify the **firstboot.cust** file.

Note: Ensure that all additional adapters listed previously are configured before performing the following operations:

- Node installation
- mksysb install
- Node migration
- Node customization

During the preceding operations, **psspfb_script** is run which unconfigures and reconfigures all adapters found in the SDR. If additional adapters are **not** registered in the Adapter class of the SDR, they will not be configured after **psspfb_script** completes.

This requirement also includes any ATM LAN Emulator adapters that are defined as enX. Those adapters must also be defined in the SDR, otherwise **psspfb_script** will unconfigure them during the preceding operations.

Configuring the switch adapters

To configure your switch adapters for use with the RS/6000 SP system, use SMIT or issue the **spadaptrs** command. *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* contains additional information on IP addressing for the switch.

If using:	**Do this:**
Perspectives	**SELECT** **smit config_data on CWS** from the Launch Pad. • The SP Configuration Database Management menu appears. **SELECT** Enter Database Information • The Enter Database Information menu appears. **SELECT** Node Database Information • The Node Database Information menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.

If using:	Do this:

SMIT	**TYPE** **smit node_data**
	• The Node Database Information menu appears.
	SELECT Additional Adapter Information
	• The Additional Adapter Database Information window appears.
	TYPE The data in the fields. Refer to your worksheet as needed.
	PRESS **Ok** to store the data.
	PRESS **Cancel** to exit SMIT.
	
	The default css adapter attributes are:
	• `Skip IP Addresses for Unused Slots? no`
	• `Enable ARP for the css Adapter? yes`
	• `Use Switch Node Numbers for css IP Addresses? yes`
	Note: The default is **yes** for SP Switch systems and **no** for SP Switch2 systems.
	
	If you want to select **yes** next to `Skip IP Addresses for Unused Slots?`, you must set the `Use Switch Node Numbers` to **no**.
	
	If you set the `Use Switch Node Numbers` to **no**, you must set `Enable ARP` to **yes**.
	Note: You cannot set `Enable ARP` to **no** with SP Switch2 systems.
	
	For css adapters, if you select **yes** next to `Use Switch Node Numbers for css IP Addresses?`, you must use the `Start Frame`, `Start Slot`, and `Node Count` fields. `Start Slot` must be set to **1**.
	
	For css adapters, you must specify the adapter name, **css0** or **css1**, in the `Adapter Name` field.
	
	Enter the following information for each consecutive block of nodes.
	1. Start Frame, Start Slot, and Node Count **OR** Node Group **OR** Node List
	2. Adapter Name (**css0** or **css1**)
	3. Starting Node's IP Address or Hostname
	4. Netmask
	5. Token Ring Data Rate (required only when configuring a token-ring adapter.)
	6. Skip IP Addresses for Unused Slots?
	7. Enable ARP for the css Adapter?
	8. Use Switch Node Numbers for css IP Addresses?
	
	Starting Slot is always relative to the frame and not to the system. This means that the first slot in the second, third, and fourth frames is still slot 1 rather than slots 17, 33, and 49. For example, for the first frame you might enter:
	`Start Frame 1`
	`Start Slot  1`
	
	and for a second frame, you might enter:
	`Start Frame 2`
	`Start Slot  1`
	
	Node List is used to specify a group of node numbers separated by commas. Node numbers can be referenced for systems with more than one frame. (Node number 17 would be used for frame 2 slot 1.) For example:
	`1,5,7,9,15,17,19`
	
	You can also specify a file that contains a single line of data containing the node list, separated by commas. Enter the full-path name, unless the file is in your current directory. For example, if you have a list of nodes in **/tmp/node_list**, enter the following in the **Node List** field:
	`/tmp/`*`node_list`*

If using:	Do this:
spadaptrs	This example adds SDR information for a css (SP Switch and SP Switch2) network of 30 nodes (frame 1 slot 1 to frame 2 slot 16, with a wide node as the first node in each frame and the rest thin nodes, and a switch on each frame) with IP addresses from 129.33.34.1 to 129.33.34.30, and a netmask of 255.255.255.0. The IP addressing corresponds to the slots in the frame, with each wide node incrementing by 2 and each thin node incrementing by 1, and each high node by 4. If you specify the **-s** flag to skip IP addresses when you are setting the css switch addresses, you must also specify **-n no** to not use switch numbers for IP address assignment, and **-a yes** to use ARP. `spadaptrs -s yes -n no -a yes 1 1 30 css0 129.33.34.1 255.255.255.0`

Note: On systems containing an SP Switch2, PSSP will only install the **ssp.css** file set on nodes that have CSS switch adapters that are defined in the SDR. Nodes with no CSS switch adapters defined in the SDR will not have the **ssp.css** file set installed.

Configuring other additional adapters

To configure other additional adapters, for example Ethernet (en), token ring (tr), or FDDI (fi), you must select the `Additional Adapter Database Information`. For these adapters you can select either the Start Frame, Start Slot, and Node Count fields, or the Node List field.

Notes:

1. When using the token ring (tr) adapter, you must select the token ring rate (4 MB or 16 MB).
2. To ensure proper operation, exit SMIT and return to the Additional Adapter panel for each different type of adapter. This clears any extraneous values left behind in the panel.
3. Enter a correct value for Ethernet speed (10, 100, 1000, or auto), Duplex (full, half, or auto), and Type (bnc, dix, tp, fiber, or NA) for every Ethernet adapter on each node.
4. For pSeries 690 servers, specifying the adapter by its physical location code and adapter type is suggested, especially if there is more than one adapter of that type present in the node. The physical location code for an adapter can be determined in one of the following ways:
   - Run the **spadapter_loc** command for the node. The command will return a list of all of the SP supported adapters installed on the node. If you ran this command as part of “Step 37: Enter the required node information” on page 62 and saved the results of the command in a file, you may already have this information available to you.
   - Visually locate the adapter on your server and look up the physical location code in the hardware publications distributed with the server.
   - If AIX is installed and running on the server, run the AIX **lscfg -p** command to list the physical location codes of all of the hardware devices installed on your system.

The distribution of your IP addresses determines how many times you perform this step. You may have to do it more than once if:

- There are gaps in your IP addresses not caused by wide nodes or high nodes or SP-attached servers.
- You want to set up alternate default routes or netmasks for certain IP address ranges

If using:	Do this:
Perspectives	**SELECT** **smit config_data on CWS** from the Launch Pad. • The SP Configuration Database Management menu appears. **SELECT** Enter Database Information • The Enter Database Information menu appears. **SELECT** Node Database Information • The Node Database Information menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit node_data** • The Node Database Information menu appears. **SELECT** Additional Adapter Information • The Additional Adapter Database Information window appears. **TYPE** The data in the fields. Refer to your worksheet as needed. **PRESS** **Ok** to store the data. **PRESS** **Cancel** to exit SMIT. `Start Slot` must be set to **1**. Enter the following information for each consecutive block of nodes. 1. Start Frame, Start Slot, and Node Count **OR** Node Group **OR** Node List 2. Either adapter name or a physical location code with the adapter type (for models other than the IBM @server pSeries 690 server, you must specify the adapter name) 3. Starting Node's IP Address or Hostname 4. Netmask 5. Default Route Hostname or IP Address 6. Additional IP Addresses 7. Ethernet Adapter Type (bnc, dix, tp, fiber, or NA) 8. Duplex (full, half, or auto) 9. Ethernet Speed (10, 100, 1000, or auto) 10. Token Ring Data Rate (required only when configuring a token-ring adapter) 11. Skip IP Addresses for Unused Slots? 12. Enable ARP for the css Adapter? 13. Use Switch Node Numbers for css IP Addresses?

spadaptrs	This example adds SDR information for an fi0 (FDDI adapter) network of 30 nodes (frame 1 slot 1 to frame 2 slot 16, with a wide node as the first node in each frame and the rest thin nodes) with IP addresses from 129.33.34.1 to 129.33.34.30, and a netmask of 255.255.255.0. The IP addressing corresponds to the slots in the frame, with each wide node incrementing by 2 and each thin node incrementing by 1.
	`spadaptrs -s yes 1 1 30 fi0 129.33.34.1 255.255.255.0`
	This example adds SDR information for a tr0 (token ring adapter) for node 1 with IP address 129.33.35.1 and a netmask of 255.255.255.0, and references the node list field.
	`spadaptrs -l 1 -r 16 tr0 129.33.35.1 255.255.255.0`

If using:	Do this:

spadaptrs (continued)	This example adds SDR information for an additional Ethernet adapter for the second logical partition in a pSeries 690 server. The adapter is a twisted pair Ethernet adapter with duplex, the speed set to auto-negotiate, and is not the SP Ethernet adapter for the node. The IP address is 129.33.35.66 with a netmask of 255.255.255.0. The pSeries 690 server is represented as frame 5, the node is assigned slot 2, and the adapter is located at the physical location U1.9-P2-I2/E4.
	`spadaptrs -P U1.9-P2-I2/E4 -t tp -d auto -f auto 5 2 1 en \`
	`129.33.35.66 255.255.255.0`

Step 41: Configure the aggregate IP interface for nodes (SP Switch2 only)

To use the **ml0** interface for running jobs over the switch, use Perspectives, SMIT, or issue the **spaggip** command.

Note: This step is optional for an SP Switch2 system.

If using:	Do this:
Perspectives	**SELECT** **smit config_data on CWS** from the Launch Pad. • The SP Configuration Database Management menu appears. **SELECT** Enter Database Information • The Enter Database Information menu appears. **SELECT** Node Database Information • The Node Database Information menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit node_data** • The Node Database Information menu appears. **SELECT** Aggregate IP Information • The Aggregate IP Information menu appears. **TYPE** The values in the entry fields. **PRESS** **Ok** after making your changes.
spaggip	For example, to add an aggregate IP address of 9.114.66.20 and a network mask of 255.255.255.0 for device ml0 on node 7, enter: `spaggip -i css0,css1 -l 7 9.114.66.20 255.255.255.0`

Step 42: Configure the number of switch planes (SP Switch2 only)

You only need to perform this step if you are configuring more than one plane on your system. The default is automatically set to one plane. To configure an SP Switch2 two plane system, use Perspectives, SMIT, or issue the **spswplane** command. To check the number of planes defined, issue:

```
splstdata -e
```

If using:	Do this:	
Perspectives	**SELECT**	The **smit cluster_mgmt** icon • The RS/6000 SP Cluster Management menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE**	**smit cluster_mgmt** • The RS/6000 SP Cluster Management menu appears.
	SELECT	Perform Switch Operations • The Perform Switch Operations menu appears.
	SELECT	Specify Number of Switch Planes (SP Switch2 Only) • The Specify Number of Switch Planes menu appears.
	ENTER	The number of switch planes.
spswplane	For example, to configure two switch planes on your SP system, enter: `spswplane -p 2`	

Step 43: Configure initial host names for nodes

Do this step if:

- You do not want the default host name to match the SP Ethernet administrative LAN adapter name. The SP Ethernet administrative LAN adapter name is the default.
- You are using short host names. The default is long host names.

This step changes the default host name information in the SDR Node Objects used during customization to set up the host name on each node, and allows you to indicate how you want to name your RS/6000 SP nodes. The default is the long form of the SP Ethernet administrative LAN adapter host name, which is how the **spadaptrs** command processes defaulted host names.

You can indicate an adapter name other than the SP Ethernet administrative LAN for the node host names to be used, as well as whether the long or short form should be used. When determining whether you want the nodes' host name to be either in long or short form, be consistent with the host name resolution on the control workstation. If the **host** command returns the short form of a host name, you should choose the short form for the node's initial host name.

Multibyte host names are not supported on the SP.

If using:	Do this:	
Perspectives	**SELECT**	**smit config_data on CWS** from the Launch Pad. • The SP Configuration Database Management menu appears.
	SELECT	Enter Database Information • The Enter Database Information menu appears
	SELECT	Node Database Information • The Node Database Information menu appears
	SELECT	Hostname Information
		From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE**	**smit node_data** • The Node Database Information menu appears.
	SELECT	Hostname Information • The Hostname Information window appears.
	TYPE	Start Frame, Start Slot, Node Count **OR** Node List **OR** Node Group
		Also type the adapter name or adapter physical location for initial hostname and whether you are using short or long hostnames.
	PRESS	**Ok** to store the data.
sphostnam	This command indicates that the host name of each node is the long (fully qualified) form of the host name of the **css0** adapter, for a system with two frames and 32 nodes. `sphostnam -a css0 1 1 32`	

RS/6000 SP security installation and configuration

The following list outlines the steps necessary to configure and customize the SP selected authentication and authorization methods:

- Select Security Capabilities Required on Nodes
- Create DCE hostnames
- Update SDR with DCE Master Security and CDS Server Hostnames
- Configure DCE Clients (Admin portion)
- Configure SP Trusted Services to use DCE authentication
- Create SP Trusted Services DCE Keyfiles
- Select Authorization Methods for AIX Remote Commands
- Enable Authentication Methods for AIX Remote Commands
- Enable Authentication Methods for SP Trusted Services

Step 44: Select security capabilities required on nodes

This step sets the security capabilities to be installed on the nodes. If **dce** is selected, the DCE file sets will be installed on the nodes, and the security, CDS, clients, and RPC will be configured and started. The DCE file sets must be located in **/spdata/sys1/install/***name***/lppsource** on the control workstation to be installed automatically.

If **k4** is selected, various Kerberos V4 configuration files will be installed.

Note: By default, AIX standard authentication is part of the AIX BOS and, therefore, no installation is required on the node.

If using:	Do this:
SMIT	**TYPE** **smit spauth_config** • The RS/6000 SP Security menu appears. **SELECT** Select Security Capabilities Required on Nodes • The Select Security Capabilities Required on Nodes menu appears. **SELECT** System Partition Name • Press **List** (F4) and then move the cursor to the desired system partition and press Enter. **SELECT** Authentication Methods • Press **List** (F4) and then select one or more authentication methods and press Enter.
spsetauth	For example, enter: `spsetauth -p partition1 -i dce`

Step 45: Create DCE hostnames (required for DCE)

If you selected DCE as an authentication method, you must set a DCE hostname for each node in the SDR. This step uses the nodes' reliable hostname as the DCE hostname if a DCE hostname does not already exist.

If using:	Do this:
SMIT	**TYPE** **smit spauth_config** • The RS/6000 SP Security menu appears. **SELECT** Create DCE hostnames
create_dcehostname	For example, enter: `create_dcehostname`

Step 46: Update the SDR with DCE master security and CDS server hostnames (required for DCE)

This step updates the SDR with DCE master security and CDS server hostnames.

If using:	Do this:
SMIT	**TYPE** **smit spauth_config** • The RS/6000 SP Security menu appears. **SELECT** Update SDR with DCE Master Security and CDS Server Hostnames • The Update SDR with DCE Master Security and CDS Server Hostnames menu appears. **SELECT** Master Security Server hostname • Enter the full name of the host containing your DCE Master Security Server. **SELECT** CDS Server hostname • Enter the full name of the host containing your DCE Initial Directory Server.
setupdce	For example, enter: `setupdce -u -s c186cw.pok.ibm.com -d c186cw.pok.ibm.com`

Step 47: Configure admin portion of DCE clients (required for DCE)

This step configures the admin portion of DCE clients.

Note: You will be prompted for the cell administrator password.

If using:	Do this:
SMIT	**TYPE** **smit spauth_config** • The RS/6000 SP Security menu appears. **SELECT** Configure DCE Clients (Admin portion) • The Configure DCE Clients (Admin portion) menu appears. **SELECT** Cell Administrator id • Enter a DCE principal having cell administration privileges. The default is **cell_admin**. **SELECT** Lan Profile id • Enter your DCE Lan profile path name. The DCE default is /.:/**lan_profile**.

setupdce	For example, enter:
	`setupdce -c cell_admin -l /.:/lan_profile`
	Note: To run this command off of the SP, you must set the SP_NAME environment variable on a remote workstation to point to the SDR of the SP system being configured. The value must be a resolvable address. For example:
	`export SP_NAME=spcws.abc.com`

Step 48: Configure SP Trusted Services to use DCE authentication (required for DCE)

This step configures SP Trusted Services into the DCE database. Data is entered into both the DCE registry and the Security Server database. You must have cell administrator authority to run this step.

This step creates SP Trusted Services principals and accounts. It uses the **/usr/lpp/ssp/config/spsec_defaults** and **/spdata/sys1/spsec/spsec_overrides** files described in “Step 22.1: Update the spsec_overrides file (optional)” on page 45.

If using:	Do this:
SMIT	**TYPE** **smit spauth_config** • The RS/6000 SP Security menu appears. **SELECT** Configure SP Trusted Services to use DCE Authentication
config_spsec	For example, enter: `config_spsec -v` **Note:** To run this command off of the SP, you must set the SP_NAME environment variable on a remote workstation to point to the SDR of the SP system being configured. Refer to the **config_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.

Step 49: Create SP Trusted Services DCE keyfiles (required for DCE)

Note: You must be root on the control workstation with default credentials to perform this step.

This step creates SP Trusted Services keyfiles. It uses the **/usr/lpp/ssp/config/spsec_defaults** and **/spdata/sys1/spsec/spsec_overrides** files described in “Step 22.1: Update the spsec_overrides file (optional)” on page 45.

If using:	Do this:
SMIT	**TYPE** **smit spauth_config** • The RS/6000 SP Security menu appears. **SELECT** Create SP Trusted Services Keyfiles
create_keyfiles	For example, enter: `create_keyfiles -v`

After running the **create_keyfiles** command, you should reacquire SP administrative credentials as described in “Step 24: Obtain credentials” on page 46.

Step 50: Select authorization methods for AIX remote commands

This step sets the authorization methods that will be used for AIX remote commands. It also calls **updauthfiles** to update security-related files such as **/.k5login**, **/.rhosts**, and **/.klogin** (as appropriate).

A new option of **none** has been added to this menu. If **none** is selected, no other authorization methods can be selected at the same time for the selected system partition. The **none** option can be selected only if all nodes are at PSSP 3.4 or later.

To enable **none** on any system partition, the secure remote command method must have been enabled on the Site Environment Menu along with the required restricted root access enablement.

When **none** is selected, no PSSP entries are automatically put in the **.k5login**, **.rhosts**, and **.klogin** files by **updauthfiles** to enable root remote command access for that system partition. If **none** is chosen for all system partitions, there will be no PSSP entries in these files. If some system partitions have authorization methods for AIX remote commands defined, the **.k5login**, **.rhosts**, and **.klogin** files will be created for each of the authorizations enabled.

Boot/install server nodes still require the **rsh** or **rcp** capability if they are not the control workstation to use NIM services. If **none** is selected as the AIX authorization method for remote commands, it will be up to the administrator to add the authorization methods necessary for boot/install activities.

GPFS, IBM Virtual Shared Disk, Problem Management, and some LoadLeveler functions (for example, `llctl -g start` or `llctl -h start`) will not function with **none** enabled.

Note: You must select at least one method in this step.

If using:	Do this:
SMIT	**TYPE** **smit spauth_config** • The RS/6000 SP Security menu appears. **SELECT** Select Authorization Methods for AIX Remote Commands • The Select Authorization Methods for AIX Remote Commands menu appears. **SELECT** System Partition name • Press **List** (F4) then move the cursor to the desired system partition and press Enter. **SELECT** Authorization Methods • Press **List** (F4) then select one or more authorization methods and press Enter.
spsetauth	For example, enter: `spsetauth -d -p partition1 dce`

Step 51: Enable authentication methods for AIX remote commands

This step enables the authentication methods that will be used for AIX remote commands.

Notes:

1. If the authentication methods enabled for use by SP Trusted Services includes DCE, the authentication methods enabled for use by the AIX remote commands must include Kerberos V5.
2. If the authentication methods enabled for use by SP Trusted Services includes compatibility, the authentication methods enabled for use by the AIX remote commands must include Kerberos V4.

If using:	Do this:
SMIT	**TYPE** **smit spauth_config** • The RS/6000 SP Security menu appears. **SELECT** Enable Authentication Methods for AIX Remote Commands • The Enable Authentication Methods for AIX Remote Commands menu appears. **SELECT** Enable on Control Workstation Only • Set to **yes** to enable the control workstation only. **Note:** You cannot specify **yes** to both **Enable on Control Workstation Only** and **Force change on nodes**. **SELECT** Force change on nodes • Set to **no** to not force a change on the nodes. **SELECT** System Partition name • Press **List** (F4) then move the cursor to the desired system partition and press Enter. **SELECT** Authentication Methods • Press **List** (F4) then select one or more authentication methods and press Enter.
chauthpar	For example, enter: `chauthpar -c -p partition1 k5 std`

Step 52: Enable authentication methods for SP Trusted Services

Note: You can skip this step if you only have AIX standard security enabled.

This step enables the authentication methods that will be used for SP Trusted Services.

If using:	Do this:	
SMIT	**TYPE**	**smit spauth_config** • The RS/6000 SP Security menu appears.
	SELECT	Enable Authentication Methods for SP Trusted Services • The Enable Authentication Methods for SP Trusted Services menu appears.
	SELECT	Enable on Control Workstation Only • Set to **yes** to enable the control workstation only. **Note:** You cannot specify **yes** to both **Enable on Control Workstation Only** and **Force change on nodes**.
	SELECT	Force change on nodes • Set to **no** to not force a change on the nodes.
	SELECT	System Partition name • Press **List** (F4) then move the cursor to the desired system partition and press Enter.
	SELECT	Authentication Methods • Press **List** (F4) then select one or more authentication methods and press Enter.
chauthpts	For example, enter: `chauthpts -c -p partition1 dce`	

Step 53: Start the key management daemon (required for DCE)

If you selected DCE as an authentication method and enabled DCE in the previous security steps, you must start the key management daemon on the control workstation. The key management daemon manages the DCE passwords associated with the SP Trusted Services. This daemon is started automatically on a node that is configured to use DCE authentication.

To start the key management daemon, issue:

```
/usr/lpp/ssp/bin/spnkeyman_start
```

Step 54: Add an extension node (optional)

At this point, you can optionally add an extension node to your system. Refer to “Chapter 10. Installing extension nodes” on page 281 for more information.

Step 55: Start RSCT subsystems

The PSSP installation code sets up a single default system partition that includes all nodes in the system. This system partition is created automatically and is called the default partition because it always exists, even on a system that cannot be partitioned.

At this time you need to add and start the RSCT subsystems. Topology Services (hats), host response (hr) are examples of RSCT subsystems. RSCT subsystems are managed by the **syspar_ctrl** command and are listed in the file **/usr/lpp/ssp/config/cmi/syspar_subsystems**. For a more complete description of RSCT subsystems, refer to the “Managing system partition-sensitive subsystems using syspar_ctrl” section in *PSSP: Administration Guide*.

If using:	Do this:
Perspectives	**SELECT** **syspar_ctrl -A** from the Launch Pad. • The **syspar_ctrl -A** command is run.
syspar_ctrl	`syspar_ctrl -A`

Step 56: Verify that RSCT subsystems have started

If using:	Do this:
Perspectives	**SELECT** **syspar_ctrl -E** from the Launch Pad. • The **syspar_ctrl -E** command is run.
syspar_ctrl	Enter: `syspar_ctrl -E`

Before continuing with the install, verify that the following subsystems have been started and have an “active” state.

- haem
- hags
- hats
- hr

To see if these subsystems have been successfully started, issue the following command:

```
lssrc -a | grep default_syspar_name
```

For example, if your default system partition name is k22s, issue:

```
lssrc -a | grep k22s
```

The preceding command returns the following output:

```
hags.k22s        hags            17134   active
hats.k22s        hats            22266   active
hr.k22s          hr              18228   active
haem.k22s        haem            21128   active
hagsglsm.k22s    hags            21338   active
haemaixos.k22s   haem            41000   active
Emonitor.k22s    emon                    inoperative
```

To continue with the install, the subsystems hags, hats, hr, and haem should all be active. If the subsystems are inactive, they should become active in a few minutes. Wait 3 minutes and check again.

If a single subsystem is inactive, simply try starting that particular subsystem by issuing:

```
syspar_ctrl -s subsystem_name
```

For example, if the subsystem is hags, issue:

```
syspar_ctrl -s hags
```

If more than one subsystem is inactive, stop and delete all of the RSCT subsystems by issuing:

```
syspar_ctrl -D
```

Then try to add and start all of the RSCT subsystems by issuing:

```
syspar_ctrl -A
```

If you still have inactive RSCT subsystems, refer to *PSSP: Diagnosis Guide* for further information.

Step 57: Set up nodes to be installed

- Do this step if you want to change the default installation settings for any of the nodes. To find out the default settings of your nodes, use the **splstdata** command.
- Be aware that the root password is not set if you are installing from a minimal mksysb. For more information on setting a root password when installing from a minimal mksysb, refer to “Step 61: Perform additional node customization” on page 90.
- If the boot/install server will be forwarding packets from the control workstation to a client node, the boot/install server is acting as a gateway to the control workstation. Therefore, ipforwarding must be correctly enabled. To turn ipforwarding on, issue:

  ```
  /usr/sbin/no -o ipforwarding=1
  ```

- You cannot export **/usr** or any directories below **/usr** because an NFS export problem will occur.

  If you have exported the **/spdata/sys1/install/image** directory or any parent directory, you must unexport it using the **exportfs -u** command before running **setup_server**. You need to do this because NIM attempts to export **/spdata/sys1/install/images/bos.obj.ssp.***, where **bos.obj.ssp.*** is the install image during **setup_server** processing. If you do not perform this task, you will receive an error. See the “Diagnosing NIM problems” chapter in *PSSP: Diagnosis Guide* for more information.
- Everything that is required in PSSP is installed on the nodes automatically, regardless of whether you use your own mksysb or the SP minimal mksysb.

This step does the following:

- Changes the default boot/install information for the node objects in the SDR so that you can indicate a different boot/install server configuration to the RS/6000 SP system
- Allows you to specify an alternate disk or disks to use when installing AIX on nodes

Using multiple boot/install servers

The default installation assumes one of the following:

- You have fewer than 40 nodes and the control workstation is configured to act as the boot/install server.
- You have more than 40 nodes, the control workstation and the first node in each frame is configured to act as the boot/install server.

Note: Do not install cascading levels of boot/install servers.

You should establish an administrative principal that has the following capabilities from the control workstation to the boot/install server nodes to run commands such as **spdelnode**:

- Root authority on the control workstation
- SP Trusted Services capability including **hardmon**
- Remote command capability

If you want different nodes to be installed by a different boot/install server, you must specify the target nodes and which node will serve as the boot/install server. For example, the first node of your second frame, node 17, will be a boot/install server for the remaining 15 nodes in the second frame. Use the **spchvgobj** command to enter this information into the SDR. The syntax used in the example specifies a start frame of 2, a starting slot of 2, and a count of 15 nodes. Perspectives or SMIT can also be used, as shown in the table at the end of this step.

```
spchvgobj -r selected_vg -n 17 2 2 15
```

After network installation is complete, if you have restricted root access (RRA) enabled, you must perform the following steps:

On the boot/install server node, you need to edit **/etc/sysctl.conf** to include the following entries:

- Include **/usr/lpp/ssp/sysctl/bin/install.cmds**
- Include **/usr/lpp/ssp/sysctl/bin/switch.cmds**
- Include **/usr/lpp/ssp/sysctl/bin/firstboot.cmds**, if initiating a node install customization

Using multiple boot/install servers in RRA is not recommended and is not automatically supported by PSSP. However, depending on the size of your system and network loads, it may not be possible to install your system with a single boot/install server.

Boot/install servers are NIM masters and, therefore, require **rsh** and **rcp** access to both the control workstation and to the nodes they serve. PSSP will not automatically create the correct entries in the authorization files to allow these commands to work.

To use additional boot/install servers, follow the following procedure to manually establish the correct authorizations on your system.

On the control workstation, the authorization files must have the following changes, depending on the setting of auth_root_rcmd:

standard An entry for the boot/install server node host name in **/.rhosts**

k4 An entry for the boot/install server node rcmd principal in **/.klogin**

dce An entry for the self-host and the spbgroot principal for the boot/install server node in **/.k5login**

Specifying your own image

The default installation assumes your nodes have not been preinstalled. If you want to have them installed with your own install image, you must specify the following:

- Which nodes you are installing with your own install image
- The name of the installation image you are using if you do not want to use the default image.

Selecting an installation disk

There are five ways you can specify the disk or disks to use for installation.

1. The hardware location format

   IBM strongly suggests that you use this format for SCSI devices. It ensures that you install on the intended disk by targeting a specific disk at a specific location. The relative location of hdisks can change depending on the hardware installed or possible hardware failures. You should always use this format when there are external disk drives present, because the manner in which the device names are defined may not be obvious. For example, to specify a single SCSI drive, enter:

   ```
   00-00-00-0,0
   ```

   or enter multiple hardware locations separated by colons:

   ```
   00-00-00-0,0:00-00-00-1,0
   ```

2. The device names format

   For example, to specify a single device name, enter:

   ```
   hdisk0
   ```

   or enter multiple device names separated by commas:

   ```
   hdisk0,hdisk1
   ```

3. A combination of the parent and connwhere attributes for SSA devices

   To specify the parent-connwhere attribute:

   ```
   ssar//0123456789ABCDE
   ```

   or to specify multiple disks, separate using colons as follows:

   ```
   ssar//0123456789ABCDE:ssar//0123456789ABCDE
   ```

   The parent-connwhere format should only be used for SSA drives.

   For more information on acquiring ssar numbers, see *AIX Kernel and Subsystems Technical Reference, Volume 2*.

4. The PVID format

   If a disk was previously configured as a physical volume in order for it to be assigned to a volume group, a physical volume identifier (PVID) was assigned to that disk by AIX. You can specify a disk by its PVID value as a string of 16 hexadecimal characters. For example:

   ```
   00d4c45202be737f
   ```

   To specify multiple disks by their PVID values, separate the specifications using colons:

   ```
   00d4c45202be737f:00d4c452eb639a2c
   ```

   Use the AIX **lspv** command to list the PVID values for the disks on your system. For more information on making an available disk a physical volume and setting its PVID, see *AIX System Management Guide: Operating System and Devices*.

5. The SAN target and logical unit identifier format

   Fibre channel attached disks are identified by a worldwide port name and a logical unit identifier (LUN ID). To specify the SAN_DISKID, combine the two values into a single string separated by *″//″*. For example, if the SAN target worldwide port name for a fibre channel attached disk is 0x50060482bfd12c9c and the LUN ID is 0x8000000000000000, the SAN_DISKID specification would be:

0x50060482bfd12c9c//0x8000000000000000

To specify multiple fibre channel disks, separate the specifications using colons:
0x50060482bfd12c9c//0x8000000000000000:0x50060482bbffd7cb//0x0

Use the AIX **lsattr -EH -l** *hdisk* command to determine the worldwide port name and LUN ID for a disk.

The hardware location, SSA parent-connwhere, PVID, and SAN_DISKID formats can be used together. Specify multiple mixed format disk values using colons to separate the specifications as follows:
00-00-09-0,1:ssar//0123456789ABCDE:00d4c45202be737f

The device names format cannot be combined with any of the other format types.

For more information on alternate root volume groups, see the “Managing root volume groups” appendix in *PSSP: Administration Guide*.

Note: The AIX **alt_disk_install** function is not related to the SP alternate root volume group support and is not supported with PSSP installation.

Mirroring the root volume group

One way to significantly increase the availability of the SP system is to set up redundant copies of the operating system on different physical disks using the AIX disk mirroring feature. Mirroring the root volume group means that there will be multiple copies of the operating system image available to a workstation or node. Mirrored system images are distributed so that a node can remain in operation even after one of the mirrored units fail.

When installing a node, you have a choice of how many copies of the root volume group you would like. AIX allows one (the original), two (the original plus one), or three (the original plus two) copies of a volume group. IBM strongly suggests that the root volume group be mirrored for a total of at least two copies. PSSP provides commands to facilitate root volume group mirroring.

You can specify how many copies and which disks to use with the **spchvgobj** command. Care should be taken when specifying disks so that no other single point of failure is introduced. For example, the specified disks should not be attached to the same adapter.

The default setting for the number of copies is based on the node type. The default is one copy for all nodes except the POWER3 Symmetric Multiprocessor (SMP) High Node, which has a default of two copies. These nodes are assumed to contain dual internal disk drives as a standard configuration. The disks will automatically be used for mirroring. If these nodes were not configured with the dual internal disks or you do not want mirroring, use the **spchvgobj** command to change the settings before installing the node.

For a complete description of how mirroring is handled by PSSP, see the “Managing root volume groups” appendix in *PSSP: Administration Guide*.

Changing volume group information in the SDR: Change the default volume group information in the SDR to specify a different boot/install server, your own installation image, a different target installation disk or disks, or mirror root volume groups.

If using:	Do this:
Perspectives	**SELECT** **smit config_data on CWS** from the Launch Pad. • The SP Configuration Database Management menu appears. **SELECT** Enter Database Information • The Enter Database Information menu appears **SELECT** Node Database Information • The Node Database Information menu appears From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit node_data** • The Node Database Information menu appears. **SELECT** Change Volume Group Information • The Change Volume Group Information window appears. **TYPE** The data in the fields. Refer to your worksheet as needed. In addition to the node range, you must supply the boot/install server node identifier. If you are not using the default image, you must supply the network install image name. If your AIX lppsource name is “default”, you must enter the correct name in the lppsource field (for example, aix433 or aix51). See “Step 12: Create the required /spdata directories” on page 22. To specify an alternate installation disk or disks, fill in the Physical Volume List field, using device names. **PRESS** **Ok** to store the data. The distribution of the nodes to be served and the number of different servers you want determines how many times you must perform this step. You may have to do it more than once if you want to define more than one server for different groups of nodes.

spchvgobj	You can use the **spchvgobj** command using the hardware location format for disk locations 00-07-00-0,0 and 00-07-00-1,0 for node 9 and set the number of copies to two. For example:
	`spchvgobj -r selected_vg -h 00-07-00-0,0:00-07-00-1,0 -l 9 -c 2`
	If you need to change the *lppsource_name* from default to a new *lppsource_name* such as aix433 for nodes 1 through 16, issue:
	`spchvgobj -r selected_vg -v aix433 1 1 16`
	If you need to change the *install_image_name* from default to a new *install_image_name* such as bos.obj.ssp.433 for nodes 17, 18, 21, 22, issue:
	`spchvgobj -r selected_vg -i bos.obj.ssp.433 -v aix433 -l 17,18,21,22`

Step 58: Verify all node information

This step verifies that all the node information has been correctly entered into the SDR.

If using:	Do this:
Perspectives	**SELECT** **smit config_data on CWS** from the Launch Pad. • The SP Configuration Database Management menu appears. **SELECT** List Database Information. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	Check each of the List Database panels for correct information. If you find any incorrect data, return to the following steps to make corrections: **Data** **Step** **Frame data** "Step 32: Enter SP or multiple NSB frame information and reinitialize the SDR" on page 54 **Node data** "Step 37: Enter the required node information" on page 62 **Additional adapters** "Step 40: Configure additional adapters for nodes" on page 68 **TYPE** **smit list_data** • The List Database Information window appears. **SELECT** Each panel listed. Verify all information is correct. **PRESS** **Done** to exit SMIT.
splstdata	**To display SDR:** **Enter:** **Site environment data** splstdata -e **Frame data** splstdata -f **Node data** splstdata -n **Adapter data** splstdata -a **Boot/install data** splstdata -b **SP expansion I/O data** splstdata -x **SP security settings** splstdata -p **Switch data** splstdata -s
If your system is large, **splstdata** returns great quantities of data. You may want to pipe the command output through a filter to reduce the amount of data you see.	

Step 59: Verify extension node information

At this point, you can optionally verify extension node information. Refer to "Chapter 10. Installing extension nodes" on page 281 for more information.

Task D. Customize the nodes

Step 60: Change the default network tunable values

When a node is installed, migrated, or customized (set to **customize** and rebooted), and that node's boot/install server does not have a **/tftpboot/tuning.cust** file, a default file of system performance tuning variable settings in **/usr/lpp/ssp/install/config/tuning.default** is copied to **/tftpboot/tuning.cust** on that node. You can override these values by following one of the methods described in the following list:

1. Select an IBM-Supplied Alternate Tuning File

IBM supplies three alternate tuning files which contain initial performance tuning parameters for three different SP environments:

a. **/usr/lpp/ssp/install/config/tuning.commercial** contains initial performance tuning parameters for a typical commercial environment.

b. **/usr/lpp/ssp/install/config/tuning.development** contains initial performance tuning parameters for a typical interactive/development environment.

c. **/usr/lpp/ssp/install/config/tuning.scientific** contains initial performance tuning parameters for a typical engineering/scientific environment.

Note: The SP-attached servers should not use the **tuning.scientific** file because of the large number of processors and the amount of traffic that they can generate.

To select one of these files for use throughout the nodes in your system, use SMIT or issue the **cptuning** command. When you select one of these files, it is copied to **/tftpboot/tuning.cust** on the control workstation and is propagated from there to each node in the system when it is installed, migrated, or customized. Each node inherits its tuning file from its boot/install server. Nodes that have as their boot/install server another node (other than the control workstation) obtain their **tuning.cust** file from that server node so it is necessary to propagate the file to the server node before attempting to propagate it to the client node. The settings in the **/tftpboot/tuning.cust** file are maintained across a boot of the node.

2. Create and Select Your Own Alternate Tuning File

   The following steps enable you to create your own customized set of network tunable values and have them propagated throughout the nodes in your system. These values are propagated to each node's **/tftpboot/tuning.cust** file from the node's boot/install server when the node is installed, migrated, or customized and are maintained across the boot of the node.

   a. On the control workstation, create the file **/tftpboot/tuning.cust**. You can choose to begin with a copy of the file located in **/usr/lpp/ssp/samples/tuning.cust** which contains a template of performance tuning settings which have been commented out. Or you may prefer to begin with a copy of one of the IBM-supplied alternate tuning files.

   b. Select the tunable values that are best for your system.

   c. Edit the **/tftpboot/tuning.cust** file by ensuring the appropriate lines are uncommented and that the tunable values have been properly set.

If using:	**Do this:**	
SMIT	**TYPE**	**smit select_tuning** • The Select System Tuning Parameters menu appears.
	SELECT	The desired tuning file

Once you have updated **tuning.cust**, continue installing the nodes. After the nodes are installed and customized, on all subsequent boots, the tunable values in **tuning.cust** will be automatically set on the nodes.

Note that each of the supplied network tuning parameter files, including the default tuning parameter file, contains the line **/usr/sbin/no -o ipforwarding=1**. IBM suggests that on non-gateway nodes, you change this line to read **/usr/sbin/no -o**

ipforwarding=0. After a non-gateway node has been installed, migrated, or customized, you can make this change in the **/tftpboot/tuning.cust** file on that node.

If you are configuring more than eight of one particular adapter type, you must change the ifsize parameter in the **tuning.cust** file.

For the latest performance and tuning information, refer to the RS/6000 Web site at:

http://www.rs6000.ibm.com/support/sp/perf

You can also access this information using the RS/6000 SP Resource Center.

Step 61: Perform additional node customization

Do this step to perform additional customization such as:

- Adding **installp** images
- Configuring host name resolution
- Setting up NFS, AFS, or NIS
- Configuring adapters that are not configured automatically
- Modifying TCP/IP configuration files
- Setting time zones

IBM provides the opportunity to run customer-supplied scripts during node installation:

script.cust This script is run from the PSSP NIM customization script (**pssp_script**) after the node's AIX and PSSP software have been installed, but before the node has been rebooted. This script is run in a limited environment where not all services are fully configured. Because of this limited environment, you should restrict your use of **script.cust** to function that must be performed prior to the post-installation reboot of the node.

firstboot.cust This script is run during the first boot of the node immediately after it has been installed. This script runs in a more "normal" environment where most all services have been fully configured. This script is a preferred location for node customization functions that do not require a reboot of the node to become fully enabled.

firstboot.cmds
When in restricted root access mode and secure remote command mode, this sysctl script is run on the control workstation during node installation to copy critical files from the control workstation to the nodes. It is enabled in the **firstboot.cust** script. See the **firstboot.cmds** and **firstboot.cust** files for information on how to set up and enable this script for sysctl.

Note: Your security environment is not set up during **script.cust** processing. If you are using AIX remote commands or SP Trusted Services, perform your customization during **firstboot.cust** processing. See "Appendix E. User-supplied node customization scripts" on page 297 for additional information.

See "Appendix E. User-supplied node customization scripts" on page 297 for more detailed information on:

- The run-time environment for each of these scripts
- How to create and where to place the scripts

"Appendix E. User-supplied node customization scripts" on page 297 also discusses migration and coexistence issues and techniques to use the same set of customization scripts across different releases and versions of AIX and PSSP.

Note: When PSSP installs a node, it uses the AIX **sysdumpdev -e** command to estimate the size of the dump for the node. PSSP creates a dump logical volume that is approximately 10 percent larger than the estimated dump size, and makes that logical volume the primary dump device. However, you may find that this dump device is not large enough to contain an entire dump due to large processes or applications running on your node.

Once your node is up and running, use:

sysdumpdev -e	To get the estimated size of the node's dump
sysdumpdev -l	To find the name of the primary dump device
lslv	To list the amount of space available in the primary dump device
extendlv	To expand the size of the dump logical volume if the estimated dump space is greater than the dump space available

There are special considerations that you must take into account if you are installing your system with the following security setup:

```
splstdata -p
List System Partition Information
...
auth_install k4
auth_root_rcmd k4
auth_methods k5:k4:std
ts_auth_methods compat
```

Because **auth_install** does not contain DCE, you must ensure that DCE is installed on the nodes before **psspfb_script** sets the authentication methods during the install process. This same requirement existed for PSSP 3.1, so you may have already implemented a process to do mksysb installs. To mksysb install a node, you will need to add code to your **/tftpboot/script.cust**. The new code in **script.cust** will need to mount the directory containing DCE and install your required DCE clients.

Installing with secure remote command methods enabled

There are special considerations to take into account if you are going to install nodes with secure remote command methods enabled. See "Step 30: Enter site environment information" on page 49 and *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for additional information.

When the node is installed, the secure remote command software must also be installed, configured, and the daemon started. The root public keys must be copied from the control workstation to the node and from the boot/install server nodes to nodes that they serve and from the boot/install server nodes to the control workstation to enable the PSSP installation and configuration scripts to be able to run the secure remote commands from the control workstation and any other BIS nodes to the nodes being installed.

To enable the secure remote command software on the nodes during the node installation, the **/tftpboot/script.cust** file must be edited to install the secure remote command software and move the root public keys to the nodes. Examples are

shipped in the **script.cust** sample file with PSSP 3.4. The **script.cust** file also adds the start of the daemon to **/etc/inittab** to ensure that the secure remote command daemon is restarted after any node reboot.

The PSSP code must be able to issue secure remote commands and copies to the nodes without being prompted for passwords or passphrases during installation and configuration.

If additional files must be copied to the nodes during the installation process with secure remote command and restricted root remote commands enabled, the **firstboot.cmds** sample file gives examples of how to enable the copy from the control workstation to the nodes in the restricted access environment and in the secure remote command enabled environment.

Step 62: Set up the switch

If you do not have a switch, skip this step and proceed to “Step 65: Set up system partitions (SP Switch or switchless systems only)” on page 97.

The optional switch connects all the nodes in the system to increase the speed of internal system communications. It supports the high volume of message passing that occurs in a parallel environment with increased bandwidth and low latency.

The switch includes software called the **Worm** which verifies the actual switch topology against an anticipated topology as specified in the switch topology file. This file tells the Worm your switch configuration. You create this file by copying one of the default topology files provided for each SP configuration.

The Worm verifies the switch connections beginning at a node designated as the primary node. By default, the primary node is the first node in the system or the partition. You can override the default and designate another node as the primary node. You *must* do this if the first node is not operational.

In addition to the primary node, a primary backup node exists that will take over for the primary node when it detects that the primary node is no longer functional. The primary backup node passively listens for activity from the primary node. When the primary backup node detects that it has not been contacted by the primary node for a specified amount of time, it assumes the role of the primary node. This takeover involves nondisruptively reinitializing the switch fabric, selecting another primary backup, and updating the SDR. By default, a node is selected from a frame that is different from the primary node. If no other frame exists (for example, a single frame system), a node is selected from a switch chip that is different from the primary node. If no other switch chip is available, any available node on the switch is selected. By default, the backup node is the last node in the system or the partition.

Step 62.1: Select a topology file

Select the correct switch topology file by counting the number of node switch boards (NSBs) and intermediate switch boards (ISBs) in your system, then apply these numbers to the naming convention. If you have an SP Switch2 two plane system, count only the number of NSBs and ISBs on one plane. The switch topology files are in the **/etc/SP** directory on the control workstation.

NSBs are switches mounted in slot 17 of frames containing nodes or SP Switch2 switches mounted in slots 2 through 16 of frames designed as multiple NSB frames. Multiple NSBs are used in systems that require a large number of switch connections for SP-attached servers or clustered enterprise server configurations.

ISBs are switches mounted in the switch frame. ISBs are used in large systems, where more than four switch boards exist, to connect many processor frames together. SP-attached servers never contain a node switch board, therefore, never include non-SP frames when determining your topology files.

The topology file naming convention is as follows:

`expected.top.`*`NSBnum`*`nsb.`*`ISBnum`*`isb.`*`type`*

where:

- *NSBnum* is the number of NSBs in the configuration
- *ISBnum* is the number of ISBs in the configuration
- *type* is the type of topology. The default type is 0.

For example, **expected.top.2nsb.0isb.0** is a file for a two frame and two switch system with no ISB switches.

The exception to this naming convention is the topology file for the SP Switch-8 configuration, which is **expected.top.1nsb_8.0isb.1**.

See the **Etopology** command in *PSSP: Command and Technical Reference* for additional information on topology file names.

Step 62.2: Managing the switch topology files

The switch topology file must be stored in the SDR. The switch initialization code uses the topology file stored in the SDR when starting the switch (**Estart**). When the switch topology file is selected for your system's switch configuration, it must be annotated with **Eannotator**, then stored in the SDR with **Etopology**. The switch topology file stored in the SDR can be overridden by having an **expected.top** file in **/etc/SP** on the primary node. **Estart** always checks for an **expected.top** file in **/etc/SP** before using the one stored in the SDR. The **expected.top** file is used when debugging or servicing the switch.

Notes:

1. Be aware that **Estart** distributes the topology file to all the nodes in the system partition on the switch. In the case of **expected.top**, this is significant because if the topology file is left on a node and the primary is changed to that node, the topology file will be used. If you have an **expected.top** file in **/etc/SP** on any of the nodes, make sure that you remove it when it is no longer needed.
2. Depending upon your configuration, the first **Estart** of the switch may take longer than subsequent **Estart**s.
3. In a two plane SP Switch2 system, the function of **/etc/SP/expected.top** is taken over by **/etc/SP/expected.top.p0** for plane 0 and by **/etc/SP/expected.top.p1** for plane 1.

Step 62.3: Annotating a switch topology file

Annotate a switch topology file before storing it in the SDR. Refer to the following table for instructions.

If using:	Do this:	
Perspectives	**SELECT**	The **smit cluster_mgmt** icon • The RS/6000 SP Cluster Management menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE**	**smit cluster_mgmt** • The RS/6000 SP Cluster Management menu appears.
	SELECT	Perform Switch Operations • The Perform Switch Operations menu appears.
	SELECT	Topology File Annotator • The Topology File Annotator menu appears.
	SELECT	Topology File Annotator using File Selection • The Topology File Annotator using File Selection menu appears.
	PRESS	**List**
	SELECT	The appropriate topology file (for example, **/etc/SP/expected.top.2nsb.0isb.0**)
	TYPE	The data in the fields, as follows: • The fully-qualified name of the file in which you want to store the annotated topology file (for example, **/etc/SP/expected.top.annotated**) • **yes** to store the topology file in the SDR
	PRESS	**Ok**

Eannotator	Use **Eannotator** to update the switch topology file's connection labels with their correct physical locations. Use the **-O yes** flag to store the switch topology file in the SDR. If **−p** is not specified, the default behavior is to perform this action on all planes. Using **Eannotator** makes the switch hardware easier to debug because the switch diagnostics information is based on physical locations.
	
	For example, to annotate a two-switch or maximum 32-node system, enter:
	`Eannotator -F /etc/SP/expected.top.2nsb.0isb.0 \`
	`-f /etc/SP/expected.top.annotated -O yes`	

Step 62.4: Storing the switch topology file in the SDR

If you entered **Eannotator -O yes** or **yes** on the Topology File Annotator menu in "Step 62.3: Annotating a switch topology file" on page 93, skip this step.

If using:	Do this:	
Perspectives	**SELECT**	The **smit cluster_mgmt** icon • The RS/6000 SP Cluster Management menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE**	**smit cluster_mgmt** • The RS/6000 SP Cluster Management menu appears.
	SELECT	Perform Switch Operations • The Perform Switch Operations menu appears.
	SELECT	Fetch / Store Topology Files • The Fetch / Store Topology Files menu appears.
	SELECT	Store a Selected Topology File • The Store a Selected Topology File menu appears.
	PRESS	**List**
	SELECT	The appropriate topology file.
	PRESS	**Ok**
Etopology	Use **Etopology** to store the switch topology file in the SDR and make sure that it has been annotated. If **–p** is not specified, the default behavior is to perform this action on all planes. For example, to store the annotated topology file expected.top.annotated in the current directory, enter: `Etopology expected.top.annotated`	

Step 63: Verify the switch primary and primary backup nodes

Frame 1, node 1 is the default oncoming primary node for the switch.

A node type exists called the primary backup node for the switch. The primary backup node passively listens for activity from the primary node. When the primary backup node detects that it has not been contacted by the primary node for a specified amount of time, it assumes the role of the primary node. This takeover involves nondisruptively reinitializing the switch fabric, selecting another primary backup, and updating the SDR. The default is the last node in the frame, not the last node slot. For partitions, the default primary is the first node and the default backup is the last node in the partition. You must override this selection if the node slot is not operational. Use SMIT or the **Eprimary** command to verify this node or change the primary to another node.

The oncoming primary and backup nodes should not both be assigned to partitions on a single pSeries 690 server, if alternatives exist.

If using:	Do this:	
Perspectives	**SELECT**	The **smit cluster_mgmt** icon • The RS/6000 SP Cluster Management menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE**	**smit cluster_mgmt** • The RS/6000 SP Cluster Management menu appears.
	SELECT	Perform Switch Operations • The Perform Switch Operations menu appears.
	SELECT	Set Primary/Primary Backup Node • The Set Primary and Primary Backup Node menu appears.
	PRESS	**Ok** to show defaults.
	PRESS	**Cancel** to return to the Set Primary and Primary Backup Node Menu.
	PRESS	**Cancel** to keep defaults
	or	
	ENTER	A different primary node
	PRESS	**Ok**
Eprimary	Enter: `Eprimary [new_primary_node] [-backup new_primary_backup_node_number]` If **–p** is not specified, the default behavior is to perform this action on all planes.	

The **Eprimary** command, without any parameters, returns the node number of the current primary node, the primary backup node, the oncoming primary node, and the oncoming primary backup node.

Step 64: Set the switch clock source for all switches (SP Switch only)

Use SMIT or the **Eclock** command to initialize the switch's clock source. The SMIT and **Eclock** interfaces require that you know the number of Node Switch Boards (NSBs) and Intermediate Switch Boards (ISBs) in your RS/6000 SP system.

Select the **Eclock** topology file from the control workstation's **/etc/SP** subdirectory, based on these numbers. For example, if your RS/6000 SP system has six node switch boards and four intermediate switch boards, you would select **/etc/SP/Eclock.top.6nsb.4isb.0** as an **Eclock** topology file.

See *PSSP: Command and Technical Reference* for the **Eclock** topology file names.

If using:	Do this:	
Perspectives	**SELECT**	The **smit cluster_mgmt** icon • The RS/6000 SP Cluster Management menu appears.
	SELECT	Perform Switch Operations From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE**	**smit cluster_mgmt** • The RS/6000 SP Cluster Management menu appears.
	SELECT	Perform Switch Operations • The Perform Switch Operations menu appears.
	SELECT	Change/Show Switch Clock Source Settings (SP Switch Only) • The Change/Show Switch Clock Source Settings menu appears.
	SELECT	Initialize Switch Clock Source Settings (SP Switch Only) • The Initialize Switch Clock Source Settings menu appears.
	PRESS	**List**
	SELECT	The correct **Eclock** file
	PRESS	**Ok**

Eclock	Use the **Eclock** command to set the switch's clock source for all switches.
	For example, if your RS/6000 SP system has six node switch boards and four intermediate switch boards, select **/etc/SP/Eclock.top.6nsb.4isb.0** as an **Eclock** topology file. Enter:
	`Eclock -f /etc/SP/Eclock.top.6nsb.4isb.0`
	This command sets the proper clock source settings on all switches within a 96-way (6 nsb, 4 isb) RS/6000 SP system.
	To verify the switch configuration information, enter:
	`splstdata -s`	

Step 65: Set up system partitions (SP Switch or switchless systems only)

This step is optional. The PSSP installation code sets up a default system partition configuration to produce an initial, single-system partition including all nodes in the system. This system partition is created automatically. If you do not want to divide your system into partitions, continue with "Step 66: Configure the control workstation as the boot/install server" on page 98.

Note: System partitioning is not supported on a switchless clustered enterprise server.

If you want to partition your system, you can select an alternate configuration from a predefined set of system partitions to implement before booting the nodes or you can use the System Partitioning Aid to generate and save a new layout. Follow the procedure described in the "Managing system partitions" chapter in *PSSP: Administration Guide* and refer to information in the "The System Partitioning Aid" section of the "Planning SP system partitions" chapter in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*. You do not have to partition your system now as part of this installation. You can partition it later.

For information on how to set a security setting in an established system partition, see “Chapter 5. Reconfiguring security” on page 179.

Step 66: Configure the control workstation as the boot/install server

This step uses the information entered in the previous steps to set up the control workstation and optional boot/install servers on nodes. It configures the control workstation as a boot/install server and configures the following options (when selected in your site environment):

- Automounter
- File Collections
- NTP
- User Management
- Accounting

You can perform this step more than once. If you encounter any errors, see *PSSP: Diagnosis Guide* for further explanation. After you correct your errors, you can start the task again.

In previous releases of PSSP, most of the installation function which configured boot/install servers and clients was performed in the single program called **setup_server** which you could run by issuing the **setup_server** command. This is still the suggested way for configuring the control workstation. For more experienced system administrators, IBM has provided a set of Perl scripts you can issue to also configure the control workstation that enable you to diagnose how the **setup_server** program is progressing. For more information, refer to “Appendix D. Boot/install server configuration commands” on page 295.

Note: If you are using AFS, you cannot run **setup_server** from Perspectives or SMIT.

If using:	**Do this:**
Perspectives	**SELECT** **smit cluster_mgmt** from the Launch Pad. • The RS/6000 SP Cluster Management menu appears. **SELECT** Run setup_server Command
SMIT	**TYPE** **smit enter_data** • The Enter Database Information window appears. **SELECT** Run setup_server Command
setup_server	Enter: `setup_server` with no parameters.
The first time **setup_server** runs, depending upon your configuration, it can take a significant amount of time to configure the control workstation as a NIM master.	

Step 67: Verify that the System Management tools were correctly installed

This step directs you to run a verification test that checks for correct installation of the System Management tools on the control workstation.

If using:	Do this:	
Perspectives	**SELECT**	**smit SP_verify on CWS** from the Launch Pad. • The RS/6000 SP Installation/Configuration Verification menu appears. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE**	**smit SP_verify** • The RS/6000 SP Installation/Configuration Verification Menu appears.
	SELECT	System Management
sysman	Enter: `SYSMAN_test`	

After the tests are run, the system creates a log in **/var/adm/SPlogs** called **SYSMAN_test.log**.

Note: After performing this step, you can ignore any messages that you receive about the number of nodes tested. Since nodes are not available during this operation, they will not be tested.

See *PSSP: Diagnosis Guide* for information about what this test does and what to do if the verification fails.

Task E. Power on and install the nodes

This section describes the steps you take to power on the nodes to be installed.

Step 68: Network boot optional boot/install servers

Note: If you have set up system partitions, do this step in each partition.

Follow the instructions in the following table to network boot the optional boot/install servers which install and customize the nodes you selected. To monitor installation progress by opening the node's read-only console, issue:

`s1term` *`frame_id slot_id`*

If you have more than eight boot/install servers on a single Ethernet segment, you should network boot those nodes in groups of eight or less. See the "IP performance tuning" section in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more information.

Notes:

1. For systems with boot/install server nodes, the **tuning.cust** file must first be propagated to the server node before attempting to propagate it to the client node.
2. For MCA nodes, the **nodecond** command remotely processes information from the initial AIX firmware menus. You should not change the language option on these menus. The language must be set to English in order for the **nodecond** command to run properly.

If using:	Do this:
Perspectives	**SELECT** The **Hardware Perspective** icon by double clicking **SELECT** The Nodes pane **SELECT** Actions → LCD and LED display • The LCD and LED display appears. **SELECT** Nodes to be netbooted **SELECT** Actions → Network Boot **SELECT** Apply • All selected nodes are booted. If you had the Hardware Perspective up before you added the required node information, you should delete and re-add the Nodes Pane. If you had the Hardware Perspective up before you partitioned your system, you should delete and re-add the CWS/System/Syspars Pane and then delete and re-add the Nodes Pane.

nodecond	Enter:
	`nodecond` *`frame_id slot_id`* `&`
	Enter:
	`spmon -Led node`*`node_number`*
	or
	`spled &`
	to check the LCD and LED display for each node.

Network installation progress

When a network installation is in progress, the LED for the nodes involved show various values. These values indicate the installation stage. Since the node installation process can be long, it is hard to determine where you are in that process. Refer to *PSSP: Diagnosis Guide* for a complete list of PSSP-specific LED values.

Step 69: Verify that System Management tools were correctly installed on the boot/install servers

Now that the boot/install servers are powered up, run the verification test from the control workstation to check for correct installation of the System Management tools on these nodes.

If using:	Do this:
Perspectives	**SELECT** **smit SP_verify on CWS** from the Launch Pad. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit SP_verify** • The RS/6000 SP Installation/Configuration Verification Menu appears. **SELECT** System Management
SYSMAN_test	Enter: `SYSMAN_test`

After the tests are run, the system creates a log in **/var/adm/SPlogs** called **SYSMAN_test.log**.

See the section on “Verifying System Management installation” in *PSSP: Diagnosis Guide* for information on what this test does and what to do if the verification test fails.

Step 70: Network boot the remaining RS/6000 SP nodes

Note: If you have set up system partitions, do this step in each partition.

Repeat the procedure used in “Step 68: Network boot optional boot/install servers” on page 99 to network boot and install, or customize the remaining nodes. You may need to ensure that all **setup_server** processes have completed on the boot/install nodes prior to issuing a network boot on the remaining nodes. Refer to the **/var/adm/SPlogs/sysman/***node***.console.log** file on the boot/install node to see if **setup_server** has completed.

If any of your boot/install servers have more than eight clients on a single Ethernet segment, you should Network Boot those nodes in groups of eight or less. See the “IP performance tuning” section in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more information.

Using a token ring-bridge gateway

If you are using a token ring through a bridge as your default gateway to your nodes and the token ring bridge is not on the same segment as your LAN, you must change the value of the broadcast field in the ODM for each node. The default value is set to **No** (confine broadcast to local token-ring) each time you install or customize a node. However, when you boot the nodes with this bridge setup, the network is unusable.

If using:	**Do this:**	
SMIT	**TYPE**	**smit chinet** • The Available Network Interfaces menu appears.
	SELECT	tr0 Token Ring Network Interface
	TYPE	`Yes` in the Confine BROADCAST to LOCAL Token-ring field.
chdev	Enter: `chdev -P -l tr0 -a allcast=off`	

Step 71: Verify node installation

If you have set up system partitions, select a global view to do this step.

Check hostResponds and powerLED indicators for each node.

If using:	Do this:
Perspectives	**SELECT** The Hardware Perspective icon by double clicking • The Hardware Perspective appears. **SELECT** The Nodes Pane **SELECT** View → Show Objects in Table View. • The Set Table Attributes dialog box appears. **SELECT** The Power table column. **SELECT** The hostResponds column while pressing the CTRL key on the keyboard. **PRESS** **Ok** • All of the table attributes should now be colored green. Check the table cells. All of the table cells in the Power and hostResponds columns should be solid green. If any are not, see the section in *PSSP: Diagnosis Guide*.
spmon	Enter: `spmon -d -G`

Step 72: Verify node expansion configuration information (optional)

Use the **splstdata -x** command to verify that the connection information for an SP expansion I/O unit is correct. For example, issue:

```
splstdata -x
```

Step 73: Enable s1_tty on the SP-attached server (SAMI hardware protocol only)

If you just installed an SP-attached server, you must ensure that the s1_tty is enabled on the server. Until the login is enabled on the tty, the **s1term** command from the control workstation to the SP-attached server will not work.

On the SP-attached server, determine which tty is mapped to 01-S1-00-00. For example, issue the following:

```
lsdev -C -c tty0
```

In response, the system displays something similar to:

```
tty0 Available 01-S1-00-00 Asynchronous Terminal
tty1 Available 01-S2-00-00 Asynchronous Terminal
```

In the previous example, tty0 is mapped to 01-S1-00-00.

Set the login to enable. For example, issue the following:

```
chdev -l tty0 -a login=enable
```

Step 74: Update authorization files in restricted mode for boot/install servers (optional)

Once the node is installed, or as part of **firstboot.cust**, the remote command authorization files on the node serviced by the non-control workstation boot/install server need to be updated depending on the setting of auth_root_rcmd. Add the following:

standard	An entry for the boot/install server node host name in **/.rhosts**
k4	An entry for the boot/install server node rcmd principal in **/.klogin**
dce	An entry for the **self-host** and the **spbgroot** principal for the **/.k5login** boot/install server node

Step 75: Run verification tests on all nodes

If you have set up system partitions, do this step in each partition.

This step directs you to run a series of verification tests that check for correct installation of your selected options on all the nodes.

If using:	Do this:	
Perspectives	**SELECT**	**smit SP_verify on CWS** from the Launch Pad. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE**	**smit SP_verify** • The RS/6000 SP Installation/Configuration Verification Menu appears.
	SELECT	System Management
SYSMAN_test	Enter: `SYSMAN_test`	

After the tests are run, the system creates a log in **/var/adm/SPlogs** called **SYSMAN_test.log**.

See the section on “Verifying System Management installation” in *PSSP: Diagnosis Guide* if the verification test fails.

Step 76: Start the optional switch

- Do this step if you have installed a switch.
- If you have set up system partitions, do this step in each partition (SP Switch only).

If using:	Do this:
Perspectives	**SELECT** The **smit cluster_mgmt** icon • The SP Cluster Management menu appears. **SELECT** Perform Switch Operations From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit switch_ops** • The Perform Switch Operations menu appears. **SELECT** Start Switch
Estart	Enter: `Estart` If **−p** is not specified, the default behavior is to perform this action on all planes.

Step 77: Verify that the switch was installed correctly

- Do this step if you have installed a switch.
- If you have set up system partitions, verify the switch from within each partition (SP Switch only).
- Check all connectors for miscabling and node communications.

Run a verification test to ensure that the switch is installed completely.

If using:	Do this:
Perspectives	**SELECT** **smit SP_verify on CWS** from the Launch Pad. From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE** **smit SP_verify** • The RS/6000 SP Installation/Configuration Verification Menu appears. **SELECT** The Communication Subsystem option **PRESS** Enter
CSS_test	Enter: `CSS_test`

After the tests are run, the system creates a log in **/var/adm/SPlogs** called **CSS_test.log**.

If the verification test fails, see the section on “Diagnosing switch problems” in *PSSP: Diagnosis Guide*.

Check the switchResponds and powerLED indicators for each node.

If using:	Do this:
Perspectives	**SELECT** The **Hardware Perspective** icon by double clicking **SELECT** The Nodes Pane • The Nodes Pane receives focus. **SELECT** View → Set Monitoring • The Set Monitoring for Nodes dialog box appears. **SELECT** The switchResponds condition **PRESS** Apply • The nodes should now be colored green. Check the node icons. All node icons should be solid green. If they are not, refer to *PSSP: Diagnosis Guide* for further information.
spmon	Enter: `spmon -d -G`

Step 78: Create DCE principals for the switch adapter host name (optional)

To allow **k5** remote command operation between nodes on the switch adapter, additional DCE configuration is needed.

1. You must be root to perform this task.
2. All adapters must have an **ftp** and a **host** account defined in the DCE database.
3. To add adapters to a DCE-configured node, perform the following steps:
   a. Login to the node or use **dsh**
   b. Run **kerberos.dce -type local**

Refer to *IBM Distributed Computing Environment for AIX: Administration Commands Reference* for more information on the **kerberos.dce** command.

Step 79: Tune the network adapters

Various models of the network adapters can have different values for transmit queue sizes. To get peak performance out of the network adapters, increase the transmit and receive queue sizes to their maximum. See “Step 5: Tune all control workstation network adapters” on page 16 for valid queue size settings.

If the adapter you are changing is also the adapter for the network you are logged in through, you will have to make the changes to the database only. Then reboot the node for the changes to become effective.

If using:	Do this:	
Perspectives	**SELECT**	**smit devices** from the Launch Pad. • The Devices menu appears From this point, you can follow the rest of the SMIT steps described in the next row of this table.
SMIT	**TYPE**	**smit devices** • The Devices menu appears
	SELECT	Communication • The Communication menu appears
	SELECT	The adapter you want to reset (for example, `Ethernet Adapter`) • The Ethernet Adapter menu appears
	SELECT	Adapter • The Adapter menu appears
	SELECT	Change / Show Characteristics of an Ethernet Adapter • An Ethernet Adapter window appears
	SELECT	An adapter from the list shown • The Change / Show Characteristics of an Ethernet Adapter window appears
	CHANGE	The HARDWARE TRANSMIT queue size. (See note 2 below.)
	CHANGE	`Apply change to DATABASE only` to **yes**. (Press F4 and select yes.)
	PRESS	**Ok** to apply the changes to the database.
chdev	Enter `chdev -P -l ent0 -a xmt_que_size=512`	

1. You must reboot the node in order for the changes to take effect.
2. To determine the name of the TRANSMIT queue size, issue:

   `lsattr -l` *`adapter_name`* `-E`

   In response, a list of all of the values for the different variables, what they are, and the name of the variable will be displayed. In the previous example, an MCA adapter with AIX 4.2.1 or later was used.

Step 80: Authorize the SP administrative principals for remote command access

If you created an SP administrative principal in “Step 22.3: Create SP administrative principals” on page 45, you may also need to authorize that principal to remotely run commands on nodes within the system. This is to ensure that SP commands which both write to the SDR and remotely access the nodes (via AIX and PSSP remote commands like **dsh** and **rsh**) are authorized to perform such actions. The SP commands that require this type of authorization include **chauthpar**, **chauthpts**, and **spsitenv**.

You must authorize the SP administrative principal to do remote commands by adding it to the **/.k5login** file on the control workstation and on the nodes. The entry format is:

```
principal@dce_cell_name
```

Note: If you are running with a secure remote command method enabled, you do not have to perform this task. See “Step 30: Enter site environment information” on page 49.

Step 81: Apply PSSP PTFs to nodes (optional)

Software maintenance (PTFs) may now be applied to the **ssp** and **rsct** file sets installed on the nodes. If you installed software maintenance to the control workstation in “Step 27: Apply PSSP PTFs (optional)” on page 48, you must install the service updates to the nodes now. Refer to “Installing updates on a per node basis” on page 264 for detailed instructions.

Run post-installation procedures

Now that your RS/6000 SP System is installed and ready to run, you can set up any additional tools you want to use to manage your parallel environment.

You can find procedures for setting up these facilities as follows:

Table 4. Information about RS/6000 SP Facilities

For RS/6000 SP Facility	**Refer to:**
Error Log	Error logging information in *PSSP: Diagnosis Guide*
LoadLeveler	*IBM LoadLeveler for AIX 5L: Using and Administering*
Login Control	“Managing user accounts” chapter in *PSSP: Administration Guide*
IBM Virtual Shared Disks	*PSSP: Managing Shared Disks*
Sysctl	“Controlling remote execution by using Sysctl” chapter in *PSSP: Administration Guide*

Chapter 3. Installing and configuring the High Availability Control Workstation

Restrictions

1. You cannot use both DCE authentication and HACWS.
2. You cannot use IPv6 aliasing with DCE, HACMP, and HACWS.
3. You cannot use HACWS when you specify **none** for AIX commands.

This chapter discusses how to install and configure the High Availability Control Workstation (HACWS).

The information in this chapter is optional—you do not have to install HACWS in order to run your SP system. Install HACWS if you want to do any of the following:

- Continue running your SP system after a control workstation failure.
- Shut down the control workstation for deferred hardware and software maintenance without having a system outage.
- Maintain the SP system function and reliability when the control workstation fails.
- Fail over the control workstation to a backup.

You can add an HACWS configuration at any time in the life of your SP system.

The HACWS installation procedure requires that you install and configure High Availability Cluster Multi-Processing for AIX (HACMP) on the primary and backup control workstations. (See “Task E. Configure High Availability Cluster Multi-Processing” on page 120.) You may use any level of HACMP that is supported with the level of AIX that you are using. Refer to the appropriate HACMP documentation to determine which levels of HACMP are supported with the level of AIX that you are using. At the time this manual was published, PSSP 3.4 was supported only with AIX 4.3.3 and AIX 5L 5.1. AIX 4.3.3 was supported with HACMP 4.4, HACMP 4.4.1, HACMP Enhanced Scalability (HACMP/ES) 4.4, and HACMP/ES 4.4.1. AIX 5L 5.1 was supported with HACMP 4.4, HACMP 4.4.1, and HACMP/ES 4.4.1.

HACMP runs on two control workstations in a two-node rotating resource configuration. There are external disks and a dual RS-232 frame supervisor card with a connection from each control workstation to each SP frame. The external disks are not accessed concurrently. The control workstation configuration provides automated detection, notification, and recovery of control workstation failures. For more information about HACMP, see *HACMP for AIX: Concepts and Facilities*.

Notes:

1. Before you begin, you should carefully plan your HACWS installation. This will help make your installation more efficient and successful. For more information about planning for HACWS installation, see the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*.
2. There is no hardware control or serial terminal cabling available from the backup control workstation to support SP-attached servers in an HACWS installation. This means that there is limited function from the backup control workstation when a failover occurs. Refer to the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more details.

Keep these points in mind as you install HACWS:

- You must install your SP system with a regular control workstation first. (The original control workstation becomes the primary control workstation in your HACWS configuration.)
- You must install HACMP on both control workstations.
- HACWS supports only a single backup control workstation.
- Standby adapters on the primary control workstation are *not* required.
- The external file system can be mounted over either the **/spdata** or the **/spdata/sys1** directories.

Finding related information

You can find more information about HACMP and HACWS in the following books:

- *HACMP for AIX: Concepts and Facilities*
- *HACMP for AIX: Planning Guide*
- *HACMP for AIX: Installation Guide*
- *HACMP for AIX: Administration Guide*
- *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*
- *PSSP: Administration Guide*
- *Implementing High Availability on RISC/6000 SP (Redbook),* SG24-4742.

If you are using HACMP/ES and HACWS, refer to the following books:

- *HACMP for AIX: Concepts and Facilities*
- *HACMP for AIX: Planning Guide*
- *HACMP for AIX: Troubleshooting Guide*
- *HACMP for AIX: Enhanced Scalability Installation and Administration Guide*
- *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*

Task A: Prepare the control workstations

This section describes the steps you take to prepare the control workstations.

Step 1: Understand the procedure

Read and become familiar with the following publications:

- *HACMP for AIX: Concepts and Facilities*
- *HACMP for AIX: Planning Guide*
- *HACMP for AIX: Installation Guide*
- *HACMP for AIX: Administration Guide*
- *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*
- *PSSP: Administration Guide*

You must be familiar with HACMP terminology, including the following:

- Service address
- Boot address
- Network interface

Step 2: Plan network configuration

Planning for your High Availability Control Workstation is essential before proceeding with the installation steps. Be sure to read the “Planning for a high availability control workstation” chapter in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*. It describes control workstation network configuration requirements which are necessary for HACWS to function correctly. It also describes a sample network configuration which is referenced by the installation steps later in this chapter.

Step 3: Install the SP system

If you have not done so already, install your SP system (single control workstation, frames, and nodes) as described in “Chapter 2. Installing and configuring a new RS/6000 SP system” on page 11). The SP system must be completely installed and operating correctly before beginning HACWS installation.

After you complete the rest of the instructions in this chapter, the original control workstation will become the primary control workstation, and the backup control workstation will be added to your system.

Step 4: Install AIX on the backup control workstation

If you have not done so already, install base AIX on the backup control workstation. For more information, refer to the *IBM AIX Installation Guide*.

Step 5: Back up the control workstations

Create mksysb images of both the primary control workstation and the backup control workstation, and keep them on hand in case the HACWS installation fails. **Restoring a mksysb image is the only way you can recover from a failed HACWS installation**.

Also back up the **/spdata** file system if it is not located on the root volume group.

Step 6: Set up the hardware

Perform all hardware setup. Refer to the *RS/6000 SP: Planning, Volume 1, Hardware and Physical Environment* for instructions. The use of a serial network with HACMP is strongly suggested. The serial connection can be either a SCSI-2 Differential bus using target mode SCSI, or a raw RS-232 serial line. Refer to the *HACMP for AIX: Installation Guide* for full details.

To set up the external disks, follow the instructions in the “Installing shared disk devices” chapter of the *HACMP for AIX: Installation Guide* or the “Planning shared disk devices” chapter of the *HACMP for AIX: Enhanced Scalability Installation and Administration Guide.*

Step 7: Configure RS-232 control lines

Each SP frame in your SP system requires a serial port on the control workstation configured to accommodate the RS-232 line. SP-attached servers require two serial ports.

You already performed this step on the primary control workstation. You need to perform the same step on the backup control workstation. Each frame must be connected to the exact same tty port on both control workstations. If not cabled correctly, the hardware monitor will not bring all of the frames online.

Refer to “Step 4: Configure RS-232 control lines” on page 15 for more instructions.

There is no hardware control or serial terminal cabling available from the backup control workstation to support SP-attached servers in an HACWS installation. This means that there is limited function from the backup control workstation when a failover occurs. Refer to the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more details.

Step 8: Set authentication methods for AIX remote commands on the backup control workstation

The authentication method on the backup control workstation must match the authentication method configured on the primary control workstation. On the primary control workstation, issue the **lsauthent** command to determine the authentication method setting. To change this setting, issue the **chauthent** command on the backup control workstation.

Step 9: Install PSSP on the backup control workstation

You must now install the PSSP file sets on the backup control workstation to match the PSSP software that is already installed on the primary control workstation. Refer to “Task B. Install PSSP on the control workstation” on page 25 for descriptions of the PSSP file sets. You can perform file set installation from the command line by issuing the **installp** command or you can use SMIT by issuing the **smit install_latest** command. Refer to the *IBM AIX Installation Guide* for more details on how to use these commands.

Notes:

1. You will complete the installation of PSSP later, so do not run any PSSP configuration commands. (Do not run the **install_cw** command).
2. You should not install the **ssp.hacws** file set at this time. You will do this later in “Step 17: Install the HACWS image on both control workstations” on page 117.

Task B. Update Kerberos V4 SP authentication services on the primary control workstation

This section describes the steps you take to update Kerberos V4 on the primary control workstation.

Step 10: Add the Kerberos V4 principal

Use this procedure to add principals for all the primary boot addresses (if the principals do not already exist.)

Some of the network interfaces configured on a regular control workstation become service addresses in the HACWS configuration. For example, a control workstation named **sp_cws** would have a network interface by the same name. When the SP system becomes an HACWS configuration, **sp_cws** becomes a service address. Since the service addresses in a rotating configuration rotate with their resource groups, the **sp_cws** network interface moves back and forth between the primary and backup control workstations.

When the **sp_cws** network interface is on the backup control workstation, the network adapter on the primary control workstation is known by an alternate name, such as **sp_cws_bt**. This alternate name is the boot address. The primary boot addresses need to be identified to Kerberos V4 so the backup control workstation can access authenticated services on the primary while the backup control workstation is acting as the system control workstation.

This example shows the procedure you should follow to add the Kerberos V4 principal **rcmd**, instance *sp_cws_bt* on the primary control workstation. Run the **/usr/kerberos/etc/kdb_edit** program as follows:

```
Opening database...
```

```
Enter Kerberos master key: kerberosMasterPassword
Previous or default values are in [brackets] ,
enter return to leave the same, or new value.

Principal name: rcmd

Instance: sp_cws_bt

<Not found>, Create [y] ? <Enter>
Principal: rcmd, Instance: sp_cws_bt, kdc_key_ver: 1

New Password: rcmdPassword Verifying, please re-enter New Password:
rcmdPassword
Principal's new key version = 1
Expiration date (enter yyyy-mm-dd) [ 2000-04-28 ] ? <Enter>
Max ticket lifetime (*5 minutes) [ 255 ] ? <Enter>
Attributes [ 0 ] ? <Enter>
Edit O.K.
Principal name: <Enter>
#

# <end_of_example>
```

Step 11: Add the Kerberos V4 rcmd service key

This example shows the procedure you should follow to add the Kerberos V4 **rcmd** service key for each primary control workstation boot address.

Run the **/usr/lpp/ssp/kerberos/bin/ksrvutil add** command as follows:

```
Name: rcmd

Instance: sp_cws_bt

Realm: XYZ.COM

Version number: 1

New principal: rcmd.sp_cws_bt@XYZ.COM; version   1

Is this correct? (y,n) <Enter>

Password: RcmdPassword Verifying, please re-enter Password: RcmdPassword

Key successfully added.

Would you like to add another key? (y,n) n

Old keyfile in /etc/krb-srvtab.old. # # <end_of_example>
```

Task C. Update Kerberos V4 SP authentication services on the backup control workstation

This section describes the steps you take to update Kerberos V4 on the backup control workstation.

Step 12: Configure the backup control workstation as a secondary Kerberos V4 authentication server or client

If the primary control workstation is either a primary or secondary Kerberos V4 authentication server, then the backup control workstation must be configured as a secondary authentication server. If the primary control workstation is an authentication client, the backup control workstation must also be configured as an authentication client.

Follow the instructions in "Step 12.1: Initializing as a secondary Kerberos V4 authentication server" or in "Step 12.2: Initializing as an authentication client system" on page 115 to configure the backup control workstation as either a secondary authentication server or as an authentication client.

Step 12.1: Initializing as a secondary Kerberos V4 authentication server

Note: If you perform this step, do not perform "Step 12.2: Initializing as an authentication client system" on page 115.

The following example illustrates the procedure you should follow to initialize the backup control workstation as a secondary authentication server.

1. Copy the **/etc/krb.conf** file from the primary authentication server to the backup control workstation.
2. Add a line to the **/etc/krb.conf** file on both control workstations, listing the backup control workstation (by its full host name) as a secondary server for the local authentication realm.

   For example, to add **sp2cw.xyz.com** as a secondary server for the authentication realm **XYZ.COM**, add this line to **/etc/krb.conf**:

   ```
   XYZ.COM sp2cw.xyz.com
   ```
3. Copy the **/etc/krb.realms** file from the primary server to the backup control workstation.
4. On the backup control workstation, run the **setup_authent** program.

   setup_authent requires you to login to the authentication service using the same administrative principal name that was defined when the primary server was set up. The remainder of the initialization of authentication services on this secondary local system takes place automatically.

   The following example shows the interaction you can expect when you run **setup_authent** when initializing as a secondary authentication server.

   ```
   #setup_authent
   <screenclear>
   ********************************************************************
   Logging into Kerberos as an admin user

    You must assume the role of a Kerberos administrator <user> to
    complete the initialization of Kerberos on the local system. The k4init
    command is invoked and will prompt you for the password. If you are
    setting up your primary server here, you have just defined it. If you
    have defined multiple administrative principals, or if your primary
    authentication server is on another system, you must first enter the
    name of an administrative principal who has root privilege (UID 0). You
    need to be authenticated as this administrator so that this program
    can create the principals and service key files for the authenticated
    services that run on the SP system.

    For more information, see the k4init man page.
   ********************************************************************
   ```

```
setup_authent: Enter name of admin user: root

Kerberos Initialization for "root.admin"

Password: rootPassword sp2cw.xyz.com: success.

sp2cw.xyz.com:       Succeeded
#
```

Note: The last two messages shown in the example above are issued by the programs that transfer the database from primary to secondary servers, to indicate that the backup database has been installed.

```
# k4list
Ticket file:    /tmp/tkt0
Principal:  root.admin@XYZ.COM

Issued           Expires          Principal
Nov 11 16:26:11  Dec 12 16:26:11  krbtgt.XYZ.COM@XYZ.COM
#
<end_of_example>
```

5. After **setup_authent** completes, add an entry for the secondary authentication server to the **/etc/krb.conf** file on all SP nodes on which you have already initialized authentication.
6. If this is the first secondary authentication server, you should create a root **crontab** entry on the primary authentication server that invokes the script **/usr/kerberos/etc/push-kprop**. This periodically propagates database changes from the primary to the secondary authentication server.

Step 12.2: Initializing as an authentication client system

Note: If you perform this step, do not perform “Step 12.1: Initializing as a secondary Kerberos V4 authentication server” on page 114.

The following example illustrates the procedure you should follow to initialize the backup control workstation as an authentication client.

1. Copy the **/etc/krb.conf** file from the primary authentication server to this system (the backup control workstation).
2. Copy the **/etc/krb.realms** file from the primary server.

   Note: If the new workstation is outside the realm of the primary server, you must add this new workstation to the **/etc/krb.realms** file on the primary server before you copy the **/etc/krb.realms** file from the primary server to the new workstation. Otherwise, the next step will fail.
3. Run the **setup_authent** program.

   setup_authent requires you to login to the authentication service using the same administrative principal name that was defined when the primary server was set up. The remainder of the initialization of authentication services on the local system takes place automatically.
4. The root **.klogin** file on a client workstation contains just the administrative principal name you used to install authentication. You may want to edit the **.klogin** file to add other principals in your configuration.
5. Enter the **/usr/lpp/ssp/kerberos/bin/k4list** command to make sure a ticket exists for the account.

The following example shows the interaction you can expect when you run **setup_authent** when initializing as an authentication client system.

Note: The initial warning message shown in the example is issued if you have installed the **ssp.authent** option on a system configured as a client rather than a server.

```
#setup_authent
<screenclear>
********************************************************************************
Logging into Kerberos as an admin user

You must assume the role of a Kerberos administrator <user>.admin to
complete the initialization of Kerberos on the local system. The k4init
command is invoked and will prompt you for the password. If you are
setting up your primary server here, you have just defined it. If you
have defined multiple administrative principals, or if your primary
authentication server is on another system, you must first enter the
name of an administrative principal who has root privilege (UID 0). You
need to be authenticated as this administrator so that this program
can create the principals and service key files for the authenticated
services that run on the SP system.

For more information, see the k4init man page.

********************************************************************************

setup_authent: Enter name of admin user: root

Kerberos Initialization for "root.admin"

Password: rootPassword

# k4list
Ticket file:    /tmp/tkt0
Principal:  root.admin@XYZ.COM

Issued           Expires          Principal
Nov 11 16:26:11  Dec 12 16:26:11  krbtgt.XYZ.COM@XYZ.COM
#

# <end_of_example>
```

Step 13: Copy Kerberos V4 keys to the backup control workstation

When control workstation services move back and forth between the two control workstations, the Kerberos V4 service keys must remain the same. The **krb_srvtab** file should be the same on both the primary and secondary authentication servers.

Enter the following commands on the backup control workstation:

```
/usr/lpp/ssp/rcmd/bin/rcp -p <primary_name>:/etc/krb-srvtab \
    /etc/krb-srvtab.primary
cp -p /etc/krb-srvtab /etc/krb-srvtab.backup
cat /etc/krb-srvtab.primary >>/etc/krb-srvtab
```

Repeat this procedure whenever you change Kerberos V4 service keys on either of the two control workstations.

Step 14: Verify Kerberos V4 data

Make sure a Kerberos V4 principal and rcmd service key exist for the network address that matches the host name of the backup control workstation.

Run the **/usr/lpp/ssp/kerberos/bin/kadmin** command on the backup control workstation as follows:

```
Welcome to the Kerberos Administration Program, version 2
Type "help" if you need it.
```

```
admin: get_entry rcmd.BackupControlWorkstation

Admin password:

Info in Database for rcmd.BackupControlWorkstation:
Max Life: 255   Exp Date: Fri Apr 28 22:59:59 2000

Attribs: 00  key: 0 0
admin:  q
Cleaning up and exiting.

# <end_of_example>
```

To verify the information, run the **/usr/lpp/ssp/kerberos/bin/ksrvutil list** command on the backup control workstation. The system will display information similar to the following:

```
Version    Principal
2     rcmd.k21sha@PPD.POK.IBM.COM
2     hardmon.k21sha@PPD.POK.IBM.COM
2     rcmd.k21shacw@PPD.POK.IBM.COM
2     hardmon.k21shacw@PPD.POK.IBM.COM
1     hardmon.k21cw@PPD.POK.IBM.COM
1     rcmd.k21cw@PPD.POK.IBM.COM
1     rcmd.k21s@PPD.POK.IBM.COM
1     hardmon.k21s@PPD.POK.IBM.COM
1     rcmd.k21cw_bt@PPD.POK.IBM.COM
1     rcmd.k21s_bt@PPD.POK.IBM.COM
#

# <end_of_example>
```

Task D. Install software

Step 15: Install HACMP or HACMP/ES on both control workstations

Follow the instructions in the “Installing HACMP for AIX software” chapter of the *HACMP for AIX: Installation Guide* or the “Installing the HACMP/ES software” chapter of the *HACMP for AIX: Enhanced Scalability Installation and Administration Guide* to install the HACMP software.

Step 16: Verify cluster software

If you are installing HACMP, follow the instructions in the “Verifying cluster software” chapter of the *HACMP for AIX: Installation Guide* to verify the cluster software. If you are installing HACMP/ES, follow the instructions in the “Installing the HACMP/ES software” chapter of the *HACMP for AIX: Enhanced Scalability Installation and Administration Guide.*

Step 17: Install the HACWS image on both control workstations

Install **ssp.hacws** on both control workstations using SMIT.

TYPE **smit install_latest**

- The Install Software window appears.

SELECT The input device (press F4 and select a device) and the **ssp.hacws** file set.

PRESS **Ok** to complete the action.

Step 18: Stop the primary control workstation

Stop control workstation services on the primary control workstation by entering the following command:

```
/usr/sbin/hacws/spcw_apps -d
```

Step 19: Configure serial network

Configure the target mode SCSI connection or the target mode SSA connection, or the raw RS-232 serial line which is to be used as the HACMP serial network. Refer to the “Configuring networks” chapter of the *HACMP for AIX: Installation Guide* for instructions.

Step 20: Configure the network

You need to configure each control workstation to use its boot addresses after it reboots. (You should have identified the boot addresses in “Step 2: Plan network configuration” on page 110.) To do this, enter the following command:

```
smit chinet
```

The Internet Address field should contain the IP address corresponding to the boot address. Do this for each boot address on both control workstations.

Note: Do not reboot until you are instructed to do so in “Step 34: Reboot control workstations” on page 125.

Step 21: Migrate the internal file system

If not previously done, the **/spdata** directory needs to reside in an external volume group so both control workstations can access it. You can accomplish this by migrating the **/spdata** files from a separate file system in an internal volume group to the external volume group.

Complete the following procedure:

1. Determine which file system contains the **/spdata** directory by entering the following command:

   ```
   df /spdata
   ```

   If the file system is separate, the mount point will be **/spdata**.
2. Unmount the /spdata file system by entering the following command:

   ```
   umount /spdata
   ```
3. Enter:

   ```
   smit chjfs
   ```
4. Select **/spdata**.
5. Set the new mount point to **/spdata.old**. Set `mount automatically at system restart` to **no**.
6. Mount the file system at its new mount point by entering the following command:

   ```
   mount /spdata.old
   ```

Step 22: Set up the external file system

The following list illustrates how the HACMP terminology maps to your HACWS setup:

HACMP Term	HACWS Equivalent

Source node	Primary control workstation
Destination node	Backup control workstation

If you have not already set up an external file system, refer to the shared logical volume configuration instructions in the “Defining shared LVM components” chapter of the *HACMP for AIX: Installation Guide*.

Follow the instructions for non-concurrent access. Use the following list with the shared logical volume configuration instructions:

1. Create a volume group on the primary control workstation. The volume group can have any name—it does not have to be **spvg**. (**spvg** is used in the examples.) Be sure to use a major number that is available on both control workstations.
2. Follow the instructions to create a file system in the new volume group.
3. Follow the instructions to rename the logical volumes. After you rename the logical volumes, check the **/etc/filesystems** file to make sure the **dev** and **log** attributes reflect the changes. If they do not, edit **/etc/filesystems** and update these attributes to reflect the changes.
4. Follow the instructions to add copies to the logical volume on the source node (primary control workstation).
5. Mount the new **/spdata** file system. Enter the following command:

   ```
   mount /spdata
   ```
6. Copy the **/spdata** files from the old **/spdata** file system to the new file system. Enter the following command:

   ```
   cp -Rph /spdata.old/* /spdata
   ```
7. Unmount both spdata file systems. Enter the following commands:

   ```
   umount /spdata.old
   umount /spdata
   ```
8. Follow the instructions in the remainder of non-concurrent access section to do the following:
   a. Vary off the volume group on the primary control workstation.
   b. Import the volume group on the backup control workstation. (Make sure to use the same major number you used on the primary control workstation.)
   c. Change the volume group so it remains dormant on the backup control workstation. If the volume group is mirrored three ways, then answer Yes to the `quorum required` question. Otherwise, answer **No** to the question. Refer to the HACMP publications for more information about quorum in an HACMP cluster.
   d. Vary off the volume group on the backup control workstation.
9. Make the **/spdata** file system available on the primary control workstation.
   a. Vary on the volume group on the primary. Enter the following command:

      ```
      varyonvg spvg
      ```
   b. Mount the **/spdata** file system. Enter the following command:

      ```
      mount /spdata
      ```
10. Issue the SMIT **chvg** command to make sure quorum status on the primary control workstation matches the quorum status that you specified for the backup control workstation.

Step 23: Complete administration tasks

Follow the AIX administration instructions in the “Additional AIX administration tasks” chapter of the *HACMP for AIX: Installation Guide*. Use the following list with the installation instructions:

1. The **install_hacws** command you will use in “Step 29: Set up the HACWS configuration” on page 123 sets I/O pacing to the HACMP recommended starting points. If you want to let **install_hacws** set these values, skip the set I/O pacing step.
2. The **/etc/hosts** file should contain both short and long host names for each network interface, so both types of references (for example, **sp_cws_bt** and **sp_cws_bt.xyz.com**) can be resolved.

Task E. Configure High Availability Cluster Multi-Processing

This section describes the steps you take to setup High Availability Cluster Multi-Processing (HACMP) to work with HACWS.

Step 24: Define the cluster environment

Follow the instructions in the “Defining the cluster topology” chapter of the *HACMP for AIX: Installation Guide* to define the cluster environment. You must configure HACMP to use hardware address swapping in conjunction with IP address takeover on any service adapter belonging to the SP Ethernet network or whose name matches the host name of the primary control workstation.

The service adapter *dutchess.xyz.com* from the sample scenario, (see the ″Planning for a high availability control workstation″ chapter in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*) falls into both categories. The “Planning TCP/IP networks” chapter of the *HACMP for AIX: Planning Guide* explains hardware address swapping.

As you define your cluster environment, you may find it helpful to read the following examples which reference the example configuration discussed in the ″Planning for a high availability control workstation″ chapter in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*.

Step 24.1: Define the cluster ID and name

Follow the instructions in the *HACMP for AIX: Installation Guide* or the *HACMP for AIX: Enhanced Scalability Installation and Administration Guide*.

Step 24.2: Define nodes to HACMP

In the sample scenario discussed in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*, the host names of the control workstations are *dutchess.xyz.com* and *ulster.xyz.com*. HACMP does not care whether you use the short or long host names, so use the short host names. The sample scenario short hostnames are `dutchess` and `ulster`. Follow the instructions in the *HACMP for AIX: Installation Guide* or the *HACMP for AIX: Enhanced Scalability Installation and Administration Guide* to define the cluster nodes to HACMP.

Step 24.3: Define adapters to HACMP

In the sample scenario discussed in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*, the service address is *dutchess.xyz.com*. The node *dutchess* has the boot address *dutchess_bt.xyz.com* and the node *ulster* has the boot address *ulster_bt.xyz.com*. The screens in the following figures show how the sample scenario adapters would be defined to HACMP. Refer to these

figures as you follow the instructions in the *HACMP for AIX: Installation Guide* to define the adapters to HACMP.

```
Add an Adapter

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

Entry Fields
* Adapter IP Label                              dutchess_bt
* Network Type                                  ether               +
* Network Name                                  hacws_en0           +
* Network Attribute                             public              +
* Adapter Function                              boot                +
Adapter Identifier                             129.40.60.21
Adapter Hardware Address
Node Name                                      dutchess            +
```

Figure 1. Adding an Adapter—Primary Boot Address

```
Add an Adapter

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

Entry Fields
* Adapter IP Label                              ulster_bt
* Network Type                                  ether               +
* Network Name                                  hacws_en0           +
* Network Attribute                             public              +
* Adapter Function                              boot                +
Adapter Identifier                             129.40.60.22
Adapter Hardware Address
Node Name                                      ulster              +
```

Figure 2. Adding an Adapter—Backup Boot Address

```
Add an Adapter

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

Entry Fields
* Adapter IP Label                              dutchess
* Network Type                                  ether               +
* Network Name                                  hacws_en0           +
* Network Attribute                             public              +
* Adapter Function                              service             +
Adapter Identifier                             129.40.60.99
Adapter Hardware Address                       0x02608c2d2a10
Node Name                                                          +
```

Figure 3. Adding an Adapter—Service Address. Note the Node Name field is blank, and the adapter hardware address field has been completed.

Step 24.4: Configure network modules

Follow the instructions in the *HACMP for AIX: Installation Guide* or the *HACMP for AIX: Enhanced Scalability Installation and Administration Guide*.

Step 24.5: Synchronize the cluster definition on all nodes

Follow the instructions in the *HACMP for AIX: Installation Guide* or the *HACMP for AIX: Enhanced Scalability Installation and Administration Guide*.

Step 25: Configure the HACWS application server

Follow the instructions in the "Configuring cluster resources" chapter of the *HACMP for AIX: Installation Guide* or the "Configuring an HACMP/ES cluster" chapter of the *HACMP for AIX: Enhanced Scalability Installation and Administration Guide* to configure the HACWS application server. Use the following definitions:

Server Name (Specify a unique name)

Start Script **/usr/sbin/hacws/spcw_apps -ua**

Stop Script **/usr/sbin/hacws/spcw_apps -di**

Step 26: Define the resource group

Follow the instructions in the "Configuring cluster resources" chapter of the *HACMP for AIX: Installation Guide* to define the resource group. If you are installing HACMP/ES, follow the instructions in the "Configuring an HACMP/ES cluster" chapter of the *HACMP for AIX: Enhanced Scalability Installation and Administration Guide*.

Adding a resource group

Use the following options:

ENTER **hacws_group1** for Resource Group Name

SELECT **rotating** for Node Relationship

ENTER The node names of the primary control workstation and backup control workstation for Participating Node Names

PRESS **Ok** to complete option selection.

Configuring a resource group

Use the following options:

ENTER The service addresses identified to HACMP for Service IP Label. You identified these service addresses to HACMP in "Step 24: Define the cluster environment" on page 120 (`dutchess` in Figure 3 on page 121).

This field can contain multiple network interfaces. You may have to type them out in order to get the SMIT screen to accept them all. You must have at least one network interface name matching the host name of the primary control workstation.

ENTER The name of your external file system for Filesystems (**/spdata** or **/spdata/sys1**).

ENTER The name of the volume group containing the **/spdata** or **/spdata/sys1** file system for Volume Groups.

ENTER The server name you used in "Step 25: Configure the HACWS application server" for Application Servers.

PRESS **Ok** to complete option selection.

Complete the rest of the instructions in the "Configuring cluster resources" chapter of the *HACMP for AIX: Installation Guide*. If you are installing HACMP/ES, finish the

instructions in the “Configuring an HACMP/ES cluster” chapter of the *HACMP for AIX: Enhanced Scalability Installation and Administration Guide*.

Step 27: Verify the cluster and node environment

Follow the instructions in the “Verifying the cluster topology” chapter of the *HACMP for AIX: Installation Guide* to verify the cluster and node environment.

Task F. Set up and test the HACWS

This section describes the steps you take to customize HACWS and verify its function.

Step 28: Make each control workstation addressable by its host name

“Step 29: Set up the HACWS configuration” requires each control workstation be addressable by its host name. If the name of the primary control workstation is *dutchess.xyz.com*, the backup control workstation must be able to communicate with the primary control workstation using the name *dutchess.xyz.com*. If the name of the backup control workstation is *ulster.xyz.com*, the primary control workstation must be able to communicate with the backup control workstation using the name *ulster.xyz.com*.

If either control workstation is not addressable by its host name, you need to enter the appropriate **ifconfig** command to cause the required name to be temporarily configured as a network interface.

Note: Do not use SMIT for this step. Using SMIT causes the change to be permanent.

An example of an **ifconfig** command is as follows:

```
ifconfig en0 dutchess.xyz.com netmask 255.255.255.192 up
```

The **ifconfig** options you need to use are specific to your site. Refer to the **ifconfig** man page for more information.

Step 29: Set up the HACWS configuration

Configure the primary and backup control workstation as an HACWS configuration.

Note: Do this step from the primary control workstation.

If using:	**Do this:.**	
SMIT	**TYPE**	**smit hacws** • The High Availability Control Workstation Management window appears.
	SELECT	**Install and Configure HACWS**.
	ENTER	The node names of the primary control workstation and backup control workstation in the HOSTNAME fields.
	SELECT	**yes** for Execute on both primary and backup?.
	PRESS	**Ok** to complete option selection and install HACWS

install_hacws	Enter:
	`/usr/sbin/hacws/install_hacws -p \`
	`primary_hostname -b backup_hostname -s`	

Step 30: Customize cluster event processing

Identify the HACMP pre- and post-event scripts provided by HACWS to HACMP.

If using:	Do this:.	
SMIT	**TYPE**	**smit hacws**
		• The High Availability Control Workstation Management window appears.
	SELECT	**Identify Event Scripts to HACMP**.
	PRESS	Enter to continue.
spcw_addevents	Enter: `/usr/sbin/hacws/spcw_addevents`	

If you want to further customize cluster event processing, follow the instructions in the “Managing a high availability control workstation” chapter in *PSSP: Administration Guide*.

Step 31: Add IP address aliases

In “Step 2: Plan network configuration” on page 110 you planned the network configuration for both control workstations. If your configuration requires IP address aliases, you need to provide the appropriate commands to HACWS. If these commands already existed somewhere else before you installed HACWS, remove them from the old location.

If you need to use an IP address alias to configure the host name of the backup control workstation as a network interface, edit the **/etc/rc.backup_cw_alias** script on the backup control workstation. This script only runs on the backup control workstation. See the comments in the script for more details. If you make any changes to the **/etc/rc.backup_cw_alias** script, you must also edit the **/etc/rc.net** script so **/etc/rc.backup_cw_alias** runs when the backup control workstation boots. Use an editor to add the following line near the end of the **/etc/rc.net** file:

```
/etc/rc.backup_cw_alias
```

If you need to use an IP address alias to configure a network interface for an SP system partition (SP Switch only), or if you need to configure an IP address alias on the active control workstation for some other reason, edit the **/spdata/sys1/hacws/rc.syspar_aliases** script. This script runs on the active control workstation before it starts the control workstation services. See the comments in the script for more details.

Step 32: Verify the HACWS configuration

Verify the HACWS configuration.

If using:	Do this:.	
SMIT	**TYPE**	**smit hacws**
		• The High Availability Control Workstation Management window appears.
	SELECT	**Verify HACWS Installation and Configuration**.

If using:	Do this:.
hacws_verify	Enter: `/usr/sbin/hacws/hacws_verify`

Step 33: Verify the hardware connections

Verify the connections between the backup control workstation and the frames. Run this step from the backup control workstation.

If using:	Do this:.
SMIT	**TYPE** **smit hacws** • The High Availability Control Workstation Management window appears. **SELECT** **Verify Frame to Control Workstation Cabling**.
spcw_verify_cabling	Type: `/usr/sbin/hacws/spcw_verify_cabling`

Step 34: Reboot control workstations

Reboot the primary control workstation. After it finishes booting, reboot the backup control workstation.

Step 35: Start cluster services on the primary control workstation

Follow the instructions in the "Starting and stopping cluster services" chapter of the *HACMP for AIX: Administration Guide* to start cluster services on the primary control workstation.

Using SMIT:

TYPE **smit clstart**

• The Start Cluster Services menu is displayed.

SELECT **now** for Start now option.

SELECT **true** for Startup Cluster Information Daemon? option.

PRESS **Ok** to complete option selection.

Step 36: Verify the HACWS installation

To make sure you installed the HACWS correctly, verify that the control workstation services exist and can move between the primary and backup control workstations.

Verify the control workstation services

Make sure the control workstation services come up on the primary control workstation. Use the following procedure to test them:

1. Service addresses should be configured on the primary control workstation. Try to telnet to a service address.
2. **/spdata** should be mounted on the primary control workstation. To verify this, use the following command:

   ```
   df /spdata
   ```

3. The SDR should be available. To verify this, use the following command:

   ```
   /usr/lpp/ssp/bin/SDRGetObjects SP
   ```

The test is successful if you receive output. You can tell whether the startup of the control workstation services has completed by issuing the command:

```
grep "SPCW_APPS COMPLETE" /tmp/hacmp.out
```

Start cluster services on the backup control workstation

Once control workstation services have completely started on the primary control workstation, start cluster services on the backup control workstation by following the procedure described in “Step 35: Start cluster services on the primary control workstation” on page 125.

Cause a failover

After cluster services have started on the primary control workstation, direct HACMP to move control workstation services to the backup control workstation.

Using SMIT:

TYPE	**smit clstop**
	• The Stop Cluster Services menu is displayed.
SELECT	**now** for Stop now option.
SELECT	**takeover** for Shutdown mode option.
PRESS	**Ok** to complete option selection.

This should stop cluster services on the primary control workstation and move control workstation services to the backup control workstation. Make sure control workstation services move to the backup. On the backup control workstation, follow the procedure described in “Verify the control workstation services” on page 125.

When you are finished testing, restart cluster services on the primary control workstation by following “Step 35: Start cluster services on the primary control workstation” on page 125.

Note: The first time you start control workstation services on the backup control workstation, the **spcw_apps** command runs and takes about a half hour to complete. You must wait for it to complete before you move control workstation services back to the primary control workstation. You can tell whether the command completed by issuing the following command:

```
grep "SPCW_APPS COMPLETE" /tmp/hacmp.out
```

Once the startup of the control workstation services has completed on the backup control workstation, you can restart cluster services on the primary control workstation and then move control workstation services back to the primary control workstation. When you have completed testing, make sure to restart cluster services on both control workstations.

Chapter 4. Migrating the software on your RS/6000 SP system

This chapter describes how you can migrate your system to PSSP 3.4 and your target level of AIX. At the time this manual was published, PSSP 3.4 was supported on AIX 4.3.3 and AIX 5L 5.1. Depending on the present level of your system, there may be a recommended minimum PTF level. Refer to the *READ THIS FIRST* document for information on required AIX file sets, PTFs, and for any additional AIX levels that PSSP 3.4 may support.

Migrating your control workstation and nodes is a complex task. Thorough planning should be completed prior to attempting to migrate your SP to a new level of AIX and PSSP. The best technique for simplifying the complexity of the task is to divide the migration of your SP system into smaller steps, where you verify that each step was successfully completed before proceeding to the next step. This technique of simplifying the complexity of the task helps to ensure the successful completion of your migration, or if necessary, allows you to recover from unexpected problems should they arise.

High-level migration steps

The process of migrating your system includes the following high-level steps. Read the high-level steps first, then refer to the appropriate section in this chapter to perform the actual steps. Once you perform the steps, refer back to this section to determine which step to address next.

Preparing to migrate

Preparing to migrate your system is a very important activity that you should not skip. It involves planning the migration carefully by examining your current configuration and the desired future configuration. Reading the migration information in the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* is a prerequisite.

Refer to “Preparing to migrate” on page 129.

Applying PTFs to nodes

Before migrating your control workstation, you may need to apply specific PTFs to your nodes. Refer to the *READ THIS FIRST* document for a list of required PTFs.

Migrating the control workstation

Prior to migrating any of your nodes, you must migrate your control workstation to the latest level of AIX and PSSP of any node you wish to serve. For example, if you plan to migrate any node to AIX 4.3.3 and PSSP 3.4, the control workstation must first be migrated to AIX 4.3.3 or later AIX level and PSSP 3.4.

To perform your control workstation migration, refer to “Migrating the control workstation to PSSP 3.4” on page 133.

PSSP 3.4 is supported on both AIX 4.3.3 and AIX 5L 5.1. If you plan to migrate to PSSP 3.4 on AIX 5L 5.1, you must first migrate your control workstation to PSSP 3.4 on AIX 4.3.3. After the control workstation migration to PSSP 3.4 on AIX 4.3.3 is complete, you can then migrate to AIX 5L 5.1.

If you are using a version of General Parallel File System (GPFS) prior to GPFS 1.3, do not enable DCE for SP Trusted Services until all instances of GPFS are upgraded to GPFS 1.3 or later. Refer to *GPFS: Concepts, Planning, and Installation Guide* for additional information.

Partitioning your system (if necessary)

Note: You cannot partition systems that have an SP Switch2 installed or systems of clustered enterprise servers.

When you read the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* and the section "Preparing to migrate" on page 129, one of the issues you needed to consider is whether you will need to partition your system. If you determined that you wanted to partition your system due to coexistence limitations, migration test purposes, or for any other reason, you should partition your system at this time. Refer to the "Managing system partitions" chapter of the *PSSP: Administration Guide* for instructions on how to partition your system. Make sure your system is working correctly before continuing. Also refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more information.

Migrating a test node to PSSP 3.4

If possible, you should migrate a single node to your target AIX level and PSSP 3.4 before migrating all or a group of your nodes. This will allow you to gain experience with the migration, resolve any problem that might occur due to the migration, and help you determine how much time you will need to migrate all or a group of your nodes. You should use the control workstation as the boot/install server for this test node.

Refer to "Paths to migrate the nodes to PSSP 3.4" on page 152 for more information.

Migrating the boot/install servers to PSSP 3.4

Prior to migrating any of your nodes, you must migrate the boot/install servers to the latest level of AIX and PSSP on any node you wish to serve. For example, if you want to migrate some nodes to AIX 4.3.3 or later AIX level and PSSP 3.4, you must first migrate these node's boot/install servers to AIX 4.3.3 and PSSP 3.4.

Refer to "Paths to migrate the nodes to PSSP 3.4" on page 152 for more information.

Migrating the nodes to PSSP 3.4

Migrating the nodes is the final activity in the migration process.

Refer to "Paths to migrate the nodes to PSSP 3.4" on page 152 for more information.

Performing post-migration activity

This section discusses the activities you should perform after a successful migration and the recovery procedures you should perform if the migration process was not successful. It also discusses how to deinstall an AIX release or AIX PTF after a migration or upgrade.

Refer to "Post-migration activity" on page 175 for more information.

Preparing to migrate

Before performing the following steps, you should:

- Read the migration information in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*
- Understand supported migration paths to PSSP 3.4
- Understand coexistence issues with PSSP 3.4
- Understand system partitioning issues

Step 1: Verify control workstation requirements

Ensure that the following requirements are met for the control workstation:

- Refer to the *READ THIS FIRST* document for the latest information on supported AIX releases and PTF levels.
- Install adequate serial ports and RS-232 cables to support all frames.
- Allocate adequate disk space for rootvg, paging, and installation. IBM requires a minimum of 4 GB of DASD made available on the control workstation. IBM suggests allocating 2 GB for rootvg, and 2 GB for the **/spdata** file system. If you plan to support mixed levels of PSSP and AIX, you may require more than 2 GB of **/spdata** file system.
- Ensure that all host names and IP addresses are resolvable on the control workstation. You should not change them for the duration of the migration.
- You must run all migration tasks from root on the control workstation. IBM suggests that you add the following directory paths in your .profile to avoid incomplete path names:

  ```
  /usr/lpp/ssp/bin
  /usr/sbin
  /usr/bin
  /usr/lpp/ssp/kerberos/bin
  ```

  The following is an example of root user path in the .profile:

  ```
  PATH=/usr/lpp/ssp/bin:/usr/sbin:/usr/lpp/ssp/kerberos/bin
  ```

Step 2: Verify boot/install server requirements

Ensure that the following requirements are met for the boot/install servers:

- The boot install server node must be at the highest level of AIX and PSSP to which it will serve.
- Allocate adequate disk space for rootvg on the boot/install server nodes taking into consideration the size of your mksysb images, the pssp.installp images, and SPOT resources.

Step 3: Understand system reconfiguration issues

Note that you should not reconfigure your system in any way during the migration process. For example, do not add a frame, switch, node, or any other piece of hardware to your system until after the migration process has completed. In addition, you should not add or delete any host names or IP addresses until after the migration process has completed. This will help to ensure a successful migration. For more information on reconfiguring your system, refer to “Chapter 6. Reconfiguring the RS/6000 SP system” on page 195.

Step 4 Understand workload management issues

As of PSSP 3.1, job management functions previously provided by the PSSP Resource Manager have been added to the LoadLeveler product. The switch table management for user space parallel jobs previously provided by the PSSP Resource Manager has been moved to a new PSSP file set **ssp.st**, Job Switch Resource Table Services. Depending on how you currently use the Resource Manager, see the *PSSP: Administration Guide* for more information on how maintain the same functionality.

The Resource Manager daemons have been removed from PSSP 3.1. The **ssp.jm** and **ssp.clients** file sets still contain the commands and library necessary to support back-level system partitions from the control workstation. The **ssp.jm** file set is no longer automatically installed by **pssp_script**. If you want to support back-level system partitions, you need to install the **ssp.jm** file set on the control workstation. Previous releases of the Resource Manager will not be automatically removed from the nodes being migrated. To regain file system space and prevent incompatibilities, it is recommended that the **ssp.jm** file set be deinstalled from all nodes.

The PSSP Job Switch Resource Table Services provide a way to load, unload, clean, and query Job Switch Resource Tables. The **ssp.st** file set is installed by **pssp_script**. LoadLeveler uses these services when scheduling and starting user space jobs.

Step 5: Understand system security issues

When migrating your system, IBM strongly suggests that you should continue using whatever authentication method you presently have. After your migration is complete, you can then change your authentication method. Refer to “Chapter 5. Reconfiguring security” on page 179 for information on changing your security settings. If you have PSSP 3.1 installed with DCE, see the exception recommended in “Step 5.1: Migrating a PSSP 3.1 system configured with DCE”.

Step 5.1: Migrating a PSSP 3.1 system configured with DCE

PSSP 3.1 supported a limited implementation of DCE at DCE 2.2 or later. If your system was set up to use this implementation, your security attributes would look similar to the following:

```
splstdata -p
List System Partition Information
...
auth_install    k4
auth_root_rcmd  k4
auth_methods    k5:k4:std
```

If your system is configured as shown in the preceding example, there are additional node migration considerations that you must take into account. PSSP supports migration of a PSSP 3.1.1 DCE implementation to PSSP 3.4 using a BOS Node Upgrade. (See “BOS node upgrade” on page 153.) However, once the migration is complete, you must take additional steps before you can perform a mksysb install of a PSSP 3.4 node. For example, if you are installing a new node or reinstalling a migrated node.

PSSP 3.2 or later requires that if **auth_methods** includes **k5**, DCE must be installed on the node before **psspfb_script** sets the authentication methods during the install process. If you choose to migrate using a “mksysb install of nodes” on page 160, this same rule applies. This requirement also existed for PSSP 3.1, so you may have already implemented a process to do mksysb installs.

If you want to maintain your PSSP 3.1 **k5** security settings with PSSP 3.4, you will need to add code to your **/tftpboot/script.cust** file to do a mksysb install of a node. This file is documented in “Step 61: Perform additional node customization” on page 90 and in “Appendix E. User-supplied node customization scripts” on page 297. The new code in **script.cust** will need to mount the directory containing DCE and install your required DCE clients.

Alternatively, you can use PSSP to install DCE during the mksysb install of the node. This is done by issuing the **spsetauth -i** command to change the **auth_install** security attribute to include DCE. You will need to ensure that DCE is at 3.1 or later on the control workstation and that your AIX lppsource contains the DCE file sets. See “mksysb install of nodes” on page 160 for detailed instructions.

Step 5.2: Migration of nodes with secure remote command methods enabled

PSSP 3.4 provides the ability to remove the dependency PSSP has on the **rsh** and **rcp** commands issued as root by enabling the use of a secure remote command method. It is the administrator’s responsibility to choose the secure remote command software and install it on the control workstation.

All nodes must be at PSSP 3.2 or later and RRA must be enabled before the secure remote command method can be enabled.

Refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for additional information.

Boot/install server nodes, other than the control workstation, require **rsh** and **rcp** capabilities to use NIM services. It is up to the administrator to add the authorization methods necessary for boot/install activities.

If additional files must be copied to the nodes during the installation process with secure remote command and restricted root remote commands enabled, the **firstboot.cmds** sample file gives examples of how to enable the copy from the control workstation to the nodes in the restricted access environment and secure remote command enabled environment.

Step 6: Understand LoadLeveler Issues

The LoadLeveler Central Manager node must be migrated to PSSP 3.4 before migrating any other LoadLeveler nodes. Refer to *Installation Memo for IBM LoadLeveler*, GI10-0642, and *IBM LoadLeveler for AIX 5L: Using and Administering* for additional information on migration considerations.

Step 7: Reserve port numbers

Some of the subsystems managed by the Syspar Controller (**syspar_ctrl**) allocate port numbers from the range 10000 to 10100, inclusive. Therefore, if any customer subsystem, such as DB2, uses port numbers in this range, such port numbers must be reserved in **/etc/services**. For details, refer to “Appendix G. Reserving ports” on page 303.

Step 8: Understand runtime prerequisite issues

PSSP 3.4 has prerequisites for several runtime files from the VisualAge C++ product. If you have this product installed, you may need to update it to the current service levels either before or during your migration. See “Step 13: Copy the PSSP images for PSSP 3.4” on page 140 for detailed information.

Step 9: Archive and verify the SDR

You should archive the SDR at its current level by issuing the **SDRArchive** command:

```
SDRArchive append_string
```

where *append_string* is the name you want to use for this archive. If you specified the *append_string*, it is appended to the name of the backup file **/spdata/sys1/sdr/archives/backup**.*JulianDate*.*HHMM*.*append_string*.

As of PSSP 3.2, the SDR is NLS enabled and allows non-ASCII data to reside in the SDR if that is how you have your site environment set up. See “Step 30: Enter site environment information” on page 49 for more information on setting up your site environment. However, if you are migrating your control workstation from PSSP 2.4 or PSSP 3.1 to PSSP 3.4, you must ensure that only ASCII data resides in the SDR and that the locale on your machine is set to **en_US** or **En_US**, otherwise the migration will fail. Only the **en_US** and **En_US** locales are supported for PSSP 2.4 and PSSP 3.1. You must remove all non-ASCII data from the SDR and set your locale to **en_US** or **En_US** before migrating from PSSP 2.4 or PSSP 3.1 to PSSP 3.4.

PSSP provides an **SDRScan** utility to determine whether non-ASCII data is present in your SDR. **SDRScan** was shipped in the **ssp.basic** file set with PSSP 2.4 PTF Set 27 (APAR IY23244) and with PSSP 3.1.1 PTF Set 21 (APAR IY23701). If this utility is installed on your system at **/usr/lpp/ssp/bin/SDRScan**, you can run the command to determine if you have non-ASCII data in your SDR:

```
SDRScan
```

All SDR objects containing non-ASCII data will be written to standard output.

If the **SDRScan** utility is not available on your system and you think you may have non-ASCII data in your SDR, you will need to manually determine the presence of this data. You can do this by running the **SDRGetObjects** command for the classes you suspect may have non-ASCII data and visually inspect the results:

```
SDRGetObjects class_name
```

All non-ASCII data must be removed from the SDR. The method to do this will depend on the attributes and classes that contain the non-ASCII data. You will need to use the appropriate PSSP commands to modify the data.

When migrating from PSSP 2.4 or PSSP 3.1 to PSSP 3.4, you must also have the default locale on your control workstation set to **en_US** or **En_US**. You can verify this by viewing the LANG stanza in the **/etc/environment** file and ensuring it is set to one of these values.

Step 10: Back up your control workstation

Before you migrate your system, you should always make a backup of your existing system. The following list provides you with a basic set of instructions for backing up your system:

1. Create a backup image of your existing root volume group of the control workstation. For example:

   ```
   mksysb -i /dev/rmt0
   ```

2. Back up files and file systems with critical data

You should keep backup copies of critical files and file systems that have configuration data in them. For example, you should back up the **/spdata** file system. Use the **tar** command to save critical files and use the **backup** command to save file systems.

Command	Example
tar	`/bin/tar -cvf /dev/rmt0 files`
backup	`/usr/sbin/backup -0cuf /dev/rmt0 /spdata`

For AIX 4.1 and later systems, you may choose to back up the volume group by issuing the **savevg** command. For example, to save the spdatavg volume group, you could issue the following command:

```
savevg -f /dev/rmt0 -i spdatavg
```

3. Export Nonroot Volume Groups

   You may want to export your nonroot volume groups immediately prior to the AIX migration. This will undefine the volume group from the system during the AIX migration. For example, you may want to export the spdatavg volume group by issuing the following commands:

   ```
   umount      /spdata   (unmount the file system)
   varyoffvg   spdatavg  (vary off volume group)
   exportvg    spdatavg  (export volume group)
   ```

   Once the AIX migration is complete, you will need to import any volume group that you previously exported. For example, to import the spdatavg volume group, issue the following command:

   ```
   importvg -y spdatavg -V major # hdisk
   ```

4. Verify Your Backups

   Whether you made the backup to a file or a tape, verify that the file exists on the medium.

Step 11: Back up your nodes

After backing up the control workstation in the last step, you should make a backup of your nodes. The following list provides you with a basic set of instructions for backing up your nodes.

1. Create a backup image of your node. You must save the mksysb file to a tape, media, or to a file system on another workstation. For example:

   ```
   mksysb -i /dev/rmt0
   ```

2. Verify your backups. Whether you make the backup to a file or a tape, verify that the file exists on the medium.

Migrating the control workstation to PSSP 3.4

Follow the steps in this section if you are upgrading your control workstation. Supported starting points for the control workstation migration include:

- PSSP 3.2 and AIX 4.3.3
- PSSP 3.1.1 and AIX 4.3.3
- PSSP 2.4 and AIX 4.2.1 or 4.3.3

Depending on the current level of AIX and PSSP installed on the control workstation, you may first need to migrate to a new level of AIX.

If you have an HACWS configuration, proceed to “HACWS migration strategy” on page 165 .

To migrate to a new PSSP level without changing your AIX level, perform “Step 2: Quiesce your system” and “Step 13: Copy the PSSP images for PSSP 3.4” on page 140 through “Step 31: Validate the control workstation” on page 146.

Keep in mind that you should always make a mksysb backup of your control workstation and a backup of the **/spdata** directory before proceeding with the migration process. For information on making a mksysb, refer to “Step 10: Back up your control workstation” on page 132.

Restricted Root Access

As of PSSP 3.2, you have the option of running your SP system with an enhanced level of security. With the restricted root access (RRA) option enabled, PSSP does not internally issue **rsh** and **rcp** commands as a root user from a node. Also, PSSP does not automatically grant authorization for a root user to issue **rsh** and **rcp** commands from a node. If you enable this option, some procedures might not work as documented. For example, to run HACMP, an administrator must grant the authorizations for a root user to issue **rsh** and **rcp** commands that PSSP otherwise grants automatically. See the ″Planning for security″ chapter in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for a description of this function and a complete list of limitations.

Step 1: Prepare to migrate and verify requirements

Before beginning, refer to “Preparing to migrate” on page 129 for information on preparing to migrate and verifying control workstation and system requirements.

Step 2: Quiesce your system

Be sure to quiesce your system as follows:

- Ensure all users are logged off nodes
- Stop all user jobs running over the switch
- Ensure that any jobs that start up automatically (for example batch submission queues), should be shut down or quiesced.
- If you are using the Switch Admin daemon for node recovery, stop it by issuing **stopsrc -s swtadmd** on SP Switch systems or **stopsrc -s swtadmd2** on SP Switch2 systems.
- Quiesce the switch using the **Equiesce** command

Step 3: Migrate to AIX 4.3.3

You need to upgrade the control workstation to the target AIX level (AIX 4.3.3). Determine which of the following methods you want to follow based upon the information described at the beginning of this chapter:

Note: For Kerberos V4 systems, you must install the **bos.net.uucp** file set. This file set contains programs required for Kerberos V4 secure transfer of keyfiles to the nodes.

- **BOS Upgrade**

  Your control workstation may already be at AIX 4.3.3. You should issue the **installp** commands to install the new AIX 4.3.3 PTFs or file sets on top of the

current AIX level to preserve the current files in the / (root), /var, /usr, and user-based file systems. To perform this operation, do “Step 3a: Upgrade”.

- **BOS Migration Install**

  The control workstation is currently at an earlier version or release of your target AIX level. For example, your control workstation is at AIX 4.2.1 and your target AIX level is 4.3.3. Use this method to preserve the current files in the / (root), /var, /usr, and user-based file systems. To perform this operation, do “Step 3b: Perform control workstation BOS migration install”.

Step 3a: Upgrade

Applying PTFs for the control workstation allows the current rootvg file systems to be preserved. This activity provides installp updates and installs necessary AIX LPs to the control workstation. You must have all the necessary AIX file sets and PTFs listed in the *READ THIS FIRST* document available during this PTF upgrade. You will issue the **installp** command where the input source can be a CD-ROM or a directory (lppsource) that contains the PTFs or file sets. Using the directory requires you to use the **bffcreate** command to copy the AIX 4.3.3 PTFs or file sets into the lppsource directory. Beware that the AIX update CD may contain PTFs for several levels of AIX. You need only to copy the PTFs pertaining to the levels of AIX installed on your system.

To create a list of LPs on the control workstation, issue:

```
lslpp -l -J >/tmp/FILE.43
```

To install and apply PTF service from a CD-ROM and commit LPs, issue:

```
installp -acNgd /dev/cd0 -f /tmp/FILE.43
```

Notes:

1. If you use this method for an upgrade, you will see messages for some file sets that were not upgraded. For example, any of the PSSP file sets.
2. Before applying any service to the control workstation, the **supper** daemon should be stopped to prevent any files from being propagated before the updates have been completed. To stop the **supper** daemon, issue the following command on the control workstation:

   ```
   stopsrc -s supfilesrv
   ```

   Once service is applied, issue the following command on the control workstation to restart the **supper** daemon:

   ```
   startsrc -s supfilesrv
   ```

After completing this step, proceed to “Step 4: Verify AIX levels” on page 136.

Step 3b: Perform control workstation BOS migration install

Attention

You cannot migrate to AIX 5L 5.1 at this point. You must complete the migration to PSSP 3.4 before migrating to AIX 5L 5.1.

This method preserves all file systems except **/tmp**, as well as the root volume group, logical volumes, and system configuration files. This activity provides **installp** updates and installs necessary AIX LPs to the control workstation. You

must have all the necessary AIX file sets and PTFs listed in the *READ THIS FIRST* document available by CD-ROM or from the NIM master.

A migration install from CD-ROM is an interactive process. You will be prompted to verify settings, continue the installation, and verify the migration.

Refer to the *AIX Installation Guide* in chapters 1-3 for more information. The following list provides some hints:

- Boot off your AIX CD-ROM or set up the control workstation as a NIM client from an appropriate NIM master.
- Select option 2 ″Change/Show Installation Settings and Install.″ Always verify that it has been set to *migrate* and that the correct target disk is assigned for root volume group.
- Select option 1 ″System Settings″ on the Installation and Settings menu.
- Select option 3 ″Migration Install″ on the Change Method of Installation menu.
- Select the proper hdisks being used for rootvg. You may want to update the hdisk configuration for the control workstation.
- Customize your control workstation based on the documented PSSP requirements. This includes installing the required AIX LPs and PTFs.

After completing this step, proceed to “Step 4: Verify AIX levels”.

Step 4: Verify AIX levels

Verify that the control workstation was successfully migrated to your target AIX level (AIX 4.3.3), by issuing the the following command:

```
oslevel
```

For example, if your target AIX level is 4.3.3 and the output of this command does not indicate AIX 4.3.3, issue the following command to return a list of AIX files not migrated to AIX 4.3.3. You may need to install AIX PTFs to migrate those file sets to AIX 4.3.3.

```
oslevel -l 4.3.3.0
```

Step 4a: Reboot the control workstation

Before rebooting the control workstation, become familiar with “Tips for installing DCE on the SP” in “Step 22: Configure DCE for the control workstation (required for DCE)” on page 44. If you plan to configure DCE after completing your migration, you may want to establish these environment variables now before rebooting.

If you just upgraded your control workstation (an AIX modification level change, for example, AIX 4.3.2 to 4.3.3,) you now need to reboot the control workstation so changes to the kernel will take effect.

Step 5: Verify the authentication value for AIX remote commands

An authentication method must be set in order for remote commands to work properly. If your SP system contains nodes at PSSP 3.1.1 or earlier, the Kerberos V4 value must be set. For SP systems containing nodes at PSSP 3.2 or later, any value is valid (Kerberos V5, Kerberos V4, or Standard AIX).

Issue the **lsauthent** command to determine if an authentication method is set. If the value is not set, use the **chauthent** command to enable Kerberos V4 or Standard AIX. Issue:

```
chauthent -k4 -std
```

Step 6: Verify the control workstation configuration

Verify the configuration for each Ethernet adapter in the control workstation.

You can verify each Ethernet adapter by *pinging* its IP address and seeing if you get a response. For example:

```
ping -c 1 129.33.34.1
```

Reconfigure the adapter if you did not receive a response.

Changing user process limits

When you first install your system, the number of processes the root user can have is set to an AIX default. You cannot continue installing your system with this default value—you must increase the number to 256.

To check the current value, enter the following command:

```
lsattr -l sys0 -E | grep maxuproc
```

To change the value, enter the following command:

```
chdev -l sys0 -a maxuproc='256'
```

Tunables and tunable values

To validate the network tunables on the control workstation, use the **no** command to view and change the network values. To list the values, enter:

```
/usr/sbin/no -a
```

To change the value of tcp_mssdflt, enter:

```
/usr/sbin/no -o tcp_mssdflt=1448
```

The following list provides the PSSP defaults to use:

```
sb_max              163840
ipforwarding             1
tcp_sendspace        65536
tcp_recvspace        65536
udp_sendspace        32768
udp_recvspace        65536
tcp_mssdflt           1448
tcp_pmtu_discover        0
udp_pmtu_discover        0
```

You need to update the **/etc/rc.net** file so that these changes will take affect.

Verifying system partition aliases (SP Switch only)

If your system is partitioned, first verify that the aliases for each of your system partitions are still defined in **/etc/rc.net**. Then, verify that the aliases are still defined by issuing the following command:

```
netstat -in
```

If the aliases are no longer defined due to the AIX migration, they must be redefined before continuing with your PSSP migration. To do this, edit the **/etc/rc.net** file looking for the template provided for changing the **inet0** to an alias. Follow the instructions in the template and edit the file to define alias IP addresses and names. For example:

```
/usr/sbin/ifconfig tr0 alias 129.40.127.101 netmask 255.255.255.0 up \
>>$logfile 2>&1
```

After editing the **/etc/rc.net** file, make sure it has execute permission, then either execute the **rc.net** script or enter the following **ifconfig** command:

```
ifconfig tr0 alias 129.40.127.101 netmask 255.255.255.0 up
```

For complete details on how to create aliases for your system partitions, refer to the section on what you need to do before you define system partitions in the "Managing system partitions" chapter in the *PSSP: Administration Guide*.

Step 7: Review space requirements for NIM boot images

Before creating a PSSP boot/install server, ensure that there is sufficient space in the root (/) file system or create a separate file system for **/tftpboot** to manage the space required for the boot images (approximately 25 MB per lppsource level supported) created by NIM.

Step 8: Import nonroot volume groups

If any nonroot volume groups were exported in "Step 10: Back up your control workstation" on page 132, you will need to import these volume groups now. For example, to import the spdatavg volume group, issue the following command:

```
importvg -y spdatavg -V major # hdisk
```

Step 9: Review space requirements for /spdata

The **/spdata** directory contains mksysb images and installp file sets. IBM suggests you create a separate volume group for the **/spdata** file system. These file sets are large and need much space (up to 2 GB per lppsource level supported). If you have not done so already, use *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* to help you estimate how much space you need to define.

Step 10: Create the required /spdata directories

You need to create the proper PSSP directory structure for PSSP 3.4.

Note: Make sure you mount the new **/spdata** file system before you create the **/spdata** directories.

You must create subdirectories on the **/spdata** file system for storing critical PSSP data. Make sure the directories have the permissions `rwxr-sr-x`. Issue the following commands to create the required directories:

- To create **/spdata/sys1/install/***name***/lppsource**, issue:

  ```
  mkdir -p /spdata/sys1/install/name/lppsource
  ```

  This is the subdirectory for the required AIX 4.3.3 file sets. You can choose any name, but the name must not contain any dots (.). If you have multiple lppsources, you should pick more than one name. By default, this is set up as the string "default". You might want to pick a more meaningful name, such as *aix433* for AIX 4.3.3 level code. If the control workstation will be the NIM server for different levels of AIX, you should create one subdirectory for each level of AIX.

- To create **/spdata/sys1/install/pssplpp/PSSP-3.4**, issue:

  ```
  mkdir /spdata/sys1/install/pssplpp/PSSP-3.4
  ```

  This is the location of installp file sets for the code version PSSP-3.4. If the control workstation will be a NIM server for nodes at other PSSP levels, also create subdirectories under pssplpp with the code versions for those levels (PSSP-2.4, PSSP-3.1.1, PSSP-3.2).

Step 11: Copy the AIX 4.3.3 LP images and other required AIX LPs and PTFs

If you have not done so already, you must now copy the AIX file sets into the **/spdata/sys1/install/***name***/lppsource** directory on your hard disk on the control workstation.

You can download all of the AIX file sets (a very large number) or only the minimal required AIX file sets (approximately 500 MB).

The following is the minimal list of AIX file sets required to perform mksysb installations: The *prefix.** syntax in the list refers to everything that starts with the *prefix*. For example, **devices.*** refers to all of the file sets starting with **devices**.

Minimal list of AIX 4.3.3 file sets:

```
Java.rte.*                  bos.diag.*
X11.apps.*                  bos.html.en_US.topnav.*
X11.base.*                  bos.mp.*
X11.compat.*                bos.net.*
X11.Dt.*                    bos.powermgt.*
X11.fnt.*                   bos.rte.*
X11.loc.*                   bos.sysmgt.*
X11.motif.*                 bos.terminfo.*
X11.vsm.*                   bos.up.*
X11.msg.*                   devices.*
bos                         perfagent
bos.64bit.*                 perl.*
bos.adt.*
```

Additional files you may want to add to your lppsource:

bos.acct.* Required if you plan to use PSSP accounting.

dce.* The DCE file sets are required only if DCE will be configured by PSSP anywhere on the system. You will need the client portion of the DCE file sets because the installation code installs the DCE client code.

Notes:

1. Download the AIX file sets and the required AIX LPs into **/spdata/sys1/install/***name***/lppsource**. The AIX file sets and required AIX LPs must exist in this directory. Links to file sets in other directories are not allowed. If you change the path name in any way, the installation fails.
2. Refer to your disk usage planning in the “Combining the space requirements” section of *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* to determine if you have allocated enough space to accomplish this task.
3. Allow at least 1-3 hours for moving all the file sets from media to disk.

The perfagent.server file set is part of the Performance Aide for AIX (PAIDE) feature of the Performance Toolbox for AIX (PTX), a separate product. Note that important PTFs for perfagent.server are distributed on the AIX Update CD-ROM. The perfagent.tools file set is part of AIX 4.3.3.

This product provides the capability to monitor your SP system’s performance, collects and displays statistical data for SP hardware and software, and simplifies runtime performance monitoring of a large number of nodes.

The perfagent.server or perfagent.tools file sets must also be copied to all of the lppsource directories on the control workstation. The level of PAIDE copied to each lppsource directory must match the level of AIX in that directory.

The required level of perfagent is dependent upon the level of AIX and PSSP as shown in the following table:

Table 5. perfagent File Sets

AIX Level	PSSP Level	Required File Sets
AIX 4.2.1	PSSP 2.4	perfagent.server 2.2.1.x, where *x* is greater than or equal to 2
AIX 4.3.3	PSSP 2.4	perfagent.server 2.2.33.*
AIX 4.3.3	PSSP 3.1.1	perfagent.tools 2.2.33.*
AIX 4.3.3	PSSP 3.2	perfagent.tools 2.2.33.*
AIX 4.3.3	PSSP 3.4	perfagent.tools 2.2.33.x
AIX 5L 5.1	PSSP 3.4	perfagent.tools 5.1.0.*

Refer to the *READ THIS FIRST* document for the latest information on PAIDE levels.

Login to the control workstation as root and run **bffcreate** using SMIT or the command line. The following example shows the product media on **cd0** and selection of all LPs. Using **all** may load unnecessary file sets into the directory.

```
bffcreate -qvX -t/spdata/sys1/install/name/lppsource -d /dev/cd0 all
```

The following warning message is issued—ignore it:

```
bffcreate:
Warning: important size information is missing from
the table of contents file.  Consequently, there may not
be enough free file system space to successfully create
the bff image(s).  Continuing anyway...
```

Step 12: Install correct level of PAIDE on the control workstation

The Performance Toolbox for AIX, Agent Component (PAIDE) is required. The correct level of AIX PAIDE (perfagent) needs to be installed on the control workstation and copied to all of the lppsource directories. The level of perfagent required is dependent upon the level of AIX. For example, issue:

```
installp -aXd /spdata/sys1/install/name/lppsource perfagent.tools
```

Step 13: Copy the PSSP images for PSSP 3.4

The RS/6000 SP package consists of several file sets that must be copied into the **/spdata/sys1/install/pssplpp/PSSP-3.4** directory using the **bffcreate** command. After copying the file sets, rename the PSSP and RSCT file sets and create the .toc file:

```
bffcreate -qvX -t /spdata/sys1/install/pssplpp/PSSP-3.4 -d /dev/cd0 all
cd /spdata/sys1/install/pssplpp/PSSP-3.4
mv ssp.3.4.0.0.I pssp.installp
mv rsct.basic.1.2.1.0.I rsct.basic
mv rsct.clients.1.2.1.0.I rsct.clients
mv rsct.core.1.2.1.0.I rsct.core
inutoc .
```

At this point, copy PSSP prerequisites to your AIX 4.3.3 lppsource. Issue:

```
cp xlC.rte.* /spdata/sys1/install/name/lppsource
cp xlC.aix43.* /spdata/sys1/install/name/lppsource
cp ipfx.* /spdata/sys1/install/name/lppsource
cp vacpp.ioc.* /spdata/sys1/install/name/lppsource
cp vacpp.cmp.* /spdata/sys1/install/name/lppsource
cd /spdata/sys1/install/name/lppsource
inutoc .
```

Refer to “Step 16: Copy the PSSP images” on page 26 for more information.

Step 14: Install the runtime prerequisites

PSSP 3.4 has prerequisites for vacpp.ioc.aix43.rte 5.0.2.0 and xlC.aix43.rte 5.0.2.0. These files and their associated prerequisites were moved to your AIX 4.3.3 lppsource **/spdata/sys1/install/***name***/lppsource** during “Step 13: Copy the PSSP images for PSSP 3.4” on page 140.

If your control workstation or nodes have parts of the xlC or vacpp products installed other than the files moved in “Step 13: Copy the PSSP images for PSSP 3.4” on page 140, your AIX lppsource must contain the files you have installed at the levels corresponding to the levels of the files shipped with PSSP or the files must be migrated before the control workstation or nodes are migrated. To determine the installed files do:

```
lslpp -L | grep xlC
lslpp -L | grep vacpp
```

If the output shows installed files different from the ones moved in “Step 13: Copy the PSSP images for PSSP 3.4” on page 140, you must obtain the maintenance corresponding to the files shipped with PSSP. Maintenance can be obtained from:

http://techsupport.services.ibm.com/rs6k/fixdb.html

The PSSP prerequisite files must be in your AIX lppsource. However, it is suggested that you place the maintenance in a directory other than your AIX lppsource and install the maintenance on both the control workstation and the nodes before beginning a migration.

To install only PSSP prerequisite runtime files, enter:

```
installp -agXd /spdata/sys1/install/name/lppsource xlC.rte \
         xlC.aix43.rte vacpp.ioc.aix43.rte
```

Alternatively, use SMIT or the **installp** command to install the maintenance from your own directory.

Step 15: Install (copy) the basic AIX (mksysb) image

The RS/6000 SP media provides a basic AIX minimal image for each level of AIX that PSSP 3.4 is supported on. You need to properly load in the AIX mksysb image to the **/spdata/sys1/install/images** directory using the **installp** command. For example, to install the AIX 4.3.3 minimal images, issue the following command:

```
installp -aXd /dev/cd0 spimg.433
```

Note: In order to reinstall your nodes, the mksysb image and the lppsource that you use must both contain the same version, release, modification, and fix levels of AIX. If you do not have a mksysb image at the same level as your lppsource, you may do one of the following:

1. Make your own updated mksysb image. In order to do this, you will need to:

a. Update an existing lppsource to the most recent maintenance level of AIX.
b. Perform a BOS node upgrade on a single node as described in "BOS node upgrade" on page 153 **or** follow the steps in "Installing updates on a per node basis" on page 264.
c. Make a mksysb image of that node as described in "Installing updates through reinstallation" on page 266.
d. Use the mksysb created in Step 1c along with your updated lppsource to install your remaining nodes.

2. Contact IBM Level 1 service to obtain an updated mksysb image.

Refer to "Step 17: Copy a basic AIX (mksysb) image" on page 28 for more information.

Step 16: Stop daemons on the control workstation and verify

Refer to the following table for the commands to issue to stop the daemons from running on the control workstation. You must stop the daemons in the order that they are presented in this table.

To stop this daemon:	Issue this command:
RSCT daemons	`syspar_ctrl -G -k`
sysctld daemon	`stopsrc -s sysctld`
splogd daemon	`stopsrc -s splogd`
hardmon daemon	`stopsrc -s hardmon`
sdrd daemons	`stopsrc -g sdr`

Issue the **lssrc -a** command to verify that the daemons are no longer running on the control workstation. The SRC objects in the table should now have a status of *inoperative* with no active process ID (PID).

Step 17: Obtain credentials

If DCE or Kerberos V4 was enabled in "Step 23: Set the authentication method for SP Trusted Services on the control workstation" on page 45, you must obtain credentials using **dce_login** or **k4init**. If DCE was selected, you should **dce_login** to the SP administrative principal created in "Step 22.3: Create SP administrative principals" on page 45. If Kerberos V4 was selected, you should use the appropriate administrative principal created in "Step 21: Initialize RS/6000 SP Kerberos V4 (optional)" on page 38.

Step 18: Install PSSP on the control workstation

You now need to install the PSSP 3.4 level code on the control workstation. The PSSP 3.4 file sets are packaged to be installed on top of previously-supported releases. If in the original installation of PSSP on the control workstation, all of the file sets available in the PSSP 3.4 package were installed, you should install all of them in this step also. For example, to install all of the file sets, enter:

```
installp -aXd /spdata/sys1/install/pssplpp/PSSP-3.4 all
```

If in the original installation, optional file sets were omitted, you can omit them in this step. For more information, refer to "Step 19: Install PSSP on the control workstation" on page 31.

You can use **installp** to install multiple file sets. For example:

```
installp -a -g -d /spdata/sys1/install/pssplpp/PSSP-3.4 -X ssp rsct
```

Step 19: Set authentication methods for SP Trusted Services

If your starting PSSP level is earlier than PSSP 3.2, the authentication method for SP Trusted Services may not have been set. To determine if the authentication method was previously set, issue the **lsauthts** command. If the response is a value of **compat** or **dce**, you can skip this step. If the authentication method for SP Trusted Services is not set to **compat** and you are migrating from a version of PSSP earlier than PSSP 3.2, issue the **chauthts** command. Issue:

```
chauthts compat
```

Step 20: Complete PSSP installation on the control workstation

To properly set up the PSSP 3.4 control workstation for the SDR, Hardmon, and other SP-related services, issue the following command:

```
install_cw
```

The **install_cw** command runs **SDR_init** which logs information in **/var/adm/SPlogs/SDR/SDR_config.log**.

For more information, see “Step 25: Complete system support installation on the control workstation” on page 46.

Step 21: Verify the authentication values in the SDR

Use the **splstdata** command to verify the authentication values for **auth_install**, **auth_root_rcmd**, **ts_auth_methods**, and **auth_methods**.

```
splstdata -p
```

Notes:

1. If you migrated from PSSP 3.2 or later, these settings should remain unchanged across the migration.
2. If you migrated from PSSP 3.1.1 or earlier, you should see the following authentication values:

   ```
   auth_install k4
   auth_root_rcmd k4
   ts_auth_methods compat
   auth_methods k4:std
   ```

Step 22: Run SDR and System Monitor verification test

Run verification tests that check for correct installation of the SDR, the System Monitor, and correct configuration of the System Monitor.

```
SDR_test     (SDR verification)
spmon_itest  (SP mon verification)
```

Step 23: Configure PSSP services and set up the site environment

You must ensure that the AIX level on the LP source (indicated by the **cw_lppsource_name**) matches the AIX level installed on your control workstation. To change any of the site environments, issue the **spsitenv** command:

```
spsitenv cw_lppsource_name=name
```

Note: If any of your nodes are running a version of PSSP earlier than PSSP 3.2, only ASCII data may be written to the SDR.

The system management environments on the control workstation are started by running **services_config**. Issue the following command:

```
/usr/lpp/ssp/install/bin/services_config
```

Step 24: Update the state of the supervisor microcode

Check which supervisors need to be updated by issuing:

```
spsvrmgr -G -r status all
```

If action is required, update the microcode by issuing:

```
spsvrmgr -G -u frame_number:slot_number
```

You can update all of the microcode at once by issuing:

```
spsvrmgr -G -u all
```

Note: When using the **spled** application, a node in the process of having the microcode on its supervisor card updated will not be displayed in the window.

For more information, see “Step 34: Update the state of the supervisor microcode” on page 60.

Step 25: Start RSCT subsystems and verify

There are RSCT subsystems that you need to add to the system. Even if you are not partitioning your system, you need to do this since you still have one default partition. Do the following:

1. Remove old subsystems (for example, hats, hags) by issuing the following command:

   ```
   syspar_ctrl -c -G
   ```

2. Add new subsystems (for example, hats, hags) by issuing the following command:

   ```
   syspar_ctrl -A -G
   ```

3. In order to monitor any new PSSP 3.4 hardware and to utilize the latest configuration database for Perspectives, you should perform the following steps at this time:

   a. Stop the Event Manager daemons in the systems partitions that contain the new hardware.

      Issue **/usr/sbin/rsct/bin/haemctrl -k** on the control workstation and on each of the nodes in the system partition. For PSSP 2.4 or earlier, issue **/usr/lpp/ssp/bin/haemctrl -k** on the control workstation and on each of the nodes in the system partition.

      You can use the **dsh** or **Sysctl** commands to run the command on multiple nodes from the control workstation. For more information on using these commands, see the “Parallel Management commands” chapter in *PSSP: Administration Guide*.

   b. Verify that all of the Event Manager daemons in each system partition have stopped.

      On the control workstation, issue the **lssrc -s haem.***domain_name* command. On the nodes, issue the **lssrc -s haem** command.

      You can use the **dsh** or **Sysctl** commands to run the command on multiple nodes from the control workstation. For more information on using these commands, see the “Parallel Management commands” chapter in the *PSSP: Administration Guide*.

The status of each daemon should indicate that it is inactive.

c. Restart the Event Manager daemons in each system partition.

   Issue **/usr/sbin/rsct/bin/haemctrl -s** on the control workstation and on each of the nodes in the system partition. For PSSP 2.4 or earlier, issue **/usr/lpp/ssp/bin/haemctrl -s** on the control workstation and on each of the nodes in the system partition.

   You can use the **dsh** or **Sysctl** commands to run the command on multiple nodes from the control workstation. For more information on using these commands, see the "Parallel Management commands" chapter in *PSSP: Administration Guide*.

To verify that the system-partition sensitive subsystems have been properly started, refer to "Step 56: Verify that RSCT subsystems have started" on page 82.

Step 26: Refresh the pmand daemons (PSSP 2.4 only)

This step is only required when migrating the control workstation from PSSP 2.4.

The **pmand** daemons out on the nodes need to be stopped and restarted in order to recognize changes that were made to the SDR in "Step 20: Complete PSSP installation on the control workstation" on page 143. You can accomplish this by running the following commands:

```
dsh -avG stopsrc -s pman

dsh -avG startsrc -s pman
```

Step 27: Start the switch and any quiesced applications

Any of the applications that you quiesced prior to migrating your control workstation should be started now if they have not already been automatically started. If you have an SP Switch, for each system partition issue the following command to restart your SP Switch:

```
Estart
```

Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd** on SP Switch systems or **startsrc -s swtadmd2** on SP Switch2 systems before issuing the **Estart** command.

Step 28: Run verification tests

You should validate that the PSSP 3.4 software has been properly installed on the control workstation by issuing the following commands:

```
SYSMAN_test
spmon_ctest
CSS_test              * run only if you have a switch
spverify_config       * run only if your system is partitioned
st_verify             * run only if Job Switch Resource Table Services
                        is installed
```

Note: If you are migrating nodes in more than one system partition, you need to run **CSS_test** in each of the system partitions.

Verify that **host_responds** and **switch_responds** , if you have a switch, are set to **yes** by issuing the following command:

```
spmon -d -G
```

Step 29: Update node description information

Using the **spgetdesc** command, you can obtain description information from the nodes and place it in the SDR for use by Perspectives and other applications. The **spgetdesc** command requires the nodes to be up in order to obtain the information from the node.

To obtain descriptions for all nodes and place it in the SDR, issue:

```
spgetdesc -au
```

If any nodes are not up when you run the **spgetdesc** command, you can reissue the command when those nodes are up by specifying a node list. For example:

```
spgetdesc -ul 1,3
```

See the **spgetdesc** command in the *PSSP: Command and Technical Reference* for more information.

Step 30: Validate the network adapters

If you are migrating from a level of PSSP 2.4 or earlier, the network adapters may have to be retuned. With PSSP 3.1, there are two new attributes that must be defined; otherwise, they will be assigned default values. The two new attributes are **duplex** and **Ethernet speed**.

To view the assigned values for these attributes, from the control workstation, issue:

```
splstdata -a
```

If the values in the **enet_rate** and **duplex** columns are correct, you do not have to do anything. To change adapter attributes, use the **spadaptrs** command. For a complete description of this command, refer to *PSSP: Command and Technical Reference*.

Step 31: Validate the control workstation

Your control workstation should now function at the PSSP 3.4 level. All nodes should now be able to communicate to the control workstation, if the proper level of PTF service has been applied to the PSSP nodes.

There are occasions based on customer changes made to the SDR and Authentication services, that may require the PSSP nodes to be rebooted and possibly recustomized by the boot/install server (BIS) nodes.

You can now add other AIX, PSSP, and customer-owned LPs and files on to the control workstation. Be careful to make sure that older applications and LPs will work properly with the target AIX level and PSSP 3.4.

Software maintenance (PTFs) may now be applied to the PSSP file sets installed on the control workstation. Refer to “Installing program updates” on page 262 for planning considerations. Follow the instructions in “Preparing the control workstation” on page 263 to install the PTFs.

Note: At this point, refer back to “High-level migration steps” on page 127 to determine your next step in the migration process.

Migrating the control workstation from AIX 4.3.3 to AIX 5L 5.1

Notes:

1. Before changing your AIX level from AIX 4.3.3 to AIX 5L 5.1, you must first migrate to PSSP 3.4.
2. Make sure your system is stable with PSSP 3.4 before migrating to AIX 5L 5.1.

Follow the steps in this section to migrate your control workstation from AIX 4.3.3 to AIX 5L 5.1.

Keep in mind that you should always make a mksysb backup of your control workstation and a backup of the **/spdata** directory before proceeding with the migration process. For information on making a mksysb, refer to “Step 10: Back up your control workstation” on page 132.

Step 1: Prepare to migrate and verify requirements

Before beginning, refer to “Preparing to migrate” on page 129 for information on preparing to migrate and verifying control workstation and system requirements.

Step 2: Quiesce your system

Be sure to quiesce your system as follows:

- Ensure all users are logged off nodes
- Stop all user jobs running over the switch
- Ensure that any jobs that start up automatically (for example batch submission queues), should be shut down or quiesced.
- If you are using the Switch Admin daemon for node recovery, stop it by issuing **stopsrc -s swtadmd** on SP Switch systems or **stopsrc -s swtadmd2** on SP Switch2 systems.
- Quiesce the switch using the **Equiesce** command

Step 3: Perform control workstation BOS migration install to AIX 5L 5.1

This method preserves all file systems except **/tmp**, as well as the root volume group, logical volumes, and system configuration files. This activity provides **installp** updates and installs necessary AIX LPs to the control workstation. You must have all the necessary AIX file sets and PTFs listed in the *READ THIS FIRST* document available by CD-ROM or from the NIM master. Note that you may need certain PTFs on your nodes if coexistence requirements exist between AIX 4.3.3 and AIX 5L 5.1.

Migration install from a CD-ROM is an interactive process. You will be prompted to verify settings, continue the installation, and verify the migration.

Refer to the *AIX Installation Guide* in chapters 1-3 for more information. The following list provides some hints:

- Boot off of your AIX CD-ROM or set up the control workstation as a NIM client from an appropriate NIM master.
- Select option 2 ″Change/Show Installation Settings and Install.″ Always verify that it has been set to *migrate* and that the correct target disk is assigned for root volume group.
- Select option 1 ″System Settings″ on the Installation and Settings menu.
- Select option 3 ″Migration Install″ on the Change Method of Installation menu.

- Select the proper hdisks being used for rootvg. You may want to update the hdisk configuration for the control workstation.
- Customize your control workstation based on the documented PSSP requirements. This includes installing the required AIX LPs and PTFs.

Step 4: Install the runtime prerequisites

PSSP has a prerequisite for vacpp.ioc.aix50.rte 5.0.1.0 when running with AIX 5L 5.1. This file is not part of the AIX installation package. The vacpp.ioc and vacpp.cmp installation files were put in your AIX 4.3.3 lppsource **/spdata/sys1/install/***name***/lppsource** during "Step 13: Copy the PSSP images for PSSP 3.4" on page 140. The runtime library associated with AIX 5L 5.1 must be installed now. Issue:

```
installp -aXd /spdata/sys1/install/name/lppsource vacpp.ioc.aix50.rte
```

Step 5: Verify AIX levels

Verify that the control workstation was successfully migrated to your target AIX level (AIX 5L 5.1), by issuing the the following command:

```
oslevel
```

For example, if your target AIX level is AIX 5L 5.1 and the output of this command does not indicate 5.1.0.0, issue the following command to return a list of AIX files not migrated to AIX 5L 5.1. You may need to install AIX PTFs to migrate those file sets to AIX 5L 5.1.

```
oslevel -l 5.1.0.0
```

Step 6: Review space requirements for NIM boot images

Before creating a PSSP boot/install server, ensure that there is sufficient space in the root (/) file system or create a separate file system for **/tftpboot** to manage the space required for the boot images (approximately 25 MB per lppsource level supported) created by NIM.

Step 7: Review space requirements for /spdata

The **/spdata** directory contains mksysb images and installp file sets. IBM suggests you create a separate volume group for the **/spdata** file system. These file sets are large and need much space (up to 2 GB per lppsource level supported). If you have not done so already, use *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* to help you estimate how much space you need to define.

Step 8: Create the required /spdata directories

You need to create the proper directory structure for AIX file sets.

Note: Make sure you mount the new **/spdata** file system before you create the **/spdata** directories.

You must create subdirectories on the **/spdata** file system for storing critical PSSP data. Make sure the directories have the permissions `rwxr-sr-x`. Issue the following commands to create the required directories:

To create **/spdata/sys1/install/***name***/lppsource**, issue:

```
mkdir -p /spdata/sys1/install/name/lppsource
```

This is the subdirectory for the required AIX 5L 5.1 file sets. You can choose any name, but the name must not contain any dots (.). If you have multiple lppsources,

you should pick more than one name. By default, this is set up as the string "default". You might want to pick a more meaningful name, such as *aix510* for AIX 5L 5.1 level code. If the control workstation will be the NIM server for different levels of AIX, you should create one subdirectory for each level of AIX.

Step 9: Copy the AIX LP images and other required AIX LPs and PTFs

In AIX 5L 5.1, changes were made to allow installation with installers other than **installp** to allow new installation media formats. With AIX 5L 5.1, two new commands (**geninstall** and **gencopy**) were introduced, which call **installp** or **bffcreate**, or other commands as appropriate. Additional subdirectories have also been added into the NIM LPP_SOURCE with AIX 5L 5.1. For NIM, instead of just putting everything in the LPP_SOURCE directory, appropriate subdirectories are created by the **gencopy** and **bffcreate** commands and the images are copied to those subdirectories.

You must copy the AIX file sets into the **/spdata/sys1/install/***name***/lppsource** directory on your hard disk on the control workstation. The **gencopy** and **bffcreate** commands place the files into the appropriate subdirectory which is **/spdata/sys1/install/***name***/lppsource/installp/ppc**.

You can download all of the AIX file sets (a very large number) or only the minimal required AIX file sets (approximately 1 GB).

The following is the minimal list of AIX file sets required to perform mksysb installations. The *prefix.** syntax in the list refers to everything that starts with the *prefix*. For example, **devices.*** refers to all of the file sets starting with **devices**.

Minimal list of AIX 5L 5.1 file sets:

```
IMNSearch.bld.*              bos.net.*
IMMSearch.rte.*              bos.perf.*
Java130.rte.*                bos.powermgt.*
Tivoli_Management_Agent.*    bos.rte.*
X11.Dt.*                     bos.svprint.*
X11.adt.*                    bos.sysmgt.*
X11.apps.*                   bos.terminfo.*
X11.base.*                   bos.txt.*
X11.compat.*                 bos.up.*
X11.fnt.*                    devices.*
X11.loc.*                    ifor_ls.base.*
X11.motif.*                  invscout.ldb.*
X11.msg.*                    invscout.rte.*
X11.vsm.*                    perl.*
bos                          perfagent.tools.*
bos.64bit.*                  printers.rte.*
bos.adt.*                    rpm.rte.*
bos.diag.*                   rsct.*
bos.doc*                     sysmgmt.help.msg.en_US.*
bos.help.msg.en_US.*         sysmgmt.msg.en_US.websm.*
bos.html.en_US.topnav.*      sysmgmt.sguide.*
bos.iconv.*                  sysmgmt.websm.*
bos.loc.iso.*                xlC.aix50.*
bos.man.en_US.*              xlC.cpp.*
bos.mp.*                     xlC.msg.en_US.cpp.*
bos.mp64.*                   xlC.rte.*
box.msg.en_US.*
```

Additional files you may want to add to your lppsource:

bos.acct.* Required if you plan to use PSSP accounting.

bos.cpr.* Required to install LoadLeveler 3.1 or Parallel Environment 3.2.

Java130.xml4j
Required for pSeries 690 servers

CIMOM Copy from the AIX toolbox for Linux applications CD. It is labeled in the Contents as openCIMOM, with the file name of Rpms/noarch/openCIMOM-0.61-1.aix5.1.noarch.rpm

dce.* The DCE file sets are required only if DCE will be configured by PSSP anywhere on the system. You will need the client portion of the DCE file sets because the installation code installs the DCE client code.

Notes:

1. Refer to your disk usage planning in the "Combining the space requirements" section of *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* to determine if you have allocated enough space to accomplish this task.
2. Allow at least 1-3 hours for moving all the file sets from media to disk.

To copy the AIX LP images, login to the control workstation as root and run **bffcreate** using SMIT or the command line. The following example shows the product media on **cd0** and the selection of all LPs. Using **all** may load unnecessary file sets into the directory.

```
bffcreate -qvX -t/spdata/sys1/install/name/lppsource -d /dev/cd0 all
```

The following warning message is issued—ignore it:

```
bffcreate:
Warning: important size information is missing from
the table of contents file.  Consequently, there may not
be enough free file system space to successfully create
the bff image(s).  Continuing anyway...
```

Step 10: Verify the correct level of PAIDE

The perfagent.server file set is part of the Performance Aide for AIX (PAIDE) feature of the Performance Toolbox for AIX (PTX), a separate product. This product provides the capability to monitor your SP system's performance, collects and displays statistical data for SP hardware and software, and simplifies runtime performance monitoring of a large number of nodes.

The Performance Toolbox for AIX, Agent Component (PAIDE) is required. The correct level of AIX PAIDE (perfagent) needs to be installed on the control workstation and copied to all of the lppsource directories. The perfagent.tools file set is part of AIX 5L 5.1. Verify that it was migrated correctly during the AIX 5L 5.1 migration. Issue:

```
lslpp -L | grep perfagent
```

You should receive output similar to the following:

```
perfagent.tools     5.1.0.0     C     F     Local Performance Analysis
```

Step 11: Obtain additional PSSP prerequisites

PSSP has prerequisites for certain compiler related files. These files are available on the PSSP 3.4 media. They were placed in the **/spdata/sys1/install/pssplpp/PSSP-3.4** directory in "Step 13: Copy the PSSP images for PSSP 3.4" on page 140. They must be copied into the AIX 5L 5.1 lppsource directory **/spdata/sys1/install/***name***/lppsource/installp/ppc**. Issue:

```
cd /spdata/sys1/install/name/lppsource/installp/ppc
cp /spdata/sys1/install/pssplpp/PSSP-3.4/xlC.rte.* .
cp /spdata/sys1/install/pssplpp/PSSP-3.4/xlC.aix50.* .
cp /spdata/sys1/install/pssplpp/PSSP-3.4/ipfx.* .
cp /spdata/sys1/install/pssplpp/PSSP-3.4/vacpp.ioc.* .
cp /spdata/sys1/install/pssplpp/PSSP-3.4/vacpp.cmp.* .
inutoc .
```

The PSSP prerequisite files may now be removed from the PSSP lppsource directory. Issue:

```
cd /spdata/sys1/install/pssplpp/PSSP-3.4
rm xlC*
rm ipfx*
rm vacpp*
inutoc .
```

If your control workstation or nodes have parts of the xlC or vacpp products installed, other than the files moved earlier in this step, then your AIX lppsource must contain the files you have installed at the levels corresponding to levels of the files shipped with PSSP or they must be installed before the migration. to determine the installed files do:

```
lslpp -L | grep xlC
lslpp -L | grep vacpp
```

If the output shows installed files different from the ones moved earlier, you must obtain the maintenance corresponding to the files shipped with PSSP. Maintenance can be obtained from:

http://techsupport.services.ibm.com/rs6k/fixdb.html

The PSSP prerequisite files must be in your AIX lppsource. However, it is suggested that you place the maintenance in a directory other than your AIX lppsource and install the maintenance on both the control workstation and the nodes before beginning a migration.

Step 12: Install (copy) the basic AIX (mksysb) image

The RS/6000 SP media provides a basic AIX minimal image for each level of AIX that PSSP 3.4 is supported on. You need to properly load in the AIX mksysb image to the **/spdata/sys1/install/images** directory using the **installp** command. For example, to install the AIX 5L 5.1 minimal images, issue the following command:

```
installp -aXd /dev/cd0 spimg.510
```

Note: In order to reinstall your nodes, the mksysb image and the lppsource that you use must both contain the same version, release, modification, and fix levels of AIX. If you do not have a mksysb image at the same level as your lppsource, you may do one of the following:

1. Make your own updated mksysb image. In order to do this, you will need to:
   a. Update an existing lppsource to the most recent maintenance level of AIX.
   b. Perform a BOS node upgrade on a single node as described in “BOS node upgrade” on page 153 **or** follow the steps in “Installing updates on a per node basis” on page 264.
   c. Make a mksysb image of that node as described in “Installing updates through reinstallation” on page 266.

d. Use the mksysb created in Step 1c on page 151 along with your updated lppsource to install your remaining nodes.

2. Contact IBM Level 1 service to obtain an updated mksysb image.

Refer to “Step 17: Copy a basic AIX (mksysb) image” on page 28 for more information.

Paths to migrate the nodes to PSSP 3.4

You cannot migrate the nodes until you have first migrated the control workstation to your target AIX level (AIX 4.3.3 or AIX 5L 5.1) and PSSP 3.4. Remember that the control workstation must always be at the latest level of AIX and PSSP to which you plan to migrate the nodes. Refer to the beginning of this chapter for that information. Keep in mind that you should always make a mksysb of all of your nodes that you want to migrate before proceeding with the migration process. For information on making a mksysb of your nodes, refer to “Step 11: Back up your nodes” on page 133.

You can migrate the nodes to your target AIX level and PSSP 3.4 in one of the ways identified in the following list:

- **BOS Node Upgrade**

  This method applies to AIX modification level changes (for example, AIX 4.3.*x* to 4.3.*x*+1) or when the AIX level is not changing, but you are migrating to a new level of PSSP (for example, AIX 4.3 and PSSP 2.4 to AIX 4.3 and PSSP 3.4). This method preserves the current rootvg and installs AIX PTF updates using the **installp** command. See “BOS node upgrade” on page 153.

- **BOS Node Migration Install**

  This method preserves all file systems except **/tmp**, as well as the root volume group, logical volumes, and system configuration files. This method requires the setup of AIX NIM on the new PSSP 3.4 control workstation. This applies only to migrations when an AIX version or release is changing (for example, AIX 4.3.3 to AIX 5L 5.1). See “BOS node migration install” on page 157.

  Use this method when migrating to AIX 5L 5.1 even if the PSSP level is already at PSSP 3.4.

- **BOS mksysb install**

  This method erases all existence of the current rootvg and installs your target AIX level and PSSP 3.4 using either an AIX 4.3 or AIX 5L 5.1 mksysb image for the node. This installation requires the setup of AIX NIM on the new PSSP 3.4 control workstation. See “mksysb install of nodes” on page 160.

Typically, a customer would migrate a node using either the migration install or upgrade methods depending upon the node’s current AIX and PSSP levels. If some nodes are identical, you could create a mksysb of the node and then migrate these nodes using the mksysb install method. Supported starting points for node migration include:

- PSSP 3.4 and AIX 4.3.3
- PSSP 3.2 and AIX 4.3.3
- PSSP 3.1.1 and AIX 4.3.3
- PSSP 2.4 and AIX 4.3.3

You can migrate the nodes from any of the starting points shown in the previous list to PSSP 3.4 either on AIX 4.3.3 or or AIX 5L 5.1.

Important

If the starting point for a node migration is PSSP 2.4 and AIX 4.2.1, you must first migrate to PSSP 3.4 and AIX 4.3.3. A second migration is required to go from PSSP 3.4 and AIX 4.3.3 to PSSP 3.4 and AIX 5L 5.1. The two migrations cannot be combined.

If you have xlC or vacpp installed on your nodes, it is suggested that you upgrade your installed files to the level corresponding to the levels of the files shipped with PSSP before migrating PSSP. The appropriate service should have been obtained in “Step 13: Copy the PSSP images for PSSP 3.4” on page 140 or “Step 11: Obtain additional PSSP prerequisites” on page 150. Maintenance for vacpp and xlC products have a wide ifreq web and if the required file sets are not available in the lppsource or installed on the node before the migration, the migration will fail.

If you do not have xlC or vacpp installed, or if the only pieces installed are those required by PSSP, (for example, your node was previously installed or migrated to PSSP 3.4 and AIX 4.3.3 and your xlC and vacpp files were installed during that process), the files do not need to be upgraded before the migration.

Before you migrate

If you choose the mksysb install method, any AIX corrective service applied to the mksysb must also be placed in the **lppsource** directory and the Shared Product Object Tree (SPOT) must be updated. Refer to the procedure documented in “Task E: Update the SPOT when installing AIX BOS service updates” on page 265.

If you choose the migration or upgrade methods, you need to make sure that the SPOT is up-to-date. Perform “Task E: Update the SPOT when installing AIX BOS service updates” on page 265.

BOS node upgrade

This method is used primarily when you need to:

- Migrate to a new level of PSSP without changing AIX levels.
- Migrate to a new level of PSSP and upgrade to a new modification level of AIX.
- To perform a BOS Upgrade only, perform “Step 1: Apply AIX 4.3.3 upgrade on the node”, “Step 2: Verify AIX migration” on page 154, and “Step 3: Reboot the node” on page 154.

This method applies AIX PTF Service and preserves the current rootvg disk configuration. It installs and upgrades the AIX BOS file sets. The **pssp_script** installs and updates the PSSP 3.4 LPs on top of current PSSP LPs.

Notes:

1. If you use this method, only the required PSSP file sets will be upgraded. Optional file sets must be upgraded manually.
2. For the following tasks, you can elect not to use **dsh** and issue commands on the node.

Step 1: Apply AIX 4.3.3 upgrade on the node

Issue the NFS **mount** command from the control workstation using **dsh** to mount the lppsource directory on the control workstation to the node.

```
dsh -w node "/usr/sbin/mount \
CWS:/spdata/sys1/install/lppsource_name/lppsource /mnt"
```

Issue the **lslpp** and **installp** commands to update all listed LPs on the node with AIX service found in the lppsource directory. You can issue the commands directly from the node.

```
dsh -w node "/usr/bin/lslpp -l -J  > /tmp/FILE.43"
dsh -w node "/usr/sbin/installp -acNgXd /mnt -f /tmp/FILE.43"
```

Step 2: Verify AIX migration

Verify that the node has successfully been migrated to the target AIX level by issuing the following command:

```
dsh -w node "/bin/oslevel"
```

For example, if the target AIX level is 4.3.3 and the output of this command does not indicate AIX 4.3.3, issue the following command to return a list of AIX files not migrated to AIX 4.3.3. You may need to install AIX PTFs to migrate those file sets to AIX 4.3.3.

```
dsh -w node "/bin/oslevel -l 4.3.3.0"
```

Step 3: Reboot the node

If you just upgraded your node (an AIX modification level change, for example, AIX 4.3.2 to 4.3.3), you now need to reboot the node so changes to the kernel will take effect. To do this, issue:

```
cshutdown -rGFN node_list
```

Step 4: Enter node configuration data

You need to set the appropriate SDR node object attributes for *lppsource_name*, *code_version*, *bootp_response*, and *pv_list* for each node being upgraded. Use the **spchvgobj** and **spbootins** commands to update these fields. If you are migrating nodes in more than one system partition, you will need to issue these commands in each system partition. For a complete description of the flags associated with these commands, refer to *PSSP: Command and Technical Reference*.

For example, to migrate nodes 1 and 2 to AIX 4.3.3 and PSSP 3.4 where lppsource was placed in **/spdata/sys1/install/aix433/lppsource**, issue the following two commands:

```
spchvgobj -r selected_vg -p PSSP-3.4 -v aix433 -h 00-00-00-0,0 \
          -l 1,2 -i bos.obj.ssp.433

spbootins -s no -r customize -l 1,2
```

Note: If you update the mksysb name for your system partition, you must reapply the system partition configuration using the **spapply_config** command. For more information on system partitioning, refer to the “Managing system partitions” chapter in *PSSP: Administration Guide*.

Step 5: Verify installation settings

Make sure that the SDR has the appropriate values specified in the following attributes for each of the nodes. Issue the following command to display the values:

```
splstdata -G -b
```

- response

  Nodes being migrated should be set to **customize**.

- lppsource_name

  The name of the lppsource of the AIX level to be used for the migration (for example, *aix433*).
- pssp_ver

  The *code_version* of PSSP should be set to "PSSP-3.4" for nodes you are migrating.
- pv_list

  The disk to install on, preferably in the hardware location format.

Make sure that **/tftpboot/script.cust** and **/tftpboot/firstboot.cust** have been properly updated for PSSP 3.4 modifications. See “Appendix E. User-supplied node customization scripts” on page 297 for additional information.

Step 6: Refresh RSCT subsystems

The SDR has now been updated to reflect the new nodes that will run PSSP 3.4. You now need to refresh the RSCT subsystems on the control workstation and all nodes to pick up these changes. Run **syspar_ctrl** on the control workstation to refresh the subsystems on both the control workstation and on the nodes.

```
syspar_ctrl -r -G
```

Note: You can ignore any messages that you receive from the nodes you are migrating at this point because the migration process is not yet complete. Once the process is complete, you should no longer receive error messages.

Step 7: Run setup_server to configure the changes

The **setup_server** command must be run to properly set up NIM on the control workstation by issuing the following command:

```
setup_server 2>&1 | tee /tmp/setup_server.out
```

The output will be saved in a log file called `setup_server.out`.

If you have a node defined as a boot/install server you must also run **setup_server** out on that node.

```
dsh -w boot/install_node "/usr/lpp/ssp/bin/setup_server \
    2>&1" | tee /tmp/setup_server.boot/install_node.out
```

Step 8: Disable nodes from the switch

If you do not have a switch in your SP system, skip this step.

If you want to bring the switch down for all nodes, issue the **Equiesce** command for each system partition.

Note: If you are using the Switch Admin daemon for node recovery, stop it by issuing **stopsrc -s swtadmd** on SP Switch systems or **stopsrc -s swtadmd2** on SP Switch2 systems before issuing the **Equiesce** command.

If you use the **Equiesce** command, you will need to later restart the switch using the **Estart** command. Issue the **Estart** command prior to the step where you “Verify the nodes.”

If you are migrating a few nodes, you must disable these nodes from the switch (if appropriate, first reassign the primary node or primary backup node). To determine if one of the nodes you are migrating is a primary or primary backup node, issue

the **Eprimary** command. If you need to reassign the primary or primary backup node, issue the **Eprimary** command with appropriate options. Then issue the **Estart** command to make your choices effective. You must then issue the **Efence** command to disable the nodes you are migrating from the switch.

```
Efence -G node_number node_number
```

Step 9: Copy the PSSP 3.4 pssp_script to nodes' /tmp

Copy the PSSP 3.4 version of **pssp_script** from the control workstation to the **/tmp** directory on each of the nodes you are migrating. For example:

```
pcp -w node /spdata/sys1/install/pssp/pssp_script /tmp/pssp_script
```

where *node* is the host name of one or more nodes.

Step 10: Execute the pssp_script on the node

Execute the **pssp_script** that you copied to **/tmp** on all of the nodes you are migrating:

```
dsh -w node /tmp/pssp_script
```

You need to wait until the script has completed before proceeding to the next step. To determine if **pssp_script** is still running on the nodes, monitor the 3 digit LEDs and wait for them to turn blank. Check the **bootp_response** and ensure that it is set to disk by issuing the following command:

```
splstdata -G -b
```

You must also update any optional PSSP file sets that you have installed on the node.

Step 11: Reboot the node

If your system contains a switch, you must reboot the node at this time so that changes to the kernel can take affect. Any other kernel change that may have occurred when you upgraded AIX would also require that you reboot the node at this time.

IBM suggests that you reboot the node at this time. To do this, issue:

```
cshutdown -rGFN node_list
```

Step 12: Rejoin the nodes to the switch network

If you disabled all nodes in "Step 8: Disable nodes from the switch" on page 155 using the **Equiesce** command, you must now issue the **Estart** command in each system partition to rejoin the nodes to the current switch network.

Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd** on SP Switch systems or **startsrc -s swtadmd2** on SP Switch2 systems before issuing the **Estart** command.

If you disabled only a few nodes using the **Efence** command, you must now issue the **Eunfence** command to bring those nodes back to the switch network.

Step 13: Start RSCT subsystems

If you rebooted your node in "Step 11: Reboot the node", you can skip this step because this step would have been automatically performed when the node was rebooted.

To start the subsystems, do the following:

1. Remove the old subsystems by issuing:

   ```
   dsh -w node /usr/lpp/ssp/bin/syspar_ctrl -c
   ```

2. Add and start the new subsystems by issuing:

   ```
   dsh -w node /usr/lpp/ssp/bin/syspar_ctrl -A
   ```

To verify that the subsystems are running, refer to “Step 56: Verify that RSCT subsystems have started” on page 82.

Step 14: Run verification tests

Verify that the nodes are running properly by issuing the following commands:

```
SYSMAN_test
CSS_test              * run only if you have a switch
spverify_config       * run only if your system is partitioned
st_verify             * run only if Job Switch Resource Table Services
                        is installed
```

Verify that **host_responds** and **switch_responds**, if you have a switch, are set to **yes** by issuing the following command:

```
spmon -d -G
```

Note: If you are migrating nodes in more than one system partition, you need to run **CSS_test** in each of the system partitions.

Step 15: Apply PSSP PTFs to nodes (optional)

Software maintenance (PTFs) may now be applied to the PSSP file sets installed on the nodes. If you installed software maintenance on the control workstation before the node migration, you must install the service updates to the nodes now. Refer to “Installing updates on a per node basis” on page 264 for detailed instructions.

Note: At this point, refer back to “High-level migration steps” on page 127 to determine your next step in the migration process.

BOS node migration install

Before proceeding with the steps in this section, be sure you have selected the correct migration path. A BOS node migration is used when you are changing your AIX version or release, for example, from AIX 4.3.3 to AIX 5L 5.1.

Step 1: Enter node configuration data

You need to properly set the appropriate SDR node object attributes for *lppsource_name*, *code_version*, *bootp_response*, and *pv_list* for each node being migrated. Use the **spchvgobj** and **spbootins** commands to update these fields. If you are migrating nodes in more than one system partition, you will need to issue these commands in each system partition. For a complete description of the flags associated with these commands, refer to *PSSP: Command and Technical Reference*.

Note: Even though it is not necessary to specify **–i** *install_image* flag on the **spchvgobj** command during migration, it is suggested at this point to avoid future problems.

For example, to migrate nodes 1 and 2 to AIX 5L 5.1 and PSSP 3.4 where lppsource was placed in **/spdata/sys1/install/aix510/lppsource**, issue the following two commands:

```
spchvgobj -r selected_vg -p PSSP-3.4 -v aix510 -h 00-00-00-0,0 \
          -l 1,2 -i bos.obj.ssp.510

spbootins -s no -r migrate -l 1,2
```

For example, to migrate nodes 1 and 2 to AIX 5L 5.1 where lppsource was placed in **/spdata/sys1/install/aix510/lppsource**, issue the following two commands. Note that because this example does not include the **−p** flag, the nodes must already be migrated to PSSP 3.4. This sets up for an AIX migration only.

```
spchvgobj -r selected_vg -v aix510 -h 00-00-00-0,0 -l 1,2 \
          -l 1,2 -i bos.obj.ssp.510

spbootins -s no -r migrate -l 1,2
```

Step 2: Verify installation settings

Make sure that the SDR has the appropriate values specified in the following attributes for each of the nodes. Issue the following command to display the values:

```
splstdata -G -b -l node_list
```

- response

  Nodes being migrated should be set to *migrate*.

- lppsource_name

  The name of the lppsource of the AIX level to be used for the migration (for example, *aix510*).

- pssp_ver

  The *code_version* of PSSP should be set to "PSSP-3.4" for the nodes you are migrating.

- pv_list

  The disk to install on, preferably in the hardware location format.

Make sure that **/tftpboot/script.cust** and **/tftpboot/firstboot.cust** have been properly updated for PSSP 3.4 modifications. See “Appendix E. User-supplied node customization scripts” on page 297 for additional information.

Step 3: Run setup_server to configure the changes

Note: Do not run **setup_server** on any nodes defined as a boot/install server until the boot/install servers have been migrated.

The **setup_server** command must be run to properly set up NIM on the control workstation by issuing the following command:

```
setup_server 2>&1 | tee /tmp/setup_server.out
```

The output will be saved in a log file called `setup_server.out`.

If you have a node defined as a boot/install server, you must also run **setup_server** out on that server node.

```
dsh -w boot/install_node "/usr/lpp/ssp/bin/setup_server \
    2>&1" | tee /tmp/setup_server.boot/install_node.out
```

Step 4: Refresh RSCT subsystems

The SDR has now been updated to reflect the new nodes that will run PSSP 3.4. You now need to refresh the RSCT subsystems on the control workstation and all nodes to pick up these changes. Run **syspar_ctrl** on the control workstation to refresh the subsystems on both the control workstation and on the nodes.

```
syspar_ctrl -r -G
```

Note: You can ignore any messages that you receive from the nodes you are migrating at this point because the migration process is not yet complete. Once the migration is complete, you should no longer receive error messages.

Step 5: Disable nodes from the switch

If you do not have a switch in your SP system, skip this step.

If you want to bring the switch down for all nodes, issue the **Equiesce** command for each system partition.

Note: If you are using the Switch Admin daemon for node recovery, stop it by issuing **stopsrc -s swtadmd** on SP Switch systems or **stopsrc -s swtadmd2** on SP Switch2 systems before issuing the **Equiesce** command.

If you use the **Equiesce** command, you will need to later restart the switch using the **Estart** command. Issue the **Estart** command prior to the step where you "Verify the nodes."

If you are migrating a few nodes, you must disable these nodes from the switch (if appropriate, first reassign the primary node or primary backup node). To determine if one of the nodes you are migrating is a primary or primary backup node, issue the **Eprimary** command. If you need to reassign the primary or primary backup node, issue the **Eprimary** command with appropriate options. Then issue the **Estart** command to make your choices effective. You must then issue the **Efence** command to disable the nodes you are migrating from the switch.

```
Efence -G node_number node_number
```

Step 6: Shut down the node

To shut down the node gracefully, issue the following command:

```
cshutdown -F -G -N node_number
```

Step 7: Network boot the node

Notes:

1. If you have any boot/install servers in your system, you need to migrate them before migrating their clients. You should not netboot more than eight nodes with the same server at a time.
2. For MCA nodes, the **nodecond** command remotely processes information from the initial AIX firmware menus. You should not change the language option on these menus. The language must be set to English in order for the **nodecond** command to run properly.

Network boot each node that you are migrating by using Perspectives or by using the **nodecond** command.

```
nodecond -G frame_id slot_id &
```

You should notice that the node has been properly installed when the LED's become blank, and the host_responds is active.

Verify that the **bootp_response** has been set to disk by issuing the following command:

```
splstdata -G -b
```

Step 8: Rejoin the nodes to the switch network

If you disabled all nodes in "Step 5: Disable nodes from the switch" on page 159 using the **Equiesce** command, you must now issue the **Estart** command in each system partition to rejoin the nodes to the current switch network.

Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd** on SP Switch systems or **startsrc -s swtadmd2** on SP Switch2 systems before issuing the **Estart** command.

If you disabled only a few nodes using the **Efence** command, you must now issue the **Eunfence** command to bring those nodes back to the switch network.

Step 9: Run verification tests

Verify that the nodes are running properly by issuing the following commands:

```
SYSMAN_test
CSS_test              * run only if you have a switch
spverify_config       * run only if your system is partitioned
st_verify             * run only if Job Switch Resource Table Services
                        is installed
```

Verify that **host_responds** and **switch_responds**, if you have a switch, are set to **yes** by issuing the following command:

```
spmon -d -G
```

Note: If you are migrating nodes in more than one system partition, you need to run **CSS_test** in each of the system partitions.

Step 10: Apply PSSP PTFs to nodes (optional)

Software maintenance (PTFs) may now be applied to the PSSP file sets installed on the nodes. If you installed software maintenance on the control workstation before the node migration, you must install the service updates to the nodes now. Refer to "Installing updates on a per node basis" on page 264 for detailed instructions.

Note: At this point, refer back to "High-level migration steps" on page 127 to determine your next step in the migration process.

mksysb install of nodes

Before proceeding with the steps in this section, be sure you have selected the correct migration path. This method erases all existence of the current rootvg and installs your target AIX level and PSSP 3.4 using either an AIX 4.3 or AIX 5L 5.1 mksysb image for the node.

Step 1: Enter node configuration data

You need to set the appropriate SDR node object attributes for *lppsource_name*, *code_version*, *bootp_response*, *next_install_image*, and *pv_list* for each node being

migrated. Use the **spchvgobj** and **spbootins** commands to update these fields. If you are migrating nodes in more than one system partition, you will need to issue these commands in each system partition. For a complete description of the flags associated with these commands, refer to *PSSP: Command and Technical Reference*.

Note: The lppsource, PSSP code version, and install image you select in this step must be available on the control workstation. See "Migrating the control workstation to PSSP 3.4" on page 133.

For example, to migrate nodes 1 and 2 to AIX 4.3.3 and PSSP 3.4 where lppsource was placed in **/spdata/sys1/install/aix433/lppsource**, issue the following two commands:

```
spchvgobj -r selected_vg -p PSSP-3.4 -v aix433 -h 00-00-00-0,0 \
          -l 1,2 -i bos.obj.ssp.433

spbootins -s no -r install -l 1,2
```

Note: In order to reinstall your nodes, the mksysb image and the lppsource that you use must both contain the same version, release, modification, and fix levels of AIX. If you do not have a mksysb image at the same level as your lppsource, you may do one of the following:

1. Make your own updated mksysb image. In order to do this, you will need to:
   a. Update an existing lppsource to the most recent maintenance level of AIX.
   b. Perform a BOS node upgrade on a single node as described in "BOS node upgrade" on page 153 **or** follow the steps in "Installing updates on a per node basis" on page 264.
   c. Make a mksysb image of that node as described in "Installing updates through reinstallation" on page 266.
   d. Use the mksysb created in Step 1c along with your updated lppsource to install your remaining nodes.
2. Contact IBM Level 1 service to obtain an updated mksysb image.

Step 2: Verify installation settings

Make sure that the SDR has the appropriate values specified in the following attributes for each of the nodes. Issue the following command to display the values:

```
splstdata -G -b
```

- response

  Nodes being migrated should be set to **install**.
- lppsource_name

  The name of the lppsource of the AIX level to be used for the migration (for example, *aix433*).
- pssp_ver

  The *code_version* of PSSP should be set to "PSSP-3.4" for nodes you are migrating.
- next_install_image

  Should be set to the appropriate AIX mksysb image (for example, *bos.obj.ssp.433*).
- pv_list

  The disk to install on, preferably in the hardware location format.

Make sure that **/tftpboot/script.cust** and **/tftpboot/firstboot.cust** have been properly updated for PSSP 3.4 modifications. See “Appendix E. User-supplied node customization scripts” on page 297 for additional information.

Step 3: Run setup_server to configure the changes

The **setup_server** command must be run to properly set up NIM on the control workstation by issuing the following command:

```
setup_server 2>&1 | tee /tmp/setup_server.out
```

The output will be saved in a log file called `setup_server.out`.

If you have a node defined as a boot/install server, you must also run **setup_server** out on that server node.

```
dsh -w boot/install_node "/usr/lpp/ssp/bin/setup_server \
    2>&1" | tee /tmp/setup_server.boot/install_node.out
```

Step 4: Refresh RSCT subsystems

The SDR has now been updated to reflect the new nodes that will run PSSP 3.4. You now need to refresh the RSCT subsystems on the control workstation and all nodes to pick up these changes. Run **syspar_ctrl** on the control workstation to refresh the subsystems on both the control workstation and on the nodes.

```
syspar_ctrl -r -G
```

Note: You can ignore any messages that you receive from the nodes you are migrating at this point because the migration process is not yet complete. Once the process is complete, you should no longer receive error messages.

Step 5: Disable nodes from the switch

If you do not have a switch in your SP system, skip this step.

If you want to bring the switch down for all nodes, issue the **Equiesce** command for each system partition.

Note: If you are using the Switch Admin daemon for node recovery, stop it by issuing **stopsrc -s swtadmd** on SP Switch systems or **stopsrc -s swtadmd2** on SP Switch2 systems before issuing the **Equiesce** command.

If you use the **Equiesce** command, you will need to later restart the switch using the **Estart** command. Issue the **Estart** command prior to the step where you “Verify the nodes.”

If you are migrating a few nodes, you must disable these nodes from the switch (if appropriate, first reassign the primary node or primary backup node). To determine if one of the nodes you are migrating is a primary or primary backup node, issue the **Eprimary** command. If you need to reassign the primary or primary backup node, issue the **Eprimary** command with appropriate options. Then issue the **Estart** command to make your choices effective. You must then issue the **Efence** command to disable the nodes you are migrating from the switch.

```
Efence -G node_number node_number
```

Step 6: Shut down the node

Nodes should be shut down gracefully using the following command:

```
cshutdown -F -G -N node_number
```

Step 7: Unconfigure DCE-related information for the node (required for DCE)

Issue the **splstdata -p** command and examine the security settings for the system partition containing the nodes to be migrated. If **auth_install** includes DCE, you must remove any DCE-related principles and objects from the DCE registry before issuing the **nodecond** command.

Note: You must have cell administrator authority to perform this step.

1. On the control workstation, use the **rm_spsec -t admin** *node_dce_hostname* command for each node being reinstalled.

   Note: To run this command remotely off of the SP, you must set the SP_NAME environment variable to point to the SDR you want to access. Refer to the **rm_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.

2. Do a DCE Admin unconfigure for the node (**smit rmdce**).

   Note: To remove any additional principals related to the node using the SMIT panels, enter the host name of the adapter to be deleted. For example, on the “Admin unconfiguration for another machine” panel in the “Machine’s name or TCP/IP address” field, enter the host name for the additional adapters.

3. For the nodes being removed, verify that all DCE principals have been deleted from the DCE registry. Issue:

   ```
   dcecp -c principal catalog -simplename
   ```

You must now create new DCE information for the node by performing the following steps:

1. Run the **setupdce** command.

   Notes:

   a. You will be prompted for the cell administrator’s password when you issue this command.

   b. To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being configured. The value must be a resolvable address. For example:

      ```
      export SP_NAME=spcws.abc.com
      ```

2. As an ID with cell administrator authority, run the **config_spsec -v** command.

   Note: To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being configured. Refer to the **config_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.

PSSP 3.1.1 and DCE exception

If at the start of your control workstation migration your system contained PSSP 3.1.1 and DCE with **auth_methods** set to **k5:k4:std**, you will need to do one of the following:

1. Automatically install DCE during the mksysb installation by putting code in your **script.cust** to do the install. Because the node was previously installed, you should review DCE documentation for additional unconfiguration and reconfiguration steps that will be required. This is the same process you should have developed for doing a mksysb install of a DCE node with PSSP 3.1.

2. Have PSSP automatically install DCE and configure the DCE clients by doing the following:
   a. Ensure that DCE is at 3.1 or later on the control workstation and that your AIX lppsource contains the DCE file sets.
   b. Use **spsetauth -i** to set **auth_install** to include DCE.

      ```
      spsetauth -p partition1 -i dce k4
      ```

   c. Define DCE host names for the control workstation and for all the nodes:

      ```
      create_dcehostname
      ```

   d. Update the SDR with DCE Master Security and CDS Server host names:

      ```
      setupdce -u -s master_security_server_host -d CDS_primary_server_host
      ```

   e. Remove existing self-host principles for the nodes being installed from the DCE database. This command may need to be reissued for each adaptor on each node that is being reinstalled to PSSP 3.4.

      ```
      /bin/unconfig.dce -config_type admin -dce_hostname old_dcehostname \
      -host_id adapter_host_name all
      ```

   f. Add the nodes being installed back into the DCE database. You will need DCE cell administrator authority to run the **setupdce** command.

      ```
      setupdce -v
      ```

Step 8: Network boot the node

Notes:

1. If you have any boot/install servers in your system, you need to migrate them before migrating their clients. You should not netboot more than eight nodes with the same server at a time.
2. For MCA nodes, the **nodecond** command remotely processes information from the initial AIX firmware menus. You should not change the language option on these menus. The language must be set to English in order for the **nodecond** command to run properly.

Network boot each node that you are migrating by using Perspectives or by using the **nodecond** command.

```
nodecond -G frame_id slot_id &
```

You should notice that the node has been properly installed when the LED's become blank, and the host_responds is active.

Verify that the **bootp_response** has been set to disk by issuing the following command:

```
splstdata -G -b
```

Step 9: Rejoin the nodes to the switch network

If you disabled all nodes in "Step 5: Disable nodes from the switch" on page 162 using the **Equiesce** command, you must now issue the **Estart** command in each system partition to rejoin the nodes to the current switch network.

Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd** on SP Switch systems or **startsrc -s swtadmd2** on SP Switch2 systems before issuing the **Estart** command.

If you disabled only a few nodes using the **Efence** command, you must now issue the **Eunfence** command to bring those nodes back to the switch network.

Step 10: Run verification tests

Verify that the nodes are running properly by issuing the following commands:

```
SYSMAN_test
CSS_test              * run only if you have a switch
spverify_config       * run only if your system is partitioned
st_verify             * run only if Job Switch Resource Table Services
                        is installed
```

Verify that **host_responds** and **switch_responds**, if you have a switch, are set to **yes** by issuing the following command:

```
spmon -d -G
```

Notes:

1. If you are migrating nodes in more than one system partition, you need to run **CSS_test** in each of the system partitions.
2. At this point, refer back to "High-level migration steps" on page 127 to determine your next step in the migration process.

HACWS migration strategy

An HACWS configuration at the PSSP 3.4 level requires the following software on both control workstations:

- PSSP 3.4 (including the ssp.hacws 3.4.0.0 file set).
- Any level of AIX that is supported with PSSP 3.4. Refer to the *READ THIS FIRST* document to determine what levels of AIX are supported with PSSP 3.4. At the time this manual was published, PSSP 3.4 was supported with AIX 4.3.3 and AIX 5L 5.1.
- Any level of HACMP that is supported with the level of AIX that you are using. Refer to the appropriate HACMP documentation to determine what levels of HACMP are supported with the level of AIX that you are using or considering. At the time this manual was published, AIX 4.3.3 was supported with HACMP 4.4, HACMP 4.4.1, HACMP Enhanced Scalability (HACMP/ES) 4.4, and HACMP/ES 4.4.1. AIX 5L 5.1 was supported with HACMP 4.4, HACMP 4.4.1, and HACMP/ES 4.4.1.

Whether or not you need to upgrade all three of these at the same time, depends on your software levels before migration. You can choose to upgrade your HACWS configuration gradually, stopping along the way to run your system long enough to become confident that it is stable before proceeding to the next phase. Just be sure that you always run your HACWS configuration with a supported combination of AIX, HACMP, and PSSP releases. For example, you may want to upgrade AIX while remaining on the old HACMP release and the old PSSP release. When you become confident that your system is stable, you can proceed to upgrade HACMP or PSSP in a later service window, after resolving any problems related to the AIX upgrade. This strategy is allowed only if the combination of AIX, HACMP, and PSSP releases that will be installed during the interim period is a supported combination. If the combination is not supported, it may be necessary to upgrade two of these software products within a single service window. In some cases, typically when you skip a few PSSP releases, it may even be necessary to upgrade all three software products within the same service window. While there is nothing wrong with upgrading everything at once, such an approach makes it more difficult to pinpoint the cause of problems that may occur during the migration.

If you implement a gradual HACWS migration strategy (within the rules outlined previously), you may also choose to upgrade AIX or HACMP on each control

workstation separately and avoid an SP system outage. For example, you might take the backup control workstation offline to upgrade AIX or HACMP, while you keep the primary control workstation in production. After you finish migrating the backup control workstation, you could move control workstation services to it, and take the primary control workstation offline to perform the same migration of AIX or HACMP. If you use this approach to upgrade HACMP, you should find out whether the old and new HACMP software levels can coexist within the same HACMP cluster. If they cannot coexist, then you cannot perform a graceful failover to move control workstation services. Instead you have to stop HACMP on the currently active control workstation, and wait for it to stop completely before starting HACMP on the takeover control workstation. Such a coexistence limitation should also encourage you to migrate the second control workstation very soon after the first, or a failure of the currently active control workstation would require manual intervention.

You cannot upgrade the PSSP software level within your HACWS configuration until both control workstations have been migrated to supported levels of both AIX and HACMP, and while you may avoid an SP system outage by upgrading AIX and HACMP on each control workstation separately, you cannot use this technique to upgrade PSSP. You must upgrade PSSP on both control workstations at the same time, so there is no way to avoid an SP system outage while you upgrade PSSP.

At this point, you should apply these rules to your situation and plan your own strategy for migrating your HACWS configuration to PSSP 3.4.

Note: You must complete your migration to PSSP 3.4 on both control workstations before migrating to AIX 5L 5.1.

High-level HACWS migration instructions

By now you should have read “HACWS migration strategy” on page 165 and planned your HACWS migration strategy. If you have not done this, you should read the section now. Refer to the items in this section as you implement your HACWS migration strategy.

Backing up your HACWS configuration

Before you begin your HACWS migration, you should back up the root volume group on each control workstation, and you should also back up the shared volume group which contains the **/spdata** file system. If anything goes wrong with your HACWS migration, you can recover by restoring the appropriate backup images. If you separate your HACWS migration into multiple service windows, then you should backup your HACWS configuration each time, in order to make sure the backup images are current.

You may back up the root volume group by issuing the **mksysb** command directly, or you may invoke the **mksysb** command through SMIT by issuing the command **smit mksysb**. You may back up the shared volume group which contains the **/spdata** file system by issuing the **savevg** command directly, or you may invoke the savevg command through SMIT by issuing the command **smit savevg**. Refer to the AIX documentation for more detailed instructions.

AIX migration considerations

Before you can migrate your HACWS control workstations to PSSP 3.4, you must migrate to AIX 4.3.3. “Migrating the control workstation to PSSP 3.4” on page 133 describes two different methods for upgrading AIX on a non-HACWS control

workstation. You should upgrade AIX on an HACWS control workstation by either applying AIX PTF service, if this is possible, or performing a migration install of AIX. Refer to the *AIX Installation Guide* for AIX installation instructions. If you perform a migration install of AIX, verify that any changes that were made to **/etc/rc.net** on the backup control workstation during “Step 31: Add IP address aliases” on page 124 do not get lost. These changes should be preserved by the AIX migration installation, but if they are not, they will need to be re-created in order for your HACWS configuration to work correctly.

Once you have completed the installation procedure, you should verify the migration of AIX by running the following command:

```
oslevel
```

The output of this command is the current AIX software level. If you are not at the level you expected, you can also use this command to get a list of file sets that need to be updated. For example, to find out which file sets must be updated in order to reach AIX 4.3.3, issue the following command:

```
oslevel -l 4.3.3.0
```

You may need to install AIX PTFs in order to migrate those file sets to the desired AIX software level.

Whenever you make a significant change to your HACWS configuration, such as upgrading the AIX software level, you should always run the HACMP **clverify** utility to verify the HACMP cluster upon which your HACWS configuration is based. You should resolve any problems that **clverify** reports before you consider the AIX upgrade to be successful. However, if you are going to upgrade HACMP immediately after you upgrade AIX, then it makes sense to delay running **clverify** until after you upgrade HACMP. Refer to the “Verifying cluster software” and “Verifying the cluster topology” chapters of the *HACMP for AIX: Installation Guide* for more information about this utility.

In addition, you can use the **hacws_verify** command to check for problems with your HACWS configuration. You will be asked to run this command when you upgrade PSSP. However, if you upgrade AIX and then put your HACWS configuration back into production before you migrate to PSSP 3.4, then you should run this command after you start control workstation services on the primary control workstation.

If using:	**Do this:**
SMIT	**TYPE** **smit hacws** • The High Availability Control Workstation Management window appears. **SELECT** Verify HACWS Installation and Configuration
hacws_verify	Enter: `/usr/sbin/hacws/hacws_verify`

Note: The **hacws_verify** command requires that the primary control workstation must be the active control workstation, or the one which is currently providing the control workstation services. If control workstation services are being provided by the backup control workstation, you cannot run this command.

HACMP migration considerations

Before you can migrate your HACWS control workstations to PSSP 3.4, you must install any level of HACMP that is supported with the level of AIX that you are using with PSSP 3.4. Refer to the "Upgrading an HACMP cluster" chapter of the *HACMP for AIX: Installation Guide* for HACMP migration installation instructions.

After you upgrade HACMP on both control workstations, you should run the HACMP **clverify** utility to verify the HACMP cluster upon which your HACWS configuration is based. You should resolve any problems that **clverify** reports before you consider the HACMP upgrade to be successful. Refer to the "Verifying cluster software" and "Verifying the cluster topology" chapters of the *HACMP for AIX: Installation Guide* for more information about this utility.

After you upgrade HACMP on both control workstations, you need to identify to HACMP the HACMP pre- and post-event scripts that are provided by HACWS.

If using:	Do this:
SMIT	**TYPE** **smit hacws** • The High Availability Control Workstation Management window appears. **SELECT** Identify Event Scripts to HACMP **PRESS** Enter to continue.
spcw_addevents	Enter: `/usr/sbin/hacws/spcw_addevents`

You already performed this task, when you originally installed HACWS. However, a new release of HACMP may introduce more cluster events, so you need to repeat this task, so that HACWS-supplied pre- and post-event scripts get defined to HACMP for the newly introduced cluster events. Failure to do this will not hurt the functionality of your HACWS configuration, but it will cause the **hacws_verify** command to fail.

In addition to the HACMP **clverify** utility, you can use the **hacws_verify** command to check for problems with your HACWS configuration. You will be asked to run this command when you upgrade PSSP. However, if you upgrade HACMP and then put your HACWS configuration back into production before you migrate to PSSP 3.4, you should run the command after you start control workstation services on the primary control workstation.

If using:	Do this:
SMIT	**TYPE** **smit hacws** • The High Availability Control Workstation Management window appears. **SELECT** Verify HACWS Installation and Configuration
hacws_verify	Enter: `/usr/sbin/hacws/hacws_verify`

Note: The **hacws_verify** command requires that the primary control workstation must be the active control workstation, or the one which is currently providing the control workstation services. If control workstation services are being provided by the backup control workstation, you cannot run this command.

PSSP migration steps

Proceed to “Migrating an HACWS configuration to PSSP 3.4”.

Migrating an HACWS configuration to PSSP 3.4

Follow the steps in this section if you are upgrading an HACWS configuration from PSSP 2.4, PSSP 3.1.1, or PSSP 3.2 to PSSP 3.4.

Prerequisites

Before following the instructions in this section, you should:

1. Read “Preparing to migrate” on page 129 and follow the instructions which describe how to prepare to migrate your SP system.
2. Read “HACWS migration strategy” on page 165 and plan your HACWS migration strategy.
3. Read “High-level HACWS migration instructions” on page 166 and perform all of the tasks that must be done prior to upgrading the PSSP software level of your HACWS configuration.

Step 1: Verify the authentication value for AIX remote commands

An authentication method must be set in order for remote commands to work properly. If your SP system contains nodes at PSSP 3.1.1 or earlier, the Kerberos V4 value must be set. For SP systems containing nodes at PSSP 3.2 or later, any value is valid (Kerberos V5, Kerberos V4, or Standard AIX).

If the value is not set, use the **chauthent** command to enable **k4**. Issue:

```
chauthent -k4 -std
```

Perform this step for both control workstations.

Step 2: Verify network tunable values

The suggested network tunable values for your control workstations are as follows:

```
sb_max              163840
ipforwarding             1
tcp_sendspace        65536
tcp_recvspace        65536
udp_sendspace        32768
udp_recvspace        65536
tcp_mssdflt           1448
tcp_pmtu_discover        0
udp_pmtu_discover        0
```

Verify that the network tunable values on both control workstations are correct. Use the **no** command to display these values. To display all of the network tunable values at once, issue the command:

```
/usr/sbin/no -a
```

To change the value of the thewall network tunable to 16384, issue the command:

```
/usr/sbin/no -o thewall=16384
```

When you change a network tunable value, the change takes effect immediately. However, the change is not preserved across a reboot. To make the change permanent, add the required no command to the bottom of the **/etc/rc.net** script. Refer to the **no** command man page for more information about changing network tunable values.

Step 3: Review space requirements for NIM boot images

Follow the instructions in “Step 7: Review space requirements for NIM boot images” on page 138.

Step 4: Review space requirements for /spdata

Follow the instructions in “Step 9: Review space requirements for /spdata” on page 138.

Step 5: Create the required /spdata directories

Follow the instructions in “Step 10: Create the required /spdata directories” on page 138.

Step 6: Copy the AIX LP images and other required AIX LPs and PTFs

Follow the instructions in “Step 11: Copy the AIX 4.3.3 LP images and other required AIX LPs and PTFs” on page 139.

Step 7: Install the correct level of PAIDE on both control workstations

For both control workstations, follow the instructions in “Step 12: Install correct level of PAIDE on the control workstation” on page 140.

Step 8: Copy the PSSP images for PSSP 3.4

Follow the instructions in “Step 13: Copy the PSSP images for PSSP 3.4” on page 140.

Step 9: Install the runtime prerequisites

Follow the instructions in “Step 14: Install the runtime prerequisites” on page 141. Perform this step for both control workstations.

Step 10: Install (copy) the basic AIX (mksysb) image

Follow the instructions in “Step 15: Install (copy) the basic AIX (mksysb) image” on page 141.

Step 11: Start control workstation services on the primary control workstation

If you have not already done so, make sure that HACMP is running on both control workstations and the primary control workstation is the active control workstation, or the one that is providing the control workstation services.

Step 12: Stop control workstation services while HACMP is running

While HACMP is still running, you need to stop the control workstation applications on both control workstations. While the primary control workstation is the active control workstation, the backup control workstation will act as a client of the primary control workstation. You need to turn this off by issuing the following command on the backup control workstation:

```
/usr/sbin/hacws/spcw_apps -d
```

Next, you need to stop control workstation services on the primary control workstation by issuing the same command on the primary control workstation.

Step 13: Install PSSP on both control workstations

You need to install the PSSP 3.4 file sets on both control workstations. When the PSSP file sets are installed, files are copied into the **/spdata/sys1** directory. This requires special consideration for an HACWS configuration, because both control workstations share the same **/spdata** file system, which resides in a shared volume group. At this point, the **/spdata** file system should be mounted on the primary control workstation, so files copied to the **/spdata/sys1** directory on the primary control workstation get written to the shared **/spdata** file system, and files copied to the **/spdata/sys1** directory on the backup control workstation get written to the backup control workstation's / (or root) file system. Since both control workstations share the same **/spdata** file system, the files that are copied to the backup control workstation are not needed, and they will get removed later by the **install_hacws** command in "Step 16: Complete PSSP installation on both control workstations". In the meantime, you need to make sure there is at least 4 MB of free space in the / (or root) file system on the backup control workstation, or installation on the backup control workstation may fail.

You can install the PSSP 3.4 file sets on the primary control workstation from either the **/spdata/sys1/install/pssplpp/PSSP-3.4** directory, which is in the locally mounted **/spdata** file system, or you can install directly from the installation media. However, IBM suggests that you install the PSSP 3.4 file sets on the backup control workstation directly from the installation media. You should not mount the **/spdata** file system on the backup control workstation to install PSSP. Now you may proceed to install the PSSP 3.4 file sets on both control workstations. Remember that you must also install the **ssp.hacws 3.4.0.0** file set on both control workstations. You can perform file set installation from the command line by issuing the installp command, or you can use SMIT by issuing the **smit install_latest** command. Refer to the *AIX Installation Guide* for more details on how to use these commands.

Step 14: Authenticate as the Kerberos V4 administrative principal

To get the Kerberos V4 administrative principal, follow the instructions in "Step 17: Obtain credentials" on page 142.

Step 15: Set authentication methods for SP Trusted Services

If your starting PSSP level is earlier than PSSP 3.2, the authentication method for SP Trusted Services may not have been set. To determine if the authentication method was previously set, issue the **lsauthts** command. If the response is **compat**, you can skip this step. If the authentication method for SP Trusted Services is not set to **compat** and you are migrating from a version of PSSP earlier than PSSP 3.2, issue the **chauthts** command. Issue:

```
chauthts compat
```

Step 16: Complete PSSP installation on both control workstations

Run the **install_hacws** command. The **install_hacws** command runs **SDR_init** which logs information in **/var/adm/SPlogs/SDR/SDR_config.log**.

If using:	Do this:
SMIT	**TYPE** **smit hacws** • The High Availability Control Workstation Management window appears. **SELECT** Install and Configure HACWS. **ENTER** The node names of the primary control workstation and backup control workstation in the HOSTNAME fields. **SELECT** **yes** for Execute on both primary and backup? **PRESS** **Ok** to complete option selection and install HACWS.
install_hacws	Enter: `/usr/sbin/hacws/install_hacws -p primary_hostname -b backup_hostname -s`

Note: Do this step from the primary control workstation.

Step 17: Verify the authentication values in the SDR

Follow the instructions in “Step 21: Verify the authentication values in the SDR” on page 143.

Step 18: Verify the HACWS configuration

Run the **hacws_verify** command.

If using:	Do this:
SMIT	**TYPE** **smit hacws** • The High Availability Control Workstation Management window appears. **SELECT** Verify HACWS Installation and Configuration
hacws_verify	Enter: `/usr/sbin/hacws/hacws_verify`

Note: The **hacws_verify** command requires that the primary control workstation must be the active control workstation, or the one which is currently providing the control workstation services. If control workstation services are being provided by the backup control workstation, you cannot run this command.

Step 19: Run SDR and System Monitor verification test

Follow the instructions in “Step 22: Run SDR and System Monitor verification test” on page 143.

Step 20: Configure PSSP services and set up the site environment

Follow the instructions in “Step 23: Configure PSSP services and set up the site environment” on page 143.

Step 21: Update the state of the supervisor microcode

Follow the instructions in “Step 24: Update the state of the supervisor microcode” on page 144.

Step 22: Restart control workstation services

Start control workstation services on the primary control workstation by issuing the following command:

```
/usr/sbin/hacws/spcw_apps -u
```

After this command completes on the primary control workstation, get the backup control workstation to synchronize with the primary control workstation by issuing the same command on the backup control workstation.

Step 23: Recycle the Event Manager daemons

In order to monitor any new PSSP 3.4 hardware and to utilize the latest configuration database for Perspectives, you should perform the following steps at this time:

1. Stop the Event Manager daemons in the systems partitions that contain the new hardware.

   Issue **/usr/sbin/rsct/bin/haemctrl -k** on the control workstation and on each of the nodes in the system partition. For PSSP 2.4 or earlier, issue **/usr/lpp/ssp/bin/haemctrl -k** on the control workstation and on each of the nodes in the system partition.

   You can use the **dsh** or **Sysctl** commands to run the command on multiple nodes from the control workstation. For more information on using these commands, see the “Parallel Management commands” chapter in *PSSP: Administration Guide*.

2. Verify that all of the Event Manager daemons in each system partition have stopped.

   On the control workstation, issue the **lssrc -s haem.***domain_name* command. On the nodes, issue the **lssrc -s haem** command.

   You can use the **dsh** or **Sysctl** commands to run the command on multiple nodes from the control workstation. For more information on using these commands, see the “Parallel Management commands” chapter in the *PSSP: Administration Guide*.

   The status of each daemon should indicate that it is inactive.

3. Restart the Event Manager daemons in each system partition.

   Issue **/usr/sbin/rsct/bin/haemctrl -s** on the control workstation and on each of the nodes in the system partition. For PSSP 2.4 or earlier, issue **/usr/lpp/ssp/bin/haemctrl -s** on the control workstation and on each of the nodes in the system partition.

   You can use the **dsh** or **Sysctl** commands to run the command on multiple nodes from the control workstation. For more information on using these commands, see the “Parallel Management commands” chapter in *PSSP: Administration Guide*.

Step 24: Refresh the pmand daemons (PSSP 2.4 only)

Follow the instructions in Step “Step 26: Refresh the pmand daemons (PSSP 2.4 only)” on page 145.

Step 25: Run verification tests

Follow the instructions in “Step 28: Run verification tests” on page 145.

Step 26: Validate the network adapters

Follow the instructions in “Step 30: Validate the network adapters” on page 146.

Step 27: Validate the control workstations

Your migration to PSSP 3.4 is now complete. You should test your HACWS configuration by performing a control workstation failover from the primary control workstation to the backup control workstation, and then failover from the backup control workstation back to the primary control workstation. Perform additional failover testing as needed. Refer to “Step 31: Validate the control workstation” on page 146 for additional instructions.

At this point, refer back to “High-level migration steps” on page 127 to determine the next step in your SP system migration process.

Migrating an HACWS configuration to AIX 5L 5.1

Follow the steps in this section to migrate your HACWS control workstations from AIX 4.3.3 to AIX 5L 5.1.

Note: Before changing your AIX level from AIX 4.3.3 to AIX 5L 5.1, you must first migrate to PSSP 3.4.

Prerequisites

Before following the steps in this section, should:

1. Read “Preparing to migrate” on page 129 for information on preparing to migrate and verifying control workstation and system requirements.
2. Read “HACWS migration strategy” on page 165 and plan your HACWS migration strategy. In some situations, it may be possible to upgrade the AIX level of your HACWS configuration without an SP system outage.
3. Read “High-level HACWS migration instructions” on page 166 and perform all of the tasks that must be done prior to upgrading the AIX software level of your HACWS configuration. Be sure to back up your HACWS configuration before you upgrade your AIX software level.

Step 1: Perform control workstation BOS migration install to AIX 5L 5.1

For both control workstations, follow the instructions in “Step 3: Perform control workstation BOS migration install to AIX 5L 5.1” on page 147. This step should be done while the control workstation is inactive (not providing control workstation services to the SP system).

Step 2: Install the runtime prerequisites

For both control workstations, follow the instructions in “Step 4: Install the runtime prerequisites” on page 148. This step should be done while the control workstation is inactive (not providing control workstation services to the SP system).

Step 3: Verify AIX levels

For both control workstations, follow the instructions in “Step 5: Verify AIX levels” on page 148. This step should be done while the control workstation is inactive.

Step 4: Verify the correct level of PAIDE

For both control workstations, follow the instructions in “Step 10: Verify the correct level of PAIDE” on page 150. This step should be done while the control workstation is inactive.

Step 5: Review space requirements for NIM boot images

Follow the instructions in “Step 6: Review space requirements for NIM boot images” on page 148.

Step 6: Review space requirements for /spdata

Follow the instructions in “Step 7: Review space requirements for /spdata” on page 148.

Step 7: Create the required /spdata directories

Follow the instructions in “Step 8: Create the required /spdata directories” on page 148.

Step 8: Copy the AIX LP images and other required AIX LPs and PTFs

Follow the instructions in “Step 9: Copy the AIX LP images and other required AIX LPs and PTFs” on page 149.

Step 9: Obtain additional PSSP prerequisites

Follow the instructions in “Step 11: Obtain additional PSSP prerequisites” on page 150.

Step 10: Install (copy) the basic AIX (mksysb) image

Follow the instructions in “Step 12: Install (copy) the basic AIX (mksysb) image” on page 151 .

Step 11: Validate the control workstations

Your migration to AIX 5L 5.1 is now complete. You should test your HACWS configuration by performing a control workstation failover from the primary control workstation to the backup control workstation, and then failover from the backup control workstation back to the primary control workstation. Perform additional failover testing as needed.

At this point, refer back to “High-level migration steps” on page 127 to determine the next step in your SP system migration process.

Post-migration activity

This section discusses the activities you should perform after a successful migration. It also discusses recovery procedures you should perform if the migration process was not successful.

Remove obsolete files and resources

After migrating the control workstation to PSSP 3.4, you may be supporting nodes at mixed levels of AIX and PSSP for a period of time. Once all the nodes have been migrated to AIX level *n* and PSSP level *p* and none of the PSSP levels that you are still supporting are dependent on AIX level *n-1*, you can remove all NIM resources and files associated with this old level of AIX and PSSP. You may want to

back up these files prior to removing them if there is any chance that you may one day need to reinstall nodes at this level of AIX and PSSP. You can remove the files listed in the following list:

- All of the NIM resources associated with AIX level *n-1*; the lppsource, spot, and mksysb.
- All of the AIX files under **/spdata/sys1/install** associated with AIX level *n-1*.
- All of the PSSP files under **/spdata/sys1/install/pssplpp** associated with PSSP level *p-1*.

Note: If you do not plan on supporting back-level system partitions from the control workstation after migrating to PSSP 3.4, you should remove the **ssp.jm** file set from the control workstation.

For example, if you completed migrating all of your nodes from PSSP 2.4 and AIX 4.2.1 to PSSP 3.4 and AIX 4.3.3 and you do not have any nodes that are running on AIX 4.2.1, you may now remove the following NIM resources and files. Assume that your mksysb image is named **bos.obj.ssp.421** and that your lppsource is stored in the directory named AIX421.

1. Remove the NIM resources associated with AIX 4.2.1; the lppsource, spot, and mksysb. To remove the lppsource resource, issue the command:

   ```
   nim -o remove lppsource_AIX421
   ```

2. Remove the spot and all files that NIM generated for this spot by issuing the following command:

   ```
   nim -o remove spot_AIX421
   ```

3. Remove the mksysb resource by first displaying a list of all resources of type mksysb by issuing the command:

   ```
   lsnim -t mksysb -l
   ```

   Determine which mksysb is associated with bos.obj.ssp.421. For example, if after examining the output of the previous command, you see that bos.obj.ssp.421 is associated with mksysb_x, then remove the mksysb_x resource by issuing the command:

   ```
   nim -o remove mksysb_x
   ```

   If the resource removal fails, you may first need to deallocate the resources from all clients using the **unallnimres** command.

   Then, remove the AIX files associated with AIX 4.2.1 (lppsource, mksysb) and PSSP 2.4 by issuing the following commands:

   ```
   rm -r /spdata/sys1/install/AIX421
   rm -r /spdata/sys1/install/images/bos.obj.ssp.421
   rm -r /spdata/sys1/install/pssplpp/PSSP-2.4
   ```

Recovery procedures

If you encounter a migration failure which you cannot resolve, you can call an IBM service representative for additional help. In addition, refer to the following list for a basic set of instructions for recovering your system. You may use the backups you made prior to the migration.

Recovering from a PTF migration failure

- If you are recovering from a failed PTF installation on the control workstation, clean up the possible failed installation by issuing:

  ```
  /usr/sbin/installp -C
  ```

and then reject the applied PTFs on the control workstation by issuing:

```
/usr/sbin/installp -r ALL
```

- If the PTFs failed on SP nodes, clean up possible failed installation by issuing:

  ```
  dsh -w node /usr/sbin/installp -C
  ```

  and reject the applied PTFs on the nodes by issuing:

  ```
  dsh -w node /usr/sbin/installp -r ALL
  ```

Recovering from a control workstation migration failure

If you are recovering from a control workstation migration failure, perform the following steps:

1. Insert the tape, the mksysb backup, that you created in “Step 10: Back up your control workstation” on page 132.
2. Change the key to the service position. If your control workstation is a PCI-based RS/6000, press the F1 or F2 key (depending on the model) at boot time. This provides a menu. From this menu, select the tape as boot device and press Enter. If your control workstation is a microchannel-based RS/6000, follow the instructions on the screen.
3. Select the disks to install the rootvg volume group.
4. Reinstall the control workstation using the mksysb install option.
5. Wait for the control workstation to reboot.
6. Login as root user after successful completion of the restore.
7. Authenticate as the Kerberos V4 Administrative Principal by issuing the **k4init** command. For example:

   ```
   k4init root.admin
   ```

8. Complete the PSSP installation on the control workstation by issuing the command:

   ```
   install_cw
   ```

9. Check to see if your SDR is correct by issuing the **spmon -d** command and the **splstdata** command. If the SDR is corrupt, you may restore a previously archived SDR by issuing the command:

   ```
   sprestore_config SDR_archive_name
   ```

10. Restore any other files that you may have saved.

Recovering from a node migration failure

If you are recovering from an SP node migration failure, perform the following steps:

1. Place the SP node mksysb image file that you previously saved in “Step 11: Back up your nodes” on page 133 into the directory **/spdata/sys1/install/images**.
2. Update the SDR to match the proper settings for the SP node being restored. Issue the **spchvgobj** command with **-r** *volume_group_name*, **-v** *old_lppsource_name*, **-p** *old_code_version*, **-i** *old_install_image* **-l** *node_list* and **spbootins** with **-s no**, **-l** *node_list* **-c** *volume_group* **-r install** for the node.
3. SP User Management changes some AIX file attributes. These changes must be undone or errors will occur later. To undo the changes, do the following:
   - Issue the following command to determine the state of **usermgmt_config**:

     ```
     /usr/lpp/ssp/bin/splstdata -e | grep usermgmt_config
     ```

     If **usermgmt_config** is **true**, issue the following command to change the state to **false**:

     ```
     spsitenv usermgmt_config=false
     ```

4. Network boot the SP node using the **nodecond** command or Perspectives. This will reinstall the SP node back to the same AIX and PSSP level of the backup mksysb image.

   Note: For MCA nodes, the **nodecond** command remotely processes information from the initial AIX firmware menus. You should not change the language option on these menus. The language must be set to English in order for the **nodecond** command to run properly.

   ```
   nodecond -G frame_id slot_id &
   ```

5. Restore any other user volume groups or files that were saved in “Step 11: Back up your nodes” on page 133.
6. Change the state of **usermgmt_config** to what it was previously.

Chapter 5. Reconfiguring security

This chapter explains the tasks necessary to add an authentication configuration to an existing SP system. Changing security configurations on an existing system is not very complex from an administrative perspective. One major restriction is that you cannot move an existing system partition between security states without a common security configuration. You must first add and establish the new security configuration to the nodes before you can remove the security configuration that you no longer want from the system partition. The commands that follow are explained more fully in “Chapter 2. Installing and configuring a new RS/6000 SP system” on page 11, however, not all of the steps shown in that chapter are needed when adding an authentication configuration.

For example, you may want to implement DCE in a system that had been migrated to PSSP 3.4 from PSSP 3.2. This system would have been installed and migrated using Kerberos V4 security. The initial settings, as shown by the **splstdata -p** command, may look similar to the following:

```
auth_install k4
auth_root_rcmd k4
ts_auth_methods compat
auth_methods k4:std
```

As the first step in going to a DCE-only system, you will initialize DCE in the system partition by running the security setup steps on the control workstation. During this process, the security settings would be changed to:

```
auth_install dce:k4
auth_root_rcmd dce:k4
ts_auth_methods dce:compat
auth_methods k5:k4:std
```

After the nodes have completed the transition to **dce**:**k4**, Kerberos **k4** can be deleted from the system partition. You would again run the security setup steps to remove **k4** and **compat** from the attributes. The final security settings may look similar to the following:

```
auth_install dce
auth_root_rcmd dce
ts_auth_methods dce
auth_methods k5
```

The following steps refer to a system partition. In some cases, they will need to be repeated for each system partition on the system.

As with many of the PSSP commands, you must have the appropriate authority or credentials to use these commands. See “Step 24: Obtain credentials” on page 46.

Adding DCE to the SP system

Restrictions

1. You cannot use both DCE authentication and HACWS.
2. You cannot use IPv6 aliasing with DCE, HACMP, and HACWS.

Notes:

1. If you currently have a level of DCE installed on your system that is earlier than DCE 3.1, you will need to migrate from that level to DCE 3.1 if you plan to configure SP Trusted Services to use DCE. Refer to *IBM Distributed Computing Environment for AIX: Quick Beginnings* for more information on how to migrate DCE.
2. For "Tips for installing DCE on the SP," see "Step 22: Configure DCE for the control workstation (required for DCE)" on page 44.

Perform the following steps to add DCE to your SP system:

1. If you do not want the DCE primary server to run on the control workstation, it must be accessible on some external system.
2. Install the DCE security client, directory client, and RPC (and the servers, if desired) on the control workstation. You must ensure that DCE is properly configured and running on the control workstation before further configuration of DCE in the system.
3. To indicate that DCE security should be installed and configured on the nodes, issue:

   ```
   spsetauth -p partition1 -i dce k4
   ```

   The preceding example assumes that Kerberos V4 was the current setting.
4. To define DCE host names for the control workstation and for all of the nodes, issue:

   ```
   create_dcehostname
   ```

   If **create_dcehostname** was run previously, it is not necessary to run it again unless new nodes were added to the system.
5. To update the SDR with DCE Master Security and CDS Server host names, issue:

   ```
   setupdce -u -s master_security_server_host -d cds_primary_server_host
   ```
6. In this step, you will be prompted to enter the cell administrator's password. You do not need to be root to run this command. Optionally, you can use the **-c** and **-l** flags or you can accept the defaults for the cell administrator ID and the LAN profile ID. To configure the "admin" portion of the nodes' DCE clients, issue:

   ```
   setupdce
   ```

   Notes:

   a. You can stop at this point in the configuration if you only want to install and configure DCE clients on the node without enabling the SP system to use DCE services. The DCE clients will be installed the next time the nodes are rebooted. You will need to continue with the remaining steps to enable DCE usage.

   b. To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being configured. The value must be a resolvable address. For example:

      ```
      export SP_NAME=spcws.abc.com
      ```
7. To configure SP Trusted Services to use DCE authentication, issue:

   ```
   config_spsec -v
   ```

   Notes:

   a. You must be logged in as the cell administrator to perform this task.

b. To run this command remotely off of the SP, you must set the SP_NAME environment variable to point to the SDR you want to access. Refer to the **config_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.

8. To create SP Trusted Services keyfiles and keytab objects, issue:

   ```
   create_keyfiles -v
   ```

   Note: You must be root with default DCE credentials to perform this task.

9. To select DCE as an authorization method for AIX remote commands, issue:

   ```
   spsetauth -p partition1 -d dce k4
   ```

   This step generates the necessary authorization files for each selected method and removes files or entries that are not needed. When adding **dce**, you will need to add it to the current setting. This implies that if **k4** was previously set, it must also be set now.

   Note: To enable DCE for authenticated remote commands, but not for SP Trusted Services, you can skip steps involving SP Trusted Services (Step 10 and Step 14) and continue to Step 11.

10. To start the Key Management daemon on the control workstation, issue:

    ```
    /usr/lpp/ssp/bin/spnkeyman_start
    ```

11. All affected nodes must be shut down. Use the **cshutdown** command (without the **-r** flag because the nodes should not be rebooted at this time).
12. To create and authorize a DCE SP administrative principal, follow the instructions in “Step 22.3: Create SP administrative principals” on page 45. You should obtain credentials for the DCE SP administrative principal before performing the next step.
13. To enable authentication methods for AIX remote commands, issue:

    ```
    chauthpar -c -p partition1 k5 k4
    ```

14. To enable authentication methods for SP Trusted Services, issue:

    ```
    chauthpts -c -p partition1 dce compat
    ```

15. Reboot all affected nodes.

    Before rebooting all affected nodes, see Note 1 on page 180.

16. Run **updauthfiles** on all nodes in system partitions not changed during the transition to DCE. Issue:

    ```
    export SP_NAME=partition_name

    dsh -av /usr/lpp/ssp/bin/updauthfiles
    ```

Your system is now configured and enabled to use DCE as an authentication method.

Adding Kerberos V4 to the SP system

This section describes the tasks involved in adding Kerberos V4 to the SP system.

Installing and configuring Kerberos V4

If any node in a partition is running a level earlier than PSSP 3.4, you must install and configure Kerberos V4 and activate it as an authentication method in that SP system partition. In addition to Kerberos V4, you can also select to install and configure DCE for each system partition.

Setting up PSSP authentication

To provide Kerberos V4 authentication services on your SP system, choose from the following three authentication implementations.

1. The authentication services provided with the SP system, based on MIT Kerberos V4. SP authentication services are designed to provide an easily-installable and configurable Kerberos for the SP system.
2. An existing implementation of MIT Kerberos V4, provided it is compatible with the SP Kerberos-authenticated services.
3. AFS version 3.3 or later.

Once you have installed and initialized Kerberos V4 authentication on your SP system, you can add other IBM RS/6000 workstations to your authentication realm. Two reasons for adding additional workstations might be:

1. If you are not using AFS authentication servers, you may want the added reliability of one or more secondary SP authentication servers. Providing secondary servers is particularly important if your SP system has to provide high availability with no single points of failure.
2. If you also want to install the SP authenticated services on office workstations. This will help you to manage the SP system, without having to login to the control workstation or the SP nodes directly.

When you use the SP authentication server support in the **ssp.authent** option, you may choose to have more than one server system. If you configure your local authentication realm with more than one server, you must designate one as the primary server. All others are secondary servers. Only the primary server has the administrative daemon, **kadmind**, that manages the content of the authentication database. The databases used by the secondary authentication servers are copies of the primary database and are updated periodically from the primary authentication server.

Additional servers can provide for greater reliability, since the authentication service will try all servers listed in the configuration file before failing a service request. If there are network problems, or if your primary server system fails, authentication requests can be handled as long as one of the configured servers is active and accessible. There may also be performance reasons to configure a secondary server. Only when there is an authentication server running on the control workstation can the switch be used for authentication protocol traffic to and from the SP nodes. In configurations where you are integrating the SP system into an existing authentication realm, where the primary authentication server is already on another workstation, you can make this possible by setting up a secondary server on the control workstation. In this way, you could have multiple SP systems in the same authentication realm, and give all SP nodes access to an authentication server across the switch.

Adding principals and changing passwords takes place in the primary database and is propagated to secondary databases through periodic updates only. Secondary databases are maintained by the **kpropd** daemon that runs only on secondary server workstations and receives the database content in encrypted form from the **kprop** program that runs on the primary server workstation. The SP authentication services include a script that you can schedule for execution in the root **crontab** file to keep your secondary authentication databases up-to-date.

Configuration files: Use the following configuration files located in the **/etc** directory to set up your system's authentication realm:

- **/etc/inetd.conf**

This file contains information used by the internet routing daemon **inetd** to route incoming requests for service to one of a large number of dynamically started daemons that are named in the file. When the SP authentication services are installed on your systems, the file is updated to route “kshell” service requests to the AIX Kerberos-authenticated remote command daemon (krshd). You should not modify this information, but you might have to resolve conflicts with names of locally installed services.

- **/etc/services**

  This file maps the names of network services to the well-known ports that they use to receive requests from their clients. This file may already contain entries relating to authentication services, since the file shipped with AIX 4.1 contains entries used by DCE. In addition to the DCE entries, many other reserved port names were added in this version of AIX, including an entry for the Kerberos V5 port, 88. The service name for this port is given as “kerberos”, which is also the name used by the standard MIT Kerberos V4 service. The port number usually assigned to the Kerberos V4 service is 750. In order to be consistent with and interoperate with AIX 3.2.5 systems running PSSP 1.2 with authentication services based on Kerberos V4, it was necessary to use the name “kerberos4”. You do not have to create an entry in the file for “kerberos4”, because the default port of 750/udp will be used if no “kerberos4” entry is found. If you are not using AFS Version 3.4 authentication servers, you should only have to modify **/etc/services** if you are using some other service that uses one or more of the traditionally used (but not formally reserved) Kerberos V4 ports. They are:

    Service name: kerberos4 Port: 750/udp
    Service name: kerberos_admin Port: 751/tcp
    Service name: krb_prop Port: 754/tcp

  You will also have to modify this file if you are using AFS Version 3.4 authentication servers. The **kaserver** in AFS Version 3.4 for AIX 4.1 accepts Kerberos V4 protocol requests using the well-defined udp port assigned to the “kerberos” service assigned to port 88 in the file distributed with the base operating system. MIT Kerberos V4, on which PSSP authentication services are based, uses a default port number of 750. PSSP authentication commands use the service name “kerberos4” to avoid this conflict with the Kerberos V5 service name, which is also used by DCE. For PSSP authentication commands to communicate with an AFS 3.4 **kaserver**, you must do either of the following:

  1. Stop the **kaserver**, redefine the udp port number for the “kerberos” service to 750 on the server system, then restart the **kaserver**.
  2. Add a statement to **/etc/services** that defines the udp port for the “kerberos4” service as 88 on the SP control workstation and on any other independent workstation that will be a client system for PSSP authenticated services.

- **/etc/krb.conf**

  This file contains the name of the local realm for your SP system and identifies the host names of all the authentication servers for all Kerberos realms known to the local realm. For more information on the content of the file, see **krb.conf** in the book *PSSP: Command and Technical Reference*.

  If you configure your realm to use AFS authentication servers, this file is built for you automatically from the corresponding AFS configuration file **CellServDB**. When you use PSSP authentication servers, you create this file or one is generated by default by the **setup_authent** script that must run on each workstation after you install the PSSP files. You must provide the file if the workstation is a client workstation or a secondary authentication server. If it is the

primary authentication server, you need to provide the file only if the system's host name contains no domain portion, following conventional network naming rules.

By convention, Kerberos uses the domain portion of a network interface name, converted to upper case, as the default realm to which the host belongs. If you want to define your own local realm name that does not follow convention, or if the host name of your primary server host has no domain portion, you must supply a file with your choice of local realm name in the first line. When using AFS authentication, the local AFS cell name, contained in the **ThisCell** file, converted to uppercase, becomes the local Kerberos realm name.

- **/etc/krb.realms**

  This file maps network interface (host) names to realms. When Kerberos needs to determine to which realm a network interface is assigned, it looks in this file. If there is no entry for the host name or for the domain part of the host name, the default is a realm name equal to the domain portion of the host name, converted to upper case. Any network interface names on systems using SP authenticated services which have a domain portion which is not the same as the local realm name, except for case, must have an entry in this file.

  Entries for all network interfaces on the SP nodes and on the control workstation that require this mapping in **/etc/krb.realms** are added automatically by either **setup_authent** initially or by the **setup_server** script. The file is kept identical on the control workstation and the nodes by the node installation customization process. There may be cases involving other work stations, however, where you will have to add entries to the file on one or more systems.

  When you set up a secondary authentication server or set up a system as an authentication client system, you are instructed to copy the configuration files from the primary server. If the new system requires mapping in **krb.realms**, the **setup_authent** script will create the entry in the local copy of the file. You will have to manually add the entry for the workstation you just configured into the control workstation's file and into the files on any other client and server workstations already installed.

Setting up the primary authentication server: The first system that you set up, when you install your PSSP software, is your primary authentication server. The primary authentication server may be your control workstation or it could be some other IBM RS/6000 workstation. You may want to have your servers entirely off of the SP system itself to provide greater physical security for them. You may want to allow logins to the control workstation that are not appropriate for the authentication servers.

The **setup_authent** script is used to configure your authentication services. If you have installed the **ssp.authent** file set on the workstation prior to running the script, it assumes that the system will be a PSSP authentication server. If you provide your own **krb.conf** file, the script will look for an entry for the local host name. If it cannot find one, it allows you to configure the system without any server (as an authentication client system).

The **setup_authent** script creates the service principals for local service instances and creates the server key file for them. It updates the **krb.realms** file as needed, and creates the root user's **.klogin** file to authorize the initial administrator principal to use remote commands.

The daemons which provide the authentication server (the KDC), and the database administration server, are started by adding entries to **/etc/inittab**.

The procedure to set up the primary authentication server is the following:

1. Install **ssp.clients** and **ssp.authent**.
2. Create a **/etc/krb.conf** file if you want or if your host names do not follow Kerberos convention.
3. Execute the **/usr/lpp/ssp/bin/setup_authent** program and follow the instructions in the prompts.

Setting up secondary authentication servers: The procedure to set up a secondary authentication server is the following:

1. Install **ssp.clients** and **ssp.authent**.
2. Copy the **/etc/krb.conf** file from the primary authentication server to this secondary server system.
3. Add a line to the **/etc/krb.conf** file listing this system as a secondary server for the local realm.
4. Copy the **/etc/krb.realms** file from the primary server to the secondary server system.
5. Execute the **/usr/lpp/ssp/bin/setup_authent** program and follow the instructions when prompted.
6. If **setup_authent** created an entry for the local **/etc/krb.realms** file, copy the file to the other systems.
7. Add an entry for the new secondary server to the **/etc/krb.conf** file on the systems on which you had previously-initialized authentication.
8. On the primary server, if this is the first secondary server, you should create a root **cron** entry that invokes the script **/usr/kerberos/etc/push-kprop** to periodically propagate database changes.

The procedure to set up your SP control workstation as a secondary authentication server is the same. For more information, see “Step 21.3: Initializing as a secondary Kerberos V4 authentication server” on page 41.

Setting up authentication client systems: The procedure to set up an authentication client system is the following:

1. Install **ssp.clients**.
2. Copy the **/etc/krb.conf** file from the primary authentication server to this system.
3. Copy the **/etc/krb.realms** file from the primary server to this system.

   Note: If the new workstation is outside the realm of the primary server, the new workstation needs to be added to the primary server’s **/etc/krb.realms** file first, before copying the **/etc/krb.realms** over to the new workstation. Otherwise, the next step, which runs the **setup_authent** program, will fail on the new client workstation.
4. Execute the **/usr/lpp/ssp/bin/setup_authent** program and follow the instructions when prompted.
5. If **setup_authent** created an entry for the local **/etc/krb.realms** file, copy the file to the other systems.

The procedure to set up your SP control workstation as an authentication client system is the same. For more information, see “Step 21.4: Initializing as an authentication client system” on page 42.

Directory path names for SP authentication services: Some of the directory paths that are used with authentication services are the following:

/usr/lpp/ssp/bin
Contains **setup_authent**

/usr/lpp/ssp/kerberos/bin
Contains commands for users of authentication services and authentication database administrators.

/usr/kerberos/bin
Contains a symbolic link for **/usr/lpp/ssp/kerberos/bin**

/usr/lpp/ssp/kerberos/etc
Contains authentication daemons and commands used by root to maintain a local SP authentication database.

/usr/kerberos/etc
Contains a symbolic link for **/usr/lpp/ssp/kerberos/etc**

/var/kerberos/database
Contains the SP authentication database on the server

/var/adm/SPlogs/kerberos
Contains error logs for the authentication servers

/etc Contains the files **krb.conf**, **krb.realms**, and **krb-srvtab**

/usr/lpp/ssp/rcmd/bin
Contains an **rsh** and **rcp** link to the AIX versions of the remote commands. On the SP, these versions support Kerberos V4 through an SP-supplied remote command library.

/usr/vice/etc
Contains AFS cell information and utilities

/usr/afsws/etc
Contains AFS executables

Configure Kerberos V4 security for each system partition

Notes:

1. When adding **k4** as an authentication method, you must ensure existing methods are also included. By not specifying a method, that method is removed. This will cause problems with configuration and the system.
2. All affected nodes must be set to **customize** to add Kerberos V4 authentication.
3. If the Kerberos V4 server will be external to the SP, it must be accessible and there must be **rsh** capability from the control workstation to the server system. You must issue the AIX **chauthent** command to allow for **rsh** access.

Perform the following steps to configure Kerberos V4 security for each system partition:

1. Kerberos V4 authentication must be set up on the control workstation before it can be selected for a system partition.

   If **setup_authent** was run previously, it is not necessary to run it again. Refer to “Step 21: Initialize RS/6000 SP Kerberos V4 (optional)” on page 38 for information on the different ways you can run **setup_authent**.
2. To indicate that Kerberos V4 security be installed and configured on the nodes, issue:

   ```
   spsetauth -p partition1 -i dce k4
   ```

The preceding example assumes that DCE was the current setting.

Note: If you only want to install and configure Kerberos V4 on the nodes, you should proceed to Step 7. This step does not enable the SP system to use Kerberos V4. You will need to continue with the remaining steps to enable Kerberos V4 usage.

3. To select Kerberos V4 as an authorization method for AIX remote commands, issue:

   ```
   spsetauth -p partition1 -d dce k4
   ```

   This step generates the necessary authorization files for each selected method and removes files or entries that are not needed. When adding **k4**, you will need to add it to the current setting. This implies that if **dce** was previously set, it must also be set now.
4. All affected nodes must be shut down. Use the **cshutdown** command (without the **-r** flag because the nodes should not be rebooted at this time).
5. To enable authentication methods for AIX remote commands, issue:

   ```
   chauthpar -c -p partition1 k5 k4
   ```

   Note: To enable Kerberos V4 for authenticated remote commands, but not for SP Trusted Services, continue to Step 7.
6. To enable authentication methods for SP Trusted Services, issue:

   ```
   chauthpts -c -p partition1 dce compat
   ```
7. To set all affected nodes to **customize**, issue:

   ```
   spbootins -r customize -s yes -l 1,3,5
   ```

   Note: In order to run the **spbootins -s yes** command, you must have SDR write authority and be authorized to perform a remote command to the target nodes.
8. Use the **cstartup** command to reboot all affected nodes.

Your system is now configured and enabled to use Kerberos V4, where appropriate, as an authentication method.

Note: The steps to add an authentication method basically work from (in order of SDR Syspar attributes) **auth_install**, **auth_root_rcmd**, **auth_methods**, **ts_auth_methods**. In order to remove a particular method, simply run the steps in reverse. For example, to remove **dce** as an authentication method for SP Trusted Services (and you have the settings **dce**.**compat**), issue:

```
chauthpts -p partition1 compat
```

This disables SP Trusted Services from using DCE as an authentication method. The steps for removing DCE or Kerberos V4 authentication are detailed in the “Changing the security configuration” chapter in *PSSP: Administration Guide*.

Enabling restricted root access (RRA)

In order to enable (activate) restricted root access (RRA), all nodes within the SP system must be at PSSP 3.2 or later. If any of your nodes are running a version of PSSP earlier than PSSP 3.2, RRA cannot be enabled. In addition, if any IBM Virtual Shared Disk adapters are defined, or if GPFS is installed, attempts to enable RRA will fail.

To complete this task, refer to:

- If you currently use nodes as boot/install servers, follow the directions in "Using multiple boot/install servers with RRA".
- "Step 61: Perform additional node customization" on page 90 and "Appendix E. User-supplied node customization scripts" on page 297
- "Step 74: Update authorization files in restricted mode for boot/install servers (optional)" on page 102

Using multiple boot/install servers with RRA

Using multiple boot/install servers in RRA is not recommended and is not automatically supported by PSSP. However, depending on the size of your system and network loads, it may not be possible to install your system with a single boot/install server.

Boot/install servers are NIM masters and, therefore, require **rsh** and **rcp** access to both the control workstation and to the nodes they serve. PSSP will not automatically create the correct entries in the authorization files to allow these commands to work.

To use additional boot/install servers, follow the following procedure to manually establish the correct authorizations on your system.

On the control workstation, the authorization files must have the following changes, depending on the setting of auth_root_rcmd:

standard
: An entry for the boot/install server node host name in **/.rhosts**

k4
: An entry for the boot/install server node rcmd principal in **/.klogin**

dce
: An entry for the self-host and the spbgroot principal for the boot/install server node

On the boot/install server node, you need to edit **/etc/sysctl.conf** to include the following entries:

- Include **/usr/lpp/ssp/sysctl/bin/install.cmds**
- Include **/usr/lpp/ssp/sysctl/bin/switch.cmds**
- Include **/usr/lpp/ssp/sysctl/bin/firstboot.cmds**, if initiating a node install customization

Enabling a secure remote command method

PSSP 3.4 gives you the ability to have the PSSP system management software use a secure remote command process in place of the AIX **rsh** and **rcp** commands when restricted root access is enabled. You can acquire, install, and configure any secure remote command software of your choice. With restricted root access and a secure remote command process enabled, the PSSP system management software has no dependency to internally issue **rsh** and **rcp** commands as a root user from the control workstation to nodes, from nodes to nodes, nor from nodes to the control workstation. Refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for additional planning information.

Note: A secure remote command method can be enabled only if all nodes are at PSSP 3.2 or later.

Steps to enable a secure remote command method include:

1. Acquire the appropriate secure remote command software. Install, configure, and start this software on your control workstation.
2. If your nodes are already installed, install, configure, and start the secure remote command software on the nodes. The sample provided in **/usr/lpp/ssp/samples/script.cust** should be reviewed. It provides an example of how to install and configure a particular secure remote command package.
3. Verify that secure remote commands are functioning properly from the control workstation to the nodes. There must be no password prompts.
4. Add an entry to **/etc/inittab** on both the control workstation and the nodes to automatically start the secure remote command daemon during reboot. It should be started immediately after **tcpip**.
5. Use the **spsitenv** command to enable PSSP use of the secure remote command method. Refer to “Step 30: Enter site environment information” on page 49.
6. For future installations, you must update the **script.cust** file to install the secure remote command software package. Refer to “Step 61: Perform additional node customization” on page 90 and “Appendix E. User-supplied node customization scripts” on page 297. Even if you choose not to install your secure remote command software package using **script.cust**, you must still modify your **script.cust** file to add an entry to **/etc/inittab** to start your secure remote command daemon. This insures that the secure remote command daemon starts early enough, even during node customization.

Configuring none for AIX remote command authorization

After you install the secure remote command software of your choice on the control workstation and enable the restricted root access and secure remote command process options, when setting your security configuration within each SP system partition, a choice of **none** is offered for AIX remote command authorization. If you choose **none**, the PSSP system management software does not generate any root entries in the **.klogin**, **.k5login**, or **.rhosts** files on the nodes in the partition. To be able to set AIX authorization for remote commands to **none** in any partition, PSSP 3.4 must be installed on all nodes in that partition.

Steps to configure **none** for AIX remote command authorization:

1. Use the **spsetauth** command to set **none** for a particular partition. Run the **spsetauth** command for each partition as needed. This command updates the root authorization files on the control workstation only and puts the setting in the SDR. For example:

   ```
   spsetauth -d -p partition_name none
   ```

2. Update the root authorization files on the nodes to remove any PSSP-generated entries. For example:

   ```
   dsh -avG /usr/lpp/ssp/bin/updauthfiles
   ```

Procedure for changing an SP system set up with dce:compat to dce only

This procedure can be used for any system partition whose security values match either system partition shown in the following example.

```
splstdata -p
```

You should receive output similar to the following:

```
List System Partition Information

System Partitions:
-------------------------------------------------
c186s
c186sp1

Syspar: c186s
-------------------------------------------------
syspar_name        c186s
...
auth_install       dce:k4
auth_root_rcmd     dce:k4
ts_auth_methods    dce:compat
auth_methods       k5:k4

Syspar: c186sp1
-------------------------------------------------
syspar_name        c186sp1
...
auth_install       dce:k4
auth_root_rcmd     dce:k4
ts_auth_methods    dce:compat
auth_methods       k5:k4
```

Note: This procedure can be used if auth_methods contains **std**. However, you will need to include **std** as an operand on the **chauthpar** and **chauthent** commands.

You should perform the following steps. In this example, **k4** and **compat** are being removed from the c186sp1 partition.

1. Ensure that there is an SP administrative principal defined to DCE. “Step 22.3: Create SP administrative principals” on page 45 describes how to create this principal in DCE. Ensure that the administrative principal is in the **/.k5login** file for all of the nodes. In the examples that follow, ″atest″ was used as the principal and ″atest″ appears in **/.k5login** as ″atest@fvtdcecell″.
2. On the control workstation, login to the administrative principal and ensure that it is working correctly.

   ```
   dce_login atest
   Enter Password:
   dsh -avG date
   ```

   If all of the nodes respond with the date, everything is defined correctly.
3. Remove **compat** from SP Trusted Services as follows:

If using:	Do this:

SMIT	**TYPE** **smit spauth_config**
	• The RS/6000 SP Security menu appears
	SELECT Enable Authentication Methods for SP Trusted Services
	• The Enable Authentication Methods for SP Trusted Services menu appears.
	`[Entry Fields]`
	`Enable on Control Workstation Only no`
	`Force change on nodes no`
	You cannot select YES for both of the previous entries.
	`System Partition Name c186s`
	`Authentication Methods dce`

chauthpts	Issue the following command: `chauthpts -p c186sp1 dce`

Verify that everything is correct. First, check the local SP Trusted Services security setting on each node, then check that the high availability daemons (**haem** specifically) have responded correctly. The security settings should no longer contain **compat**. Issue:

```
dsh -avG lsauthts
```

You should receive output similar to the following:

```
c186n01,ppd.pok.ibm.com: DCE
...
```

Issue:

```
dsh -avG "lssrc -ls haem | grep secure"
```

You should receive output similar to the following:

```
c186n01,ppd.pok.ibm.com: Daemon security:            DCE
...
```

If **k4** was removed from **all** system partitions, **chauthpts** will set the local SP Trusted Services setting on the control workstation to include just DCE. To verify, issue:

```
lsauthts
```

You should receive output similar to the following:

```
DCE
```

4. Remove **k4** from the authentication methods for AIX remote commands.

If using:	Do this:
SMIT	**TYPE** **smit spauth_config** • The RS/6000 SP Security menu appears **SELECT** Enable Authentication Methods for AIX Remote Commands • The Enable Authentication Methods for AIX Remote Commands menu appears. [Entry Fields] Enable on Control Workstation Only no Force change on nodes no You cannot select YES for both of the previous entries. System Partition Name c186s Authentication Methods k5
chauthpar	Issue the following command: `chauthpar -p c186sp1 k5`

To verify, issue:

```
dsh -avG lsauthent
c186n01.ppd.pok.ibm.com: Kerberos 5
c186n02.ppd.pok.ibm.com: Kerberos 5
...
```

5. To verify your system partition security setting, enter:

   ```
   splstdata -p
   ```

   You should receive output similar to the following:

   ```
   List System Partition Information

   System Partitions:
   -------------------------------------------------
   c186s
   c186sp1

   Syspar: c186s
   -------------------------------------------------
   syspar_name        c186s
   ...
   auth_install       dce:k4
   auth_root_rcmd     dce:k4
   ts_auth_methods    dce
   auth_methods       k5

   Syspar: c186sp1
   -------------------------------------------------
   syspar_name        c186sp1
   ...
   auth_install       dce
   auth_root_rcmd     dce
   ts_auth_methods    dce
   auth_methods       k5
   ```

6. After Kerberos 4 has been removed from all of the system partitions on the SP, it should be removed from the local settings on the control workstation. Issue:

   ```
   lsauthent
   ```

   You should receive output similar to the following:

```
Kerberos 5
Kerberos 4
```

Issue:

```
chauthent -k5
lsauthent
```

You should receive output similar to the following:

```
Kerberos 5
```

7. Once **k4** is removed from **all** of the system partitions on the SP, the Kerberos daemons can be stopped and removed from **inittab**. To determine the daemon names, issue:

   ```
   lssrc -a | grep k
   ```

   You should receive output similar to the following:

   ```
   kerberos                    15872     active
   kadmind                     11558     active
   ```

   To stop the daemons, issue:

   ```
   stopsrc -s Kerberos
   stopsrc -s kadmind
   ```

   To remove the daemons from **inittab**, issue:

   ```
   rmitab kerb
   rmitab kadmind
   ```

   The Kerberos configuration files must also be removed or renamed. These files are:

   krb-srvtab

   krb-srvtab.save

   krb.realms

   Note: These files should be saved until you are absolutely certain you will not be going back to **k4**.

   Removing the **krb-srvtab** files is particularly important because the **setup_server** command uses it as an indicator that Kerberos V4 processing is required. If the **krb-srvtab** file is present, but the Kerberos daemons are not running, **setup_server** will fail.

Chapter 6. Reconfiguring the RS/6000 SP system

This chapter explains the tasks necessary to reconfigure your RS/6000 SP system. It provides the SP commands that you should use to add, delete, and modify your hardware.

Notes:

1. It is very important that you consult *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* prior to reconfiguring your system to understand the implications of adding, deleting, or modifying any hardware.
2. If you are adding any hardware to your system, you must ensure that your control workstation is at the highest level of AIX and PSSP that will be running on the rest of your system.
3. To learn more about the implications of changing IP addresses and host names, refer to the appendix titled “IP address and host name changes for SP systems” in the *PSSP: Administration Guide*.
4. For more information about configuring LoadLeveler, refer to *IBM LoadLeveler for AIX 5L: Using and Administering*.
5. To perform these tasks, you will need write access to the SDR and administrative and control access to the affected frame and slot objects in the hardware monitor. Refer to *PSSP: Administration Guide* and *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for additional information.

Adding a frame, SP-attached server, or clustered enterprise server

When you add a frame, SP-attached server, or a clustered enterprise server to your RS/6000 SP system, you should plan how it fits into your network configuration. Consider the number of new nodes and how many interfaces they will have when planning your network configuration. Record this information in your SP configuration worksheets, located in the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*. Both the *RS/6000 SP: Planning, Volume 1, Hardware and Physical Environment* and the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* offer information to help you with these decisions. For clustered enterprise servers, there are also planning issues that you must consider if you will be adding an SP frame or switch sometime in the future.

IBM suggests that you add frames only to the end of your system, otherwise you may have to reconfigure the System Data Repository (SDR).

After you plan the configuration, follow these steps to add the frame and its nodes to your RS/6000 SP system, referring to your worksheets as necessary. Many steps include a reference to a previous chapter for more detailed information on the steps that you will be performing.

Frames that will never contain nodes, such as SP expansion I/O units or switch-only frames, can be added with frame numbers greater than 128.

See the section on using RS/6000 SP authentication services in the *PSSP: Administration Guide* for more information.

Step 1: Archive the SDR

Before reconfiguring your system, you should back up the SDR by issuing:

```
SDRArchive
```

Note the location and the name of the file created after you issue this command.

Step 2: Unpartition your system

If your existing system has multiple partitions defined and you want to add a frame that has a switch, you need to bring the system down to one partition before you can add the additional frame.

Step 2.1: Repartition your system to a single system partition

See the "Managing system partitions" chapter in the *PSSP: Administration Guide* for instructions on partitioning your SP system.

Step 3: Connect frames to your control workstation

Connect RS-232 and Ethernet cables from the SP system frames and from the SP-attached servers or clustered enterprise servers to the control workstation according to your SP Control Workstation Network Worksheet. Your IBM Customer Engineer (CE) performs this step. The CE will need access to the control workstation to run diagnostics. See *RS/6000 SP: Installation and Relocation* for instructions.

Step 4: Configure RS-232 control lines

You must allocate the appropriate number of tty ports in this step. Each SP frame in your system requires a serial port on the control workstation configured to accommodate the RS-232 line. SP-attached servers require two serial ports. For example, if you have two SP frames and one SP-attached server, configure four tty terminals. Refer to "Step 1.4: Verify the control workstation serial ports" on page 13 for valid tty port values.

Enter a command similar to the following to define and configure an RS-232 line for parent adapter sa0 on serial port 1:

```
mkdev -c tty -t tty -s rs232 -p sa0 -w 1
```

The following example configures a second port of a two-frame system:

```
mkdev -c tty -t tty -s rs232 -p sa1 -w 2
```

Step 5: Configure the Ethernet adapter (optional)

Use SMIT or the **chdev** command to configure each Ethernet adapter connecting the nodes on your frames to the control workstation. For example, enter a command similar to the following to configure an SP Ethernet administrative LAN adapter:

```
chdev -l en0 -a netaddr=129.33.41.1 -a netmask=255.255.255.0 -a state=up
```

For details on the correct use of **chdev**, see the *IBM AIX Commands Reference*, the man pages, or the online information database.

Refer to your SP Control Workstation Network Worksheet in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*.

If the adapter is not yet defined or configured, use **smit mkinet** or the **mkdev** command instead of **smit chinet** or **chdev** to specify a new IP host name and

netmask values. For example, enter a command similar to the following to define and configure an SP Ethernet administrative LAN adapter:

```
mkdev -c if -s EN -t en -a netaddr=129.33.34.1 \
-a netmask=255.255.255.0 -a state=up -q -w en0
```

If you are adding an extension node to your system, you may want to configure the adapters now. For more information, refer to “Chapter 10. Installing extension nodes” on page 281.

Step 5.1: Verify the control workstation interfaces

Verify the configuration for each Ethernet adapter in the control workstation. You can verify that the adapter is installed even if it is not cabled to the SP system yet.

Verify each Ethernet adapter by *pinging* its IP address and seeing if you get a proper response. If you do not receive a response, debug the network problem, and reconfigure the adapter.

For example:

```
ping -c 1 129.33.34.1
```

Step 6: Enter Hardware Management Console (HMC) information (optional)

You must perform this step if your SP system or clustered enterprise server system will contain IBM @server pSeries 690 servers.

The hardware control and monitor functions for the pSeries 690 server are managed through a network connection from the control workstation to the hardware management console (HMC) that is controlling the pSeries 690 server. Install the HMC for your pSeries 690 server following the instructions in *pSeries 690 Installation Guide*. The following operations must be performed locally on the HMC before the pSeries 690 server is defined to PSSP. Refer to *Hardware Management Console for pSeries Operations Guide* for details on performing these operations. Review the *READ THIS FIRST* document that accompanies the PSSP installation media for information on required HMC product and PTF levels and the corresponding pSeries 690 hardware and software product and PTF levels.

Perform the following steps for each HMC and each pSeries 690 server:

1. Ensure that the HMC is installed and configured to operate on the SP Ethernet administrative LAN network. Use the HMC *System Configuration* interface to customize the network settings. Ensure the netmask is properly assigned and that all IP addresses have been registered with your name server and can be resolved. Note the IP address assigned to this HMC SP Ethernet administrative LAN connection. This information will be required when defining the HMC to the control workstation in the next section.
2. Use the HMC *User Management* interface to define a user ID with the role of *System Administrator* and assign a password. This information will be required when defining the HMC to the control workstation in the next section.
3. Ensure that the pSeries 690 server is recognized by the HMC. Use the HMC *Partition Management* interface to determine if the server is present. Follow problem resolution procedures in *Hardware Management Console for pSeries Operations Guide* if an entry for the server is not displayed on the interface.
4. Use the HMC *Partition Management* interface to view the properties for the managed system object. If desired, change the system name from the default name set by the HMC. Note the defined system name. This information will be

required when entering the non-SP frame information for this server in "Step 8: Enter non-SP frame information and reinitialize the SDR (optional)" on page 200. If the system name is changed in the future, the new name will then also need to be changed in the non-SP frame information stored in the SDR on the control workstation.

5. Use the HMC *Partition Management* interface to select the desired power-on mode for your system: full system partition (SMP) mode, logical partition standby mode, or physical partition mode.
6. If you selected logical partition standby mode or physical partition mode, use the HMC *Partition Management* interface to create partitions and profiles as necessary. Partition objects must be created at this time, but the partitions do not need to be activated or installed before proceeding with control workstation operations. When a partition is activated from the control workstation, the default profile for the partition will be used. If you want to use a different profile, change the default profile setting for the partition using the HMC *Partition Management* interface.
7. View the properties for each partition object and note the partition ID. Each partition is represented in PSSP as a node. The partition ID will be used by PSSP to assign a corresponding SP slot number and node number to the SP node for that partition.

The following operations must be performed on the control workstation before entering non-SP frame information for your pSeries 690 servers.

1. Use the AIX **ping** command to verify that the control workstation has network connectivity to each HMC:

   ```
   ping hmc_hostname_or_ipaddr
   ```

   where *hmc_hostname_or_ipaddr* is the host name or IP address of the HMC as configured on the HMC *System Configuration* interface in Step 1 on page 197 of the previous section. If the command fails, review your network and name server configurations on both the control workstation and the HMC.
2. Define the previously-created HMC user ID to PSSP for **hardmon**. Running the following command once for each HMC:

   ```
   /usr/lpp/ssp/bin/sphmcid hmc_hostname_or_ipaddr hmc_sysadmin_userid
   ```

   where *hmc_hostname_or_ipaddr* is the host name or IP address of the HMC as configured on the HMC *System Configuration* interface in Step 1 on page 197 of the previous section and *hmc_sysadmin_userid* is the system administrator user ID created on the HMC *User Management* interface in Step 2 on page 197 of the previous section. You will be prompted to enter the password for the *hmc_sysadmin_userid*. You must run this command again anytime the password is changed for the *hmc_sysadmin_userid*.
3. Define the switch node numbers for your nodes. If this system has an SP Switch or is a switchless SP system, you must define an SP switch node number for each logical partition in the pSeries 690 server. Manually edit the **/etc/switch.info** file to include one entry for each logical partition in the attached server. See *PSSP: Command and Technical Reference* for details on editing this file. You can skip this step if the system has an SP Switch2 or is a switchless clustered enterprise server system.

   An example **/etc/switch.info** file might contain the following entries for a pSeries 690 server that will be defined as frame 5, with four LPARs, attached to switch 2 in the system:

```
# Node_number     Switch_node_number
65                16
66                17
67                18
68                19
```

If you are running your pSeries 690 server in SMP mode or will only be defining one LPAR and it is assigned partition ID 1, you can skip this operation and simply enter the switch node number when you enter the the other non-SP frame information in “Step 8: Enter non-SP frame information and reinitialize the SDR (optional)” on page 200.

Step 7: Enter SP or multiple NSB frame information and reinitialize the SDR

You must perform this step at least once for each set of frames or multiple node switch board (NSB) frames that you are adding to the system. You do not need to reinitialize the SDR until you are entering the last set of frames.

SP frames containing nodes must be numbered between 1 and 128 inclusive. This is to ensure that nodes will be numbered between 1 and 2047. Larger frame numbers, up to 250, can be used for frames that will contain only switches or SP expansion I/O units.

SP frames

This step creates frame objects in the SDR for each SP frame in your system. At the end of this step, the SDR is reinitialized, resulting in the creation of node objects for each node attached to your frames.

Enter information about your frames using Perspectives, SMIT, or the **spframe** command. You must be an authenticated administrative user to issue this command.

Specify the **spframe** command with **-r yes** to reinitialize the SDR (when running the command for the final series of frames), a starting frame number, a frame count, and the starting frame’s tty port. The following example enters information for four frames (frame 1 to frame 4) and indicates that frame 1 is connected to **/dev/tty0**, frame 2 to **/dev/tty1**, and so on, and reinitializes the SDR:

```
spframe -r yes 1 4 /dev/tty0
```

If frames are not contiguously numbered, repeat this step for each series of contiguous frames. To save time, do not specify reinitialization of the SDR until you are entering the final series of contiguous frames.

Multiple node switch board (NSB) frames (SP Switch2 only)

In PSSP 3.4, you can install multiple NSBs in an SP frame. A multiple NSB frame can contain switches only in slots 1 through 16. You cannot install SP nodes in a multiple NSB frame.

You can enter information for non-SP frames using Perspectives, SMIT, or the **spframe** command. If frames or tty ports are not all contiguously numbered, repeat this step for each series of contiguous information. To save time, do not specify the reinitialization of the SDR until you are entering the final series of contiguous frames.

Specify **spframe** command with **-r yes** to reinitialize the SDR (when running the command for final series of frames), a starting frame number, a frame count, and the starting frame's tty port.

The following example enters information for two frames (frame 1 to frame 2) and indicates that frame 1 is connected to **/dev/tty0**, frame 2 to **/dev/tty1**, and reinitializes the SDR.

```
spframe -r yes -m 1 2 /dev/tty0
```

If frames are not contiguously numbered, repeat this step for each series of contiguous frames. To save time, do not specify reinitialization of the SDR until you are entering the final series of contiguous frames.

Step 8: Enter non-SP frame information and reinitialize the SDR (optional)

If you entered SP or multiple NSB frame information in "Step 7: Enter SP or multiple NSB frame information and reinitialize the SDR" on page 199, you must reinitialize the SDR before continuing to enter frame information for non-SP frames. You must perform this step at least once for each frame protocol of non-SP frames that you are adding to the system.

SP-attached servers and clustered enterprise servers also require frame objects in the SDR. These frames are referred to as non-SP frames and one object is required for each server attached to your SP. These objects have a non-SP hardware protocol associated with them which instructs PSSP as to which method of hardware communications is to be used for controlling and monitoring the node associated with this frame object. Valid hardware protocol values of the nodes within the frame are:

HMC IBM @server pSeries 690 servers

CSP RS/6000 H80, M80, and IBM @server pSeries 660 servers (6H0, 6H1, 6M1)

SAMI RS/6000 S70, S7A, and S80 or IBM @server pSeries 680 servers

The number of tty port values you must define depends on the hardware protocol type you selected.

HMC Does not require a tty port value, but the HMC must be connected by the SP Ethernet administrative LAN

CSP Requires one tty port value

SAMI Requires two tty port values

The servers that use the SAMI hardware protocol require two tty port values to define the tty ports on the control workstation to which the serial cables connected to the server are attached. The tty port value defines the serial connection to the operator panel on these servers for hardware controls. The s1 tty port value defines the connection to the serial port on the servers for serial terminal (s1term) support.

Switch port numbers are required on SP Switch or switchless systems for each SP-attached server in your system. This information is available from your Switch Configuration Worksheet. Although switch ports are not required for switchless or SP Switch2 clustered enterprise servers, you may want to specify a switch port if you plan to add an SP frame sometime in the future. *RS/6000 SP: Planning,*

Volume 2, Control Workstation and Software Environment explains how to fill out your worksheet and provides details on assigning switch port numbers.

For pSeries 690 servers in an SP Switch or switchless system, a switch node number is required for each logical partition (LPAR). These switch node numbers must be specified to PSSP through the **/etc/switch.info** file. Manually edit the **/etc/switch.info** file to include one entry for each LPAR in the attached server. See the **switch.info** file in *PSSP: Command and Technical Reference* for details on editing this file.

An example of the **/etc/switch.info** file might contain the following entries for a pSeries 690 server that will be defined as frame 5 with four LPARs attached to switch 2 in the system:

```
# Node_number        Switch_node_number
65                          16
66                          17
67                          18
68                          19
```

If you are running your pSeries 690 server in full system partition (SMP) mode or will only be defining one LPAR and it is assigned partition ID 1, you can skip this operation. Instead just simply enter the switch node number when you enter the other non-SP frame information later in this step through the SMIT menu or the **spframe** command.

You can enter information for non-SP frames using Perspectives, SMIT, or the **spframe** command. If using SMIT, a different procedure is used for each hardware protocol. If frames, tty ports, or switch port values are not all contiguously numbered, repeat this step for each series of contiguous information. To save time, do not specify the reinitialization of the SDR until you are entering the final series of contiguous frames.

Specify the **spframe** command with the **-n** option for each series of contiguous non-SP frames. The **-n** option is not required for switchless clustered enterprise servers or SP Switch2 systems. Specify the **-r yes** option when running the command for the final series of frames. Include the starting frame number, the number of frames, the starting tty port value, and the starting switch port number for each invocation of the command.

The following example enters non-SP information for one S80 server (frame 5), one H80 server (frame 6), and one pSeries 690 server with four LPARs (frame 7).

The first server has the following characteristics:

```
Frame Number:        5
tty port for operator panel connection:  /dev/tty4
tty port for serial terminal connection: /dev/tty5
switch port number: 10
```

The second server has the following characteristics:

```
Frame Number:        6
tty port for operator panel connection:  /dev/tty6
switch port number: 11
```

The third server has the following characteristics:

```
Frame Number:        7
switch port number: 12, 13, 14, 15
```

To define the first two servers to PSSP, enter:

```
spframe -r no -p SAMI -n 10 -s /dev/tty5 5 1 /dev/tty4

spframe -r no -p CSP -n 11 6 1 /dev/tty6
```

Append the following to the **/etc/switch.info** file:

```
7,1 12
7,2 13
7,3 14
7,4 15
```

To define the third server to PSSP and reinitialize the SDR, enter:

```
spframe -r yes -p HSC -d huntley -i 129.33.32.121 7
```

Note: The SP-attached server and clustered enterprise server in your system will be represented with the node number corresponding to the frame defined in this step. For pSeries 690, each logical partition in the server will be represented as a node on that frame. PSSP will assign an SP slot number and node number that corresponds to the partition ID set by the HMC for that partition. Continue with the remaining installation steps to install the SP-attached server, clustered enterprise server, or logical partition as an SP node.

Step 9: Verify frame information

All frames must be powered up and connected to the control workstation so that the nodes are automatically detected and added to the SDR.

You should see the SP frames represented with thin, wide, or high nodes, depending on your configuration. If using Perspectives, SP-attached servers are shown as a unique SP-attached server icon. If using **spmon -d**, SP-attached servers are represented as a one node frame. The pSeries 690 servers will be represented as one frame with one node per LPAR within that frame. For multiple node switch board frames and for intermediate switch board frames, you should see the SP frames represented with switches listed in the appropriate slot locations. If your frames are not correctly represented, you may have a hardware problem, such as a misplugged RS-232 cable. See the “Diagnosing hardware and software problems” chapter in *PSSP: Diagnosis Guide* for help in correcting the error. If an error occurred, the frame must be deleted, using the **spdelfram** command, prior to reissuing the **spframe** command. After updating the RS-232 connection to the frame, you should reissue the **spframe** command.

Step 10: Add nodes

Proceed to the “Adding nodes” section to complete the installation of your system.

Adding nodes

When you add a node, SP-attached server, or clustered enterprise server to your SP system, you should plan how it fits into your configuration. Consider the number of interfaces it will contain when planning your network configuration. Record this information in your SP configuration worksheets located in the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*. You also must consider which nodes provide their boot/install service. You can configure a new boot/install server for these nodes or add them to an existing boot/install server as clients.

For more information on adding an extension node, refer to “Chapter 10. Installing extension nodes” on page 281.

Notes:

1. When you add a node, ensure that the switch adapters are compatible with the switch adapters in the rest of your system and ensure that the nodes you are adding are supported by the switch in the system. For example, the SP Switch-8 only handles eight nodes.
2. You cannot add a node in a location already defined for an SP Switch Router Adapter node or an SP-attached server. View the System Partition Map to determine which slots are valid for adding nodes.
3. Do these steps in the system partition to which you add the nodes (SP Switch only).
4. **To SP Switch-8 Users:**

   Since the switch node numbers for these switches are computed sequentially, it is possible for new nodes to be incorrectly numbered when a group of nodes is added. You can avoid this problem by having your IBM Customer Engineer (CE) connect the nodes one at a time in an ascending fashion. Between each node, you should verify that a node object has been created with the proper switch node number.
5. If you are adding a new LPAR to an existing pSeries 690 server attached to your SP or clustered enterprise server system, perform the steps in “Adding a new LPAR” on page 245 then continue with the steps in this section.

Step 1: Archive the SDR

Note: Perform this step only if you did not back up the SDR in “Step 1: Archive the SDR” on page 196.

Before reconfiguring your system, you should back up the SDR by issuing:

```
SDRArchive
```

Note the location and the name of the file created after you issue this command.

Step 2: Connect new nodes to the frame

Do this step if you are adding nodes to an existing frame. Your IBM Customer Engineer (CE) performs this step. See *RS/6000 SP: Installation and Relocation* for instructions.

When adding nodes with connected SP expansion I/O units, the I/O units must be powered on before the nodes are powered on in order for the connections to be properly recognized.

To add a new logical partition to an existing pSeries 690 attached server or clustered enterprise server, you will need to use the HMC interface to create the logical partition and assign resources to that partition. Once the partition has been created, **hardmon** will automatically find the new LPAR and a new node will be added to the SDR for that pSeries 690 frame. To add a new LPAR, perform the following steps:

1. Start the HMC interface. This can be done in one of the following ways:
   - Login directly to the HMC subsystem. This will automatically start the GUI on the local HMC display.

- Manage the HMC subsystem through a remote WebSM session running on your control workstation or other UNIX workstation that has connectivity to the HMC.
- Launch a remote WebSM session to the HMC from the Hardware Perspective running on your control workstation.

2. Select the pSeries 690 system object and follow HMC procedures for creating a new logical partition. Refer to the *Hardware Management Console for pSeries Operations Guide* for details on performing these operations. If this logical partition requires resources that are currently owned by other partitions in the server, follow the steps in "Deleting Resources from an LPAR" on page 246 before proceeding with this step.
3. View the properties for the LPAR and note the partition ID assigned by the HMC.

hardmon will automatically find the new LPAR and a new node object will be created in the SDR on the control workstation. The frame slot number of the LPAR is equivalent to the partition ID assigned to that LPAR by the HMC. This slot number in turn is used to derive the node number following standard PSSP frame-slot number to node number conversion algorithms.

Step 3: Update the state of the supervisor microcode

To ensure that you have the latest level of microcode required by the hardware on your SP system, issue the **spsvrmgr** command. For example, to get the status in report form of all of your frames, nodes, and switches, enter:

```
spsvrmgr -G -r status all
```

To update the microcode of the frame supervisor of frame 3, enter:

```
spsvrmgr -G -u 3:0
```

Note: When using the **spled** application, a node in the process of having the microcode on its supervisor card updated will not be displayed in the window.

Refer to the *PSSP: Command and Technical Reference* for more information on using the **spsvrmgr** command.

Step 4: Enter the required node information

- Verify that the node objects have been created by issuing **splstdata -n** and verify that there is an entry for each node in your system.
- Be sure to have your node configuration worksheet on hand with all the node information completed before attempting to perform this step. *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* explains how to fill out your worksheet.
- If multiple IP interfaces map to the same host name on the starting node, you must enter the Ethernet IP address for the starting node. Do not enter its host name.

  If multiple IP interfaces do not map to the same host name on the starting node and you decide to enter its host name, it must be identical to the default host name returned by the **host** command for the starting node SP Ethernet IP address. For example, if the SP Ethernet administrative LAN adapter IP address of a node is 123.45.678.90 and host 123.45.678.90 gives v64n90.xen.kry.arg.com, then this host name must be used.

- The host name of a node is case sensitive. If you choose to enter the host name for a node, it must match the format of the host name returned when you issue **/usr/bin/host** against the node's IP address.
- Enter a correct value for Ethernet speed (10, 100, or auto), Duplex (full, half, or auto), and Type (bnc, dix, tp, or NA) for the SP Ethernet adapter on each node.
- When adding nodes with connected SP expansion I/O units, verify that:
  - The node expansion objects were created by issuing **splstdata -x**
  - An entry exists for each I/O unit in your system
  - The node connection information is correct (see "Step 23: Verify the SP expansion I/O unit configuration" on page 223)
- If a node is not directly connected to an SP Ethernet adapter on the control workstation or the host name of the control workstation is not set to the name of that SP Ethernet adapter, the default route for the node must be an adapter that is automatically configured. See the **spadaptrs** command in *PSSP: Command and Technical Reference* for a list of adapter types that can be automatically configured.
- For pSeries 690 servers, specifying the SP Ethernet adapter by its physical location code is suggested, especially if there is more than one Ethernet adapter present in the node. The physical location code for an adapter can be determined in one of the following ways:
  - Run the **spadapter_loc** command for the node. The command will return a list of all the SP supported adapters installed on the node. Save the results of the command in a file. You can use the returned hardware Ethernet address for your SP Ethernet adapter in "Step 5: Acquire the hardware Ethernet addresses" on page 206 to save processing time later.
  - Visually locate the adapter on your server and look up the physical location code in the hardware publications distributed with the server.
  - If AIX is installed and running on the server, run the AIX **lscfg -p** command to list the physical location codes of all the hardware devices installed on your system.

This step adds IP address-related information to the node objects in the SDR. It also creates adapter objects in the SDR for the SP Ethernet administrative LAN adapters on your nodes. This information is used during node customization and configuration.

Note: The default route that you enter in this step is not the same as the default route on the node. The route that you enter here goes in the SDR Node Class. It is the route over which the node communicates with its boot/install server (for example, install, customize, and so on). The default route must be a valid path from the SP Ethernet administrative LAN adapter to the node's boot/install server and the control workstation.

The default route on the node is the route it will use for its network communications if there is no specific route to the destination. During the boot process, this is set to the default route in the SDR. It can be changed later on in the boot process or after the node is running, but should not be changed permanently in the SDR. For FDDI, token ring, or other Ethernet adapters, create the route in **firstboot.cust**. The following example defines a route for an Ethernet adapter. This example also saves the route into the node's ODM.

```
old_route_info=$($lsattr -E -l inet0 | $grep 'route *net,.*,
0, 0-9 . *' | $awk ' print $2; ' | $tail -n 1) #-
if -n "$old_route_info" ; then #-
```

```
$chdev -l inet0 -a delroute="$old_route_info" > /dev/null
2>&1 #-
fi #-
$chdev -l inet0 -a route="0,<route>"
```

In order for the route to remain set after customization, also set the route up in **/etc/inittab** after the line that runs **rc.sp**. For the switch, set the route up in **/etc/inittab** after the line that runs **rc.switch**.

Enter information about your nodes attached to each Ethernet adapter using Perspectives, SMIT, or the **spadaptrs** command.

The following example configures an SP Ethernet administrative LAN adapter network of 16 nodes with IP addresses ranging from 129.33.32.1 to 129.33.32.16, a netmask of 255.255.255.192, and a default route of 129.33.32.200 for a twisted-pair Ethernet using auto-negotiate for the communication transfer and rate.

```
spadaptrs -e 129.33.32.200 -t tp -d auto -f auto 1 1 16 en0
          129.33.32.1 255.255.255.192
```

The following example configures the adapter on the SP Ethernet administrative LAN adapter for the first logical partition of a pSeries 690 server. The adapter is a twisted pair Ethernet adapter with communication transfer and rate set to auto-negotiate. The IP address is 129.33.32.65 with a netmask of 255.255.255.192. The pSeries 690 server is represented as frame 5, the node is assigned slot 1, and the adapter is located at the physical location U1.9-P2-I2/E1.

```
spadaptrs -P U1.9-P2-I2/E1 -t tp -d auto -f auto 5 1 1 en
          129.33.32.65 255.255.255.192
```

If you are adding an extension node to your system, you may want to enter required node information now. For more information, refer to “Chapter 10. Installing extension nodes” on page 281.

Step 5: Acquire the hardware Ethernet addresses

- Do not do this step on a production running system because it shuts down the nodes.
- Select only the new nodes you are adding. All the nodes you select are powered off and back on.
- The nodes for which you are obtaining Ethernet addresses must be physically powered on when you perform this step. No ttys can be open in write mode.

This step gets hardware Ethernet addresses for the SP Ethernet administrative LAN adapters for your nodes, either from a file or from the nodes themselves, and puts them into the Node Objects in the SDR. That information is used to set up the **/etc/bootptab** files for your boot/install servers. This step will also **ping** the default route set for this node.

If you know the hardware Ethernet addresses, you can speed this process by putting the addresses in the **/etc/bootptab.info** file. If you are performing this step for a pSeries 690 server, you may already have the hardware Ethernet addresses available to you from “Step 4: Enter the required node information” on page 204. Create the **/etc/bootptab.info** file as follows:

- Create a file named **/etc/bootptab.info** (if it does not already exist), listing your RS/6000 SP nodes by node number, followed by a blank and the hardware Ethernet address. For example, a file containing addresses for a frame might look like this:

```
1 08005ABAB177
3 08005ABAAEAB
5 08005ABAB161
7 08005ABAB17A
9 02608CF53067
13 02608CF527F2
17 08005ABAB1A0
19 08005ABAB062
21 002035D34F7A
22 002035D34FE2
23 002035D34F3C
24 002035D34F70
25 002035D34E65
26 002035D34E5F
27 002035D34FE5
28 002035D34F68
29 02608CF55E6D
```

Note: If you are using a **bootptab.info** file, you must only place nodes in the current system partition in the file. If you have multiple system partitions, you must update the **bootptab.info** file and run **sphrdwrad** for each system partition.

The **/etc/bootptab.info** file is not required. If you do not know your hardware Ethernet addresses, and the **/etc/bootptab.info** file does not exist, use **sphrdwrad** to access the SP node and retrieve the hardware Ethernet address for you. (This makes **sphrdwrad** take longer to run.)

The following example gets all hardware Ethernet addresses for an RS/6000 SP system:

```
sphrdwrad 1 1 rest
```

This example gets all hardware Ethernet addresses for the nodes specified in the node list (the **-l** flag):

```
sphrdwrad -l 10,12,17
```

If this step fails, look for the node conditioning instructions in the *PSSP: Diagnosis Guide*.

Step 6: Verify that the Ethernet addresses were acquired

This step verifies that Ethernet addresses were placed in the SDR node object.

Attention: If your system is large, **splstdata** returns great quantities of data. You may want to pipe the command output through a filter to organize the amount of data you see.

To display SDR boot/install data, enter:

```
splstdata -b
```

Step 7: Configure additional adapters for nodes

Perform this step if you have a switch or if you require any additional adapters.

Be sure to have your switch configuration worksheet on hand with all the switch information completed before attempting to perform this step. *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* explains how to fill out your worksheet.

This step creates adapter objects in the SDR for each node. The data in the adapter objects is used during the customization or installation steps to configure the adapters on the nodes. You can configure the following adapters with this procedure:

- Ethernet (en)
- FDDI (fi)
- Token ring (tr)
- css

To configure adapters such as ESCON and PCA, you must configure the adapter manually on each node using **dsh**, or modify the **firstboot.cust** file.

Configuring the switch adapters

To configure your switch adapters for use with the RS/6000 SP system, use SMIT or issue the **spadaptrs** command. *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* contains additional information on IP addressing for the switch.

The following example adds SDR information for a **css** (SP Switch or SP Switch2) network of 30 nodes (frame 1 slot 1 to frame 2 slot 16, with a wide node as the first node in each frame and the rest thin nodes, and a switch on each frame) with IP addresses from 129.33.34.1 to 129.33.34.30, and a netmask of 255.255.255.0. The IP addressing corresponds to the slots in the frame, with each wide node incrementing by 2 and each thin node incrementing by 1, and each high node by 4.

If you specify the **-s** flag to skip IP addresses when you are setting the **css** switch addresses, you must also specify **-n no** to not use switch numbers for IP address assignment, and **-a yes** to use ARP.

```
spadaptrs -s yes -n no -a yes 1 1 30 css0 129.33.34.1 255.255.255.0
```

Configuring other additional adapters

To configure other additional adapters, for example Ethernet (en), token ring (tr), or FDDI (fi), you must select the `Additional Adapter Database Information`. For these adapters, you can select either the Start Frame, Start Slot, and Node Count fields, or the Node List field.

Notes:

1. When using the token ring (tr) adapter, you must select the token ring rate (4 MB or 16 MB).
2. For best results, exit and get back into this panel for each different type of adapter. This clears any extraneous values left behind in the panel.
3. Enter a correct value for Ethernet speed (10, 100, 1000, or auto), Duplex (full, half, or auto), and Type (bnc, dix, tp, fiber, or NA) for every Ethernet adapter on each node.
4. For pSeries 690 servers, specifying the adapter by its physical location code and adapter type is suggested, especially if there is more than one adapter of that type present in the node. The physical location code for an adapter can be determined in one of the following ways:
   - Run the **spadapter_loc** command for the node. The command will return a list of all of the SP supported adapters installed on the node. If you ran this command as part of “Step 4: Enter the required node information” on page 204 and saved the results of the command in a file, you may already have this information available to you.
   - Visually locate the adapter on your server and look up the physical location code in the hardware publications distributed with the server.

- If AIX is installed and running on the server, run the AIX **lscfg -p** command to list the physical location codes of all of the hardware devices installed on your system.

The distribution of your IP addresses determines how many times you perform this step. You may have to do it more than once if:

- There are gaps in your IP addresses not caused by:
  - Wide nodes
  - High nodes
  - SP-attached servers
  - Clustered enterprise servers
- You want to set up alternate default routes or netmasks for certain IP address ranges.

The following example adds SDR information for an fi0 (FDDI adapter) network of 30 nodes (frame 1 slot 1 to frame 2 slot 16, with a wide node as the first node in each frame and the rest thin nodes) with IP addresses from 129.33.34.1 to 129.33.34.30, and a netmask of 255.255.255.0. The IP addressing corresponds to the slots in the frame, with each wide node incrementing by 2 and each thin node incrementing by 1.

```
spadaptrs -s yes 1 1 30 fi0 129.33.34.1 255.255.255.0
```

This example adds SDR information for a tr0 (token ring adapter) for node 1 with IP address 129.33.35.1 and a netmask of 255.255.255.0, and references the node list field.

```
spadaptrs -l 1 -r 16 tr0 129.33.35.1 255.255.255.0
```

This example adds SDR information for an additional Ethernet adapter for the second logical partition in a pSeries 690 server. The adapter is a twisted pair Ethernet adapter with duplex, the speed set to auto-negotiate, and is not the SP Ethernet adapter for the node. The IP address is 129.33.35.66 with a netmask of 255.255.255.0. The pSeries 690 server is represented as frame 5, the node is assigned slot 2, and the adapter is located at the physical location U1.9-P2-I2/E4.

```
spadaptrs -P U1.9-P2-I2/E4 -t tp -d auto -f auto 5 2 1 en \
          129.33.35.66 255.255.255.0
```

For nodes running DCE: You need to perform the following steps only if you have a new adapter on a node that is running DCE.

Note: All adapters must have an **ftp** and a **host** account defined in the DCE database.

1. Login to the DCE cell as a cell administrator.
2. Run **kerberos.dce -type admin -ip_name** *hostname_of_adapter*

   Note: You must either reboot or customize the node in order for the adapter to be configured.
3. After the node has completed reboot or customization, login as root to the node and issue **kerberos.dce -type local**

Step 8: Configure initial host names for nodes

Do this step if:

- You do not want the default host name to match the SP Ethernet administrative LAN adapter name. The SP Ethernet administrative LAN adapter name is the default.

- You are using short host names. The default is long host names.

This step changes the default host name information in the SDR Node Objects used during customization to set up the host name on each node, and allows you to indicate how you want to name your RS/6000 SP nodes. The default is the long form of the SP Ethernet administrative LAN adapter host name, which is how the **spadaptrs** command processes defaulted host names.

You can indicate an adapter name other than SP Ethernet administrative LAN adapter for the node host names to be used, as well as whether the long or short form should be used. When determining whether you want the nodes' host name to be either in long or short form, be consistent with the host name resolution on the control workstation. If the **host** command returns the short form of a host name, you should choose the short form for the node's initial host name.

Hostnames containing multibyte character data are not supported on the SP.

The following example indicates that the host name of each node is the long (fully qualified) form of the host name of the **css0** adapter for a system with two frames and 32 nodes:

```
sphostnam -a css0 1 1 32
```

Step 9: Set up nodes to be installed

- Do this step if you want to change the default installation settings for any of the nodes. To find out the default settings of your nodes, use the **splstdata** command.
- Be aware that the root password is not set if you are installing from a minimal mksysb. For more information on setting a root password when installing from a minimal mksysb, refer to "Step 61: Perform additional node customization" on page 90.
- If the boot/install server will be forwarding packets from the control workstation to a client node, the boot/install server is acting as a gateway to the control workstation. Therefore, ipforwarding must be correctly enabled. To turn ipforwarding on, issue:

  ```
  /usr/sbin/no -o ipforwarding=1
  ```
- You cannot export **/usr** or any directories below **/usr** because an NFS export problem will occur.

  If you have exported the **/spdata/sys1/install/image** directory or any parent directory, you must unexport it using the **exportfs -u** command before running **setup_server**. You need to do this because NIM attempts to export **/spdata/sys1/install/image/bos.obj.ssp.***, where **bos.obj.ssp.*** is the install image during **setup_server** processing. If you do not perform this task, you will receive an error. See the "Diagnosing NIM problems" chapter in the *PSSP: Diagnosis Guide* for more information.
- If you have selected DCE as an authentication method, ensure that you have DCE file sets installed in the lppsource directory for the node.
- Everything that is required in PSSP is installed on the nodes automatically, regardless of whether you use your own mksysb or the SP minimal mksysb.
- If you selected secure remote command methods for your remote command environment (see "Step 30: Enter site environment information" on page 49, you must ensure that the secure remote command program you intend to use is

available to be installed on the node during the PSSP installation process. Refer to “Step 1.10: Install the secure remote command software” on page 14 for more information.

- **SP-attached server and clustered enterprise server notes:**
  - If you are adding an SP-attached server or clustered enterprise server and want to preserve your current software environment, you must set the node to **customize** instead of **install**. For example:

    ```
    spbootins -r customize -l 33
    ```

  - If the root volume group of the SP-attached server or clustered enterprise server has been mirrored and you want to preserve the mirroring, you must record information about the existing mirrors in the SDR. If the root volume group of the M80, for example, has two copies on two physical disks in locations 30-68-00-0,0 and 30-68-00-2,0 with quorum turned off, enter the following to preserve the mirroring:

    ```
    spchvgobj -r selected_vg -c 2 -q false -h 30-68-00-0,0:30-68-00-2,0 -l 33
    ```

    Note: Failure to record the mirroring information will result in the root volume group being unmirrored during customization.

    To verify the information, enter:

    ```
    splstdata -b -l 33
    ```

  - Make sure that the PSSP *code_version* is set to PSSP-3.4.

This step does the following:

- Changes the default boot/install information for the node objects in the SDR so that you can indicate a different boot/install server configuration to the RS/6000 SP system
- Allows you to specify an alternate disk or disks to use when installing AIX on nodes

The default installation assumes one of the following:

- You have fewer than 40 nodes and the control workstation is configured to act as the boot/install server.
- You have more than 40 nodes, the control workstation and the first node in each frame is configured to act as boot/install server.

The default installation assumes your nodes have not been preinstalled. If you want to have them installed with your own install image, you must specify the following:

- Which nodes you are installing with your own install image
- The name of the installation image you are using if you do not want to use the default image

If you want different nodes to be installed by a different boot/install server, you must specify the target nodes and which node will serve as the boot/install server.

Selecting an installation disk

There are five ways you can specify the disk or disks to use for installation.

1. The hardware location format

   IBM strongly suggests that you use this format for SCSI devices. It ensures that you install on the intended disk by targeting a specific disk at a specific location. The relative location of hdisks can change depending on the hardware installed or possible hardware failures. You should always use this format when there are external disk drives present, because the manner in which the device names

are defined, may not be obvious. Installation on external disk drives is not supported. For example, to specify a single SCSI drive, enter:

```
00-00-00-0,0
```

or enter multiple hardware locations separated by colons:

```
00-00-00-0,0:00-00-00-1,0
```

2. The device names format

   For example, to specify a single device name, enter:

   ```
   hdisk0
   ```

   or enter multiple device names separated by commas:

   ```
   hdisk0,hdisk1
   ```

3. A combination of the parent and connwhere attributes for SSA devices

   To specify the parent-connwhere attribute:

   ```
   ssar//0123456789ABCDE
   ```

   or to specify multiple disks, separate using colons as follows:

   ```
   ssar//0123456789ABCDE:ssar//0123456789ABCDE
   ```

   The parent-connwhere format should only be used for SSA drives.

   For more information on acquiring ssar numbers, see *AIX Kernel and Subsystems Technical Reference, Volume 2*.

4. The PVID format

   If a disk was previously configured as a physical volume in order for it to be assigned to a volume group, a physical volume identifier (PVID) was assigned to that disk by AIX. You can specify a disk by its PVID value as a string of 16 hexadecimal characters. For example:

   ```
   00d4c45202be737f
   ```

   To specify multiple disks by their PVID values, separate the specifications using colons:

   ```
   00d4c45202be737f:00d4c452eb639a2c
   ```

   Use the AIX **lspv** command to list the PVID values for the disks on your system. For more information on making an available disk a physical volume and setting its PVID, see *AIX System Management Guide: Operating System and Devices*.

5. The SAN target and logical unit identifier format

   Fibre channel attached disks are identified by a worldwide port name and a logical unit identifier (LUN ID). To specify the SAN_DISKID, combine the two values into a single string separated by *"//"*. For example, if the SAN target worldwide port name for a fibre channel attached disk is 0x50060482bfd12c9c and the LUN ID is 0x8000000000000000, the SAN_DISKID specification would be:

   ```
   0x50060482bfd12c9c//0x8000000000000000
   ```

   To specify multiple fibre channel disks, separate the specifications using colons:

   ```
   0x50060482bfd12c9c//0x8000000000000000:0x50060482bbffd7cb//0x0
   ```

   Use the AIX **lsattr -EH -l** *hdisk* command to determine the worldwide port name and LUN ID for a disk.

The hardware location, SSA parent-connwhere, PVID, and SAN_DISKID formats can be used together. Specify multiple mixed format disk values using colons to separate the specifications as follows:

```
00-00-09-0,1:ssar//0123456789ABCDE:00d4c45202be737f
```

The device names format cannot be combined with any of the other format types.

You can use the **spchvgobj** command using the hardware location format for disk locations 00-07-00-0,0 and 00-07-00-1,0 for node 9. For example:

```
spchvgobj -r selected_vg -h 00-07-00-0,0:00-07-00-1,0 -1 9
```

If you need to change *lppsource_name* from default to a new *lppsource_name* such as aix433 for nodes 1 through 16, issue:

```
spchvgobj -r selected_vg -v aix433 1 1 16
```

If you need to change the *install_image_name* from default to *install_image_name* such as **bos.obj.ssp.433** for nodes 17, 18, 21, 22, issue:

```
spchvgobj -r selected_vg -i bos.obj.ssp.433 -v aix433 -l 17,18,21,22
```

For more information on alternate root volume groups, see the “Managing root volume groups” appendix in *PSSP: Administration Guide*.

Note: The AIX **alt_disk_install** function is not related to the SP alternate root volume group support and is not supported with PSSP installation.

Mirroring the root volume group

One way to significantly increase the availability of the SP system is to set up redundant copies of the operating system on different physical disks using the AIX disk mirroring feature. Mirroring the root volume group means that there will be multiple copies of the operating system image available to a workstation or node. Mirrored system images are distributed so that a node can remain in operation even after one of the mirrored units fail.

When installing a node, you have a choice of how many copies of the root volume group you would like. AIX allows one (the original), two (the original plus one), or three (the original plus two) copies of a volume group. IBM strongly suggests that the root volume group be mirrored for a total of at least two copies. PSSP provides commands to facilitate root volume group mirroring.

You can specify how many copies and which disks to use with the **spchvgobj** command. Care should be taken when specifying disks so that no other single point of failure is introduced. For example, the specified disks should not be attached to the same adapter.

The default setting for the number of copies is based on the node type. The default is one copy for all nodes except the POWER3 Symmetric Multiprocessor (SMP) High Node, which has a default of two copies. These nodes are assumed to contain dual internal disk drives as a standard configuration. The disks will automatically be used for mirroring. If these nodes were not configured with the dual internal disks or you do not want mirroring, use the **spchvgobj** command to change the settings before installing the node.

You can use the **spchvgobj** command using the hardware location format for disk locations 00-07-00-0,0 and 00-07-00-1,0 for node 9 and set the number of copies to two. For example:

```
spchvgobj -r selected_vg -h 00-07-00-0,0:00-07-00-1,0 -1 9 -c 2
```

For a complete description of how mirroring is handled by PSSP, see the “Managing root volume groups” appendix in *PSSP: Administration Guide*.

Step 10: Update the security information for the new nodes

The new nodes will need to be configured to the DCE database. Perform the following steps to add the new nodes. Step 1, Step 2, and Step 3 are required for DCE. Step 4 is required for both DCE and Kerberos V4.

1. You must set a DCE host name for each node in the SDR. This step uses the nodes' reliable host name as the DCE host name if a DCE host name does not already exist. Run **create_dcehostname** to update the SDR Node class attribute dcehostname for the new nodes.
2. Run **setupdce** so the new nodes principals can be added to the DCE registry.

   Notes:

   a. You must know the cell administrator password to perform this step.
   b. To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being configured. The value must be a resolvable address. For example:

      ```
      export SP_NAME=spcws.abc.com
      ```
3. Run **config_spsec** so the new nodes service principals can be created.

   Notes:

   a. You must have cell administrator authority to perform this step.
   b. To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being configured. Refer to the **config_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.
4. All nodes in the system partition need to be updated and, therefore, the control workstation's authorization files need to be updated as well. To create the authorization files issue **updauthfiles**.

Step 11: Refresh RSCT subsystems

The SDR has now been updated to reflect the new nodes that will run PSSP 3.4. You now need to refresh the system RSCT subsystems on the control workstation and all nodes to pick up these changes. Run **syspar_ctrl** on the control workstation to refresh the subsystems on both the control workstation and on the nodes.

```
syspar_ctrl -r -G
```

Step 12: Verify all node information

This step verifies that all the node information has been correctly entered into the SDR.

If using:	Do this:	
splstdata	**To display SDR:**	**Enter:**
	Site environment data	splstdata -e
	Frame data	splstdata -f
	Node data	splstdata -n
	Adapter data	splstdata -a
	Boot/install data	splstdata -b
	SP expansion I/O data	splstdata -x
	SP security settings	splstdata -p
	Switch data	splstdata -s
If your system is large, **splstdata** returns great quantities of data. You may want to pipe the command output through a filter to reduce the amount of data you see.		

Step 13: Add the new node to the root authorization files

The new nodes must be added to the root authorization files on previously-installed nodes. Issue:

```
dsh -avG /usr/lpp/ssp/bin/updauthfiles
```

Step 14: Configure the boot/install server

Do this step if you have not run **setup_server** already using the SMIT Boot/Install Server Information window or the **spbootins** command.

This step uses the information entered in the previous steps to set up the control workstation and optional boot/install servers on nodes. It configures the control workstation as a boot/install server and configures the following options (when selected in your site environment):

- Automounter
- File Collections
- NTP
- User Management
- Accounting

You can perform this step more than once. If you encounter any errors, see the *PSSP: Diagnosis Guide* for further explanation. After you correct your errors, you can start the task again.

In previous releases of PSSP, most of the installation function which configured boot/install servers and clients was performed in the single program called **setup_server** which you could run by issuing the **setup_server** command. This is still the suggested way for configuring the control workstation. For more experienced system administrators, IBM has provided a set of Perl scripts you can issue to also configure the control workstation that enable you to diagnose how the **setup_server** program is progressing.

If you have a node defined as a boot/install server, you must also run **setup_server** out on that server node. Enter the **setup_server** command on the control workstation with no parameters. For example:

```
setup_server
```

The first time **setup_server** runs, depending upon your configuration, it can take a significant amount of time to configure the control workstation as a NIM master.

Step 15: Change the default network tunable values

When a node is installed, migrated, or customized (set to **customize** and rebooted), and that node's boot/install server does not have a **/tftpboot/tuning.cust** file, a default file of system performance tuning variable settings in **/usr/lpp/ssp/install/config/tuning.default** is copied to **/tftpboot/tuning.cust** on that node. You can override these values by following one of the methods described in the following list:

1. Select an IBM-Supplied Alternate Tuning File

   IBM supplies three alternate tuning files which contain initial performance tuning parameters for three different SP environments:

   a. **/usr/lpp/ssp/install/config/tuning.commercial** contains initial performance tuning parameters for a typical commercial environment.

   b. **/usr/lpp/ssp/install/config/tuning.development** contains initial performance tuning parameters for a typical interactive/development environment.

   c. **/usr/lpp/ssp/install/config/tuning.scientific** contains initial performance tuning parameters for a typical engineering/scientific environment.

      Note: The SP-attached servers and clustered enterprise servers should not use the **tuning.scientific** file because of the large number of processors and the amount of traffic that they can generate.

   To select one of these files for use throughout the nodes in your system, use SMIT or issue the **cptuning** command. When you select one of these files, it is copied to **/tftpboot/tuning.cust** on the control workstation and is propagated from there to each node in the system when it is installed, migrated, or customized. Each node inherits its tuning file from its boot/install server. Nodes which have as their boot/install server another node (other than the control workstation) obtain their tuning.cust file from that server node so it is necessary to propagate the file to the server node before attempting to propagate it to the client node. The settings in the **/tftpboot/tuning.cust** file are maintained across a boot of the node.

2. Create and Select Your Own Alternate Tuning File

   The following steps enable you to create your own customized set of network tunable values and have them propagated throughout the nodes in your system. These values are propagated to each node's **/tftpboot/tuning.cust** file from the node's boot/install server when the node is installed, migrated, or customized and are maintained across the boot of the node.

   a. On the control workstation, create the file **/tftpboot/tuning.cust**. You can choose to begin with a copy of the file located in **/usr/lpp/ssp/samples/tuning.cust** which contains a template of performance tuning settings which have been commented out. Or you may prefer to begin with a copy of one of the IBM-supplied alternate tuning files.

   b. Select the tunable values that are best for your system.

   c. Edit the **/tftpboot/tuning.cust** file by ensuring the appropriate lines are uncommented and that the tunable values have been properly set.

Using SMIT:

TYPE **smit select_tuning**

- The Select System Tuning Parameters menu appears.

SELECT The desired tuning file

Once you have updated **tuning.cust**, continue installing the nodes. After the nodes are installed and customized, on all subsequent boots, the tunable values in **tuning.cust** will be automatically set on the nodes.

Note that each of the supplied network tuning parameter files, including the default tuning parameter file, contains the line **/usr/sbin/no -o ipforwarding=1**. IBM suggests that on non-gateway nodes, you change this line to read **/usr/sbin/no -o ipforwarding=0**. After a non-gateway node has been installed, migrated, or customized, you can make this change in the **/tftpboot/tuning.cust** file on that node.

If you are configuring more than eight of one particular adapter type, you must change the ifsize parameter in the **tuning.cust** file.

For the latest performance and tuning information, refer to the RS/6000 Web site at:

http://www.rs6000.ibm.com/support/sp/perf

Step 16: Perform additional node customization

Do this step to perform additional customization such as:

- Adding **installp** images
- Configuring host name resolution
- Setting up NFS, AFS, or NIS
- Configuring adapters that are not configured automatically
- Modifying TCP/IP configuration files
- Setting time zones

IBM provides the opportunity to run customer-supplied scripts during node installation:

script.cust This script is run from the PSSP NIM customization script (**pssp_script**) after the node's AIX and PSSP software have been installed, but before the node has been rebooted. This script is run in a limited environment where not all services are fully configured. Because of this limited environment, you should restrict your use of **script.cust** to function that must be performed prior to the post-installation reboot of the node.

firstboot.cust This script is run during the first boot of the node immediately after it has been installed. This script runs in a more "normal" environment where most all services have been fully configured. This script is a preferred location for node customization functions that do not require a reboot of the node to become fully enabled.

firstboot.cmds
When in restricted root remote access mode and secure remote command mode, this sysctl script is run on the control workstation during node installation to copy critical files from the control workstation to the nodes. It is enabled in the **firstboot.cust** script. See the **firstboot.cmds** and **firstboot.cust** files for information on how to set up and enable this script for sysctl.

Note: Your security environment is not set up during **script.cust** processing. If you are using AIX remote commands or SP Trusted Services, perform your

customization during **firstboot.cust** processing. See “Appendix E. User-supplied node customization scripts” on page 297 for additional information.

See “Appendix E. User-supplied node customization scripts” on page 297 for more detailed information on:

- The run-time environment for each of these scripts
- How to create and where to place the scripts

“Appendix E. User-supplied node customization scripts” on page 297 also discusses migration and coexistence issues and techniques to use the same set of customization scripts across different releases and versions of AIX and PSSP.

Note: When PSSP installs a node, it uses the AIX **sysdumpdev -e** command to estimate the size of the dump for the node. PSSP creates a dump logical volume that is approximately 10 percent larger than the estimated dump size, and makes that logical volume the primary dump device. However, you may find that this dump device is not large enough to contain an entire dump due to large processes or applications running on your node.

Once your node is up and running, use:

sysdumpdev -e	To get the estimated size of the node’s dump
sysdumpdev -l	To find the name of the primary dump device
lslv	To list the amount of space available in the primary dump device
extendlv	To expand the size of the dump logical volume if the estimated dump space is greater than the dump space available

Step 17: Additional switch configuration

Frames with Switches

If you have added a frame with a switch, you will need to perform “Step 17.1: Select a topology file” through “Step 17.4: Storing the switch topology file in the SDR” on page 219. If you have an SP Switch, you will also need to perform “Step 17.5: Set the switch clock source for all switches (SP Switch only)” on page 220.

Nodes or SP-attached servers (SP switch only)

If you have only added nodes or an SP-attached server, you will need to perform only “Step 17.3: Annotating a switch topology file” on page 219 and “Step 17.4: Storing the switch topology file in the SDR” on page 219.

Step 17.1: Select a topology file

Select the correct switch topology file by counting the number of node switch boards (NSBs) and intermediate switch boards (ISBs) in your system, then apply these numbers to the naming convention. The switch topology files are in the **/etc/SP** directory on the control workstation.

NSBs are switches mounted in slot 17 of frames containing nodes or SP Switch2 switches mounted in slots 2 through 16 of frames designed as multiple NSB frames. Multiple NSBs are used in systems that require a large number of switch connections for SP-attached servers or clustered enterprise server configurations. ISBs are switches mounted in the switch frame. ISBs are used in large systems,

where more than four switch boards exist, to connect many processor frames together. SP-attached servers never contain a node switch board, therefore, never include non-SP frames when determining your topology files.

The topology file naming convention is as follows:

`expected.top.`*`NSBnum`*`nsb.`*`ISBnum`*`isb.`*`type`*

where:

- *NSBnum* is the number of NSBs in the configuration
- *ISBnum* is the number of ISBs in the configuration
- *type* is the type of topology. The default type is 0.

For example, **expected.top.2nsb.0isb** is a file for a two frame and two switch system with no ISB switches.

The exception to this naming convention is the topology file for the SP Switch-8 configuration, which is **expected.top.1nsb_8.0isb.1**.

See the **Etopology** command in *PSSP: Command and Technical Reference* for additional information on topology file names.

Step 17.2: Managing the switch topology files

The switch topology file must be stored in the SDR. The switch initialization code uses the topology file stored in the SDR when starting the switch (**Estart**). When the switch topology file is selected for your system's switch configuration, it must be annotated with **Eannotator**, then stored in the SDR with **Etopology**. The switch topology file stored in the SDR can be overridden by having an **expected.top** file in **/etc/SP** on the primary node. **Estart** always checks for an **expected.top** file in **/etc/SP** before using the one stored in the SDR. The **expected.top** file is used when debugging or servicing the switch.

Notes:

1. Be aware that **Estart** distributes the topology file to all the nodes in the system partition on the switch. In the case of **expected.top**, this is significant because if the topology file is left on a node and the primary is changed to that node, the topology file will be used. If you have an **expected.top** file in **/etc/SP** on any of the nodes, make sure that you remove it when it is no longer needed.
2. Depending upon your configuration, the first **Estart** of the switch may take longer than subsequent **Estart**s.
3. In a two plane SP Switch2 system, the function of **/etc/SP/expected.top** is taken over by **/etc/SP/expected.top.p0** for plane 0 and by **/etc/SP/expected.top.p1** for plane 1.

Step 17.3: Annotating a switch topology file

Use the **Eannotator** command to update the switch topology file's connection labels with their correct physical locations. Use the **-O yes** flag to store the switch topology file in the SDR. Using **Eannotator** makes the switch hardware easier to debug because the switch diagnostics information is based on physical locations.

For example, to annotate a two-switch or maximum 32-node system, enter:

```
Eannotator -F /etc/SP/expected.top.2nsb.0isb.0 \
           -f /etc/SP/expected.top.annotated -O yes
```

Step 17.4: Storing the switch topology file in the SDR

If you entered **Eannotator -O yes** or **yes** on the Topology File Annotator menu in "Step 17.3: Annotating a switch topology file", skip this step.

Use the **Etopology** command to store the switch topology file in the SDR and make sure that it has been annotated. For example, to store a two-switch or maximum 32-node configuration, enter:

```
Etopology /etc/SP/expected.top.2nsb.0isb.0.annotated
```

Step 17.5: Set the switch clock source for all switches (SP Switch only)

Use SMIT or the **Eclock** command to initialize the switch's clock source. The SMIT and **Eclock** interfaces require that you know the number of Node Switch Boards (NSBs) and Intermediate Switch Boards (ISBs) in your RS/6000 SP system.

Select the **Eclock** topology file from the control workstation's **/etc/SP** subdirectory, based on these numbers. For example, if your RS/6000 SP system has six node switch boards and four intermediate switch boards, you would select **/etc/SP/Eclock.top.6nsb.4isb.0** as an **Eclock** topology file.

See *PSSP: Command and Technical Reference* for the **Eclock** topology file names.

Use the **Eclock** command to set the switch's clock source for all switches.

For example, if your RS/6000 SP system has six node switch boards and four intermediate switch boards, select **/etc/SP/Eclock.top.6nsb.4isb.0** as an **Eclock** topology file. Enter:

```
Eclock -f /etc/SP/Eclock.top.6nsb.4isb.0
```

This command sets the proper clock source settings on all switches within a 96-way (6 NSB, 4 ISB) RS/6000 SP system.

To verify the switch configuration information, enter:

```
splstdata -s
```

Step 18: Redefine system partitions (SP Switch or switchless systems only)

If you want to partition your system, you can select an alternate configuration from a predefined set of system partitions to implement before booting the nodes or you can use the System Partitioning Aid to generate and save a new layout. Follow the procedure described in the "Managing system partitions" chapter in the *PSSP: Administration Guide* and refer to information in "The System Partitioning Aid" section of the "Planning SP system partitions" chapter in the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*. You do not have to partition your system now as part of this installation. You can partition it later.

Note: System partitioning is not supported on clustered enterprise servers or on systems with an SP Switch2 switch.

If you have a frame with a switch or a frame with a switchless system, you will need to redefine your system partition configuration to match the hardware. At this point, you should not move any existing nodes to different system partitions. If you want to reconfigure your system partitions after completing this task, see *PSSP: Administration Guide* for more information.

Step 19: Network boot optional boot/install servers

If you are adding a node or nodes that will function as a boot/install server, you will need to perform this step.

SP Switch note: If you have set up system partitions, do this step in each partition.

SP-attached server and clustered enterprise server note: If you do not want to reinstall your existing SP-attached server or clustered enterprise server, but want to preserve its environment, perform steps “Step 19.1: Upgrade AIX” through “Step 19.10: Reboot” on page 222.

To monitor installation progress by opening the node’s read-only console, issue:

```
s1term frame_id slot_id
```

If you have eight or more boot/install servers on a single Ethernet segment, you should network boot those nodes in groups of eight or less. See the “IP performance tuning” section in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more information.

To network boot your nodes, issue:

```
nodecond frame_id slot_id &
```

Note: For MCA nodes, the **nodecond** command remotely processes information from the initial AIX firmware menus. You should not change the language option on these menus. The language must be set to English in order for the **nodecond** command to run properly.

To check the LCD and LED display for each node, enter:

```
spmon -Led nodenode_number
```

or

```
spled &
```

Network installation progress

When a network installation is in progress, the LED for the nodes involved show various values. These values indicate the installation stage. Since the node installation process can be long, it is hard to determine where you are in that process. Refer to *PSSP: Diagnosis Guide* for a complete list of PSSP-specific LED values.

Note: Perform “Step 19.1: Upgrade AIX” through “Step 19.10: Reboot” on page 222 **only** if you want to preserve the environment of your existing SP-attached server or clustered enterprise server.

SP-attached server and clustered enterprise server installation

Perform the following steps to add an SP-attached server or clustered enterprise server and preserve your existing software environment.

Step 19.1: Upgrade AIX: If your SP-attached server or clustered enterprise server is not at AIX 4.3.3, you must first upgrade to that level of AIX before proceeding.

Step 19.2: Set up name resolution of the SP-attached server or clustered enterprise server: In order to do PSSP customization, the following must be resolvable on the SP-attached server or clustered enterprise server:

- The control workstation host name
- The name of the boot/install server’s interface that is attached to the SP-attached server or clustered enterprise’s SP Ethernet administrative LAN adapter interface

Step 19.3: Set up routing to the control workstation host name: If you have a default route set up on the SP-attached server or clustered enterprise server, you

will have to delete it. If you do not remove the route, customization will fail when it tries to set up the default route defined in the SDR. In order for customization to occur, you must define a static route to the control workstation's host name. For example, the control workstation's host name is its token ring address, such as 9.114.73.76 and your gateway is 9.114.73.256:

```
route add -host 9.114.73.76 9.114.73.256
```

Step 19.4: FTP the SDR_dest_info file: During customization, certain information will be read from the SDR. In order to get to the SDR, you must FTP the **/etc/SDR_dest_info** file from the control workstation to the **/etc/SDR_dest_info** file on the SP-attached server or clustered enterprise server and check the mode and ownership of the file.

Step 19.5: Verify perfagent: Ensure that perfagent.tools 2.2.32.x is installed on your SP-attached server or clustered enterprise server.

Step 19.6: Mount the pssplpp directory: Mount the **/spdata/sys1/install/pssplpp** directory on the boot/install server from the SP-attached server or clustered enterprise server. For example, issue:

```
mount k3n01:/spdata/sys1/install/pssplpp /mnt
```

Step 19.7: Install ssp.basic: Install **ssp.basic** and its prerequisites onto the SP-attached server or clustered enterprise server. For example, issue:

```
installp -aXgd/mnt/PSSP-3.4 ssp.basic 2>&1 | tee /tmp/install.log
```

Step 19.8: Unmount the pssplpp directory: Unmount the **/spdata/sys1/install/pssplpp** directory on the boot/install server from the SP-attached server or clustered enterprise server. For example, issue:

```
umount /mnt
```

Step 19.9: Run pssp_script: Run the **pssp_script** by issuing:

```
/usr/lpp/ssp/install/bin/pssp_script
```

Step 19.10: Reboot: Perform a reboot. For example:

```
shutdown -Fr
```

Step 20: Verify that System Management tools were correctly installed on the boot/install servers

Now that the boot/install servers are powered up, run the verification test from the control workstation to check for correct installation of the System Management tools on these nodes.

To do this, enter:

```
SYSMAN_test
```

After the tests are run, the system creates a log in **/var/adm/SPlogs** called **SYSMAN_test.log**.

See the section on "Verifying System Management installation" in the *PSSP: Diagnosis Guide* for information on what this test does and what to do if the verification test fails.

Step 21: Network boot the remaining RS/6000 SP nodes

SP Switch note: If you have set up system partitions, do this step in each partition.

Repeat the procedure used in “Step 19: Network boot optional boot/install servers” on page 220 to network boot and install, or customize the remaining nodes. You may need to ensure that all setup_server processes have completed on the boot/install nodes prior to issuing a network boot on the remaining nodes. Refer to the **/var/adm/SPlogs/sysman/***node***.console.log** file on the boot/install node to see if **setup_server** has completed.

If any of your boot/install servers have more than eight clients on a single Ethernet segment, you should Network Boot those nodes in groups of eight or less. See the “IP performance tuning” section in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more information.

Using a token ring-bridge gateway

If you are using a token ring through a bridge as your default gateway to your nodes and the token ring bridge is not on the same segment as your LAN, you must change the value of the broadcast field in the ODM for each node. The default value is set to **No** (confine broadcast to local token-ring) each time you install or customize a node. However, when you boot the nodes with this bridge setup, the network is unusable.

To do this, enter:

```
chdev -P -l tr0 -a allcast=off
```

Step 22: Verify node installation

To check the hostResponds and powerLED indicators for each node, enter:

```
spmon -d -G
```

Step 23: Verify the SP expansion I/O unit configuration

To verify that the SP expansion I/O unit is properly configured in the SDR, issue:

```
splstdata -x
```

Step 24: Enable s1_tty on the SP-attached server or clustered enterprise server (SAMI hardware protocol only)

If you just installed an SP-attached server or clustered enterprise server, you must ensure that the s1_tty is enabled on the server. Until the login is enabled on the tty, the **s1term** command from the control workstation to the SP-attached server or clustered enterprise server will not work.

On the SP-attached server or clustered enterprise server, determine which tty is mapped to 01-S1-00-00. For example, issue the following:

```
lsdev -C -c tty0
```

In response, the system displays something similar to:

```
tty0 Available 01-S1-00-00 Asynchronous Terminal
tty1 Available 01-S2-00-00 Asynchronous Terminal
```

In the previous example, tty0 is mapped to 01-S1-00-00.

Set the login to enable. For example, issue the following:

```
chdev -l tty0 -a login=enable
```

Step 25: Start the switch (optional)

Do this step if you have a switch installed in your system. If you have set up system partitions (SP Switch only), do this step in each partition.

```
Estart
```

Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd** on SP Switch systems or **startsrc -s swtadmd2** on SP Switch2 systems before issuing the **Estart** command.

Step 26: Verify that the switch was installed correctly

- Do this step if you have added an additional switch or switches to your system.
- If the switches you added span multiple system partitions, verify the switch information from within each system partition (SP Switch only).
- Check all connectors for miscabling and node communications.

Run a verification test to ensure that the switch is installed completely. To do this, enter:

```
CSS_test
```

After the tests are run, the system creates a log in **/var/adm/SPlogs** called **CSS_test.log**.

If the verification test fails, see the section on “Diagnosing switch problems” in the *PSSP: Diagnosis Guide*.

To check the switchResponds and powerLED indicators for each node, enter:

```
spmon -d -G
```

Step 27: Tune the network adapters for added nodes

Various models of network adapters can have different values for transmit and receive queue sizes. The queue setting for Micro Channel adapters is 512. For PCI adapters, the queue setting is 256 or greater.

Note: For AIX 4.2.1, the receive queue size is not tunable.

You can set these values using SMIT or the **chdev** command. If the adapter you are changing is also the adapter for the network you are logged in through, you will have to make the changes to the database only. Then reboot the nodes for the changes to become effective. To do this, enter:

```
chdev -P -l ent0 -a xmt_que_size=256
```

You must reboot the nodes in order for the changes to take effect.

Step 28: Reconfigure LoadLeveler to add the new node to the LoadLeveler cluster

If you are using LoadLeveler as your workload management system, add the new node to the LoadLeveler configuration. For more information on this step, see the “Administration tasks for parallel jobs” chapter in *IBM LoadLeveler for AIX 5L: Using and Administering*.

Adding an adapter to a node

This procedure can be used to logically add an adapter, also known as a network interface card (NIC), to a previously installed SP node. This includes Ethernet, token ring, and FDDI adapters. To configure adapters such as ATM, ESCON, and PCA, you must configure the adapter manually on each node.

Step 1: Archive the SDR

To save the current SDR attributes, run the command:

```
SDRArchive
```

This command saves SDR information for all system partitions.

Step 2: Disable nodes from the switch

If any of the following conditions exit, you can skip this step:

- You do not have a switch in your SP system
- You have already logically removed the SP Switch from your nodes
- The node you are adding the adapter to does not contain a switch adapter

If you want to bring the switch down for all nodes, issue the **Equiesce** command for each system partition. If you use the **Equiesce** command, you must later restart the switch by issuing the **Estart** command. Issue the **Estart** command prior to “Step 8: Verify installation settings” on page 227.

The Switch Admin Daemon will issue the **Estart** command under certain circumstances, thus re-establishing switch recovery and primary node takeover. To see if you are using the Switch Admin Daemon, issue the following command:

```
lssrc -a | grep swtadmd
```

or

```
lssrc -a | grep swtadmd2
```

where **swtadmd** is used for an SP Switch and **swtadmd2** for an SP Switch2.

If the response returned shows that the subsystem is active, you may want to turn off the Switch Admin Daemon before issuing the **Equiesce** command. To turn off the daemon, issue the following command:

```
stopsrc -s swtadmd
```

or

```
stopsrc -s swtadmd2
```

where **swtadmd** is used for an SP Switch and **swtadmd2** for an SP Switch2.

If you are customizing a few nodes, you must disable these nodes from the switch (if appropriate, first reassign the primary node or primary backup node). To determine if one of the nodes you are migrating is a primary or primary backup node, issue the **Eprimary** command. If you need to reassign the primary or primary backup node, issue the **Eprimary** command with appropriate options. Then issue the **Estart** command to make your choices take effect. You must then issue the **Efence** command to disable the nodes you are customizing from the switch.

```
Efence -G node_number node_number ... node_number
```

Step 3: Shutdown the node

The node must be powered off before the adapter can be installed. Stop the applications, then use SP Perspectives or the **cshutdown** command to shut down the node. For example, to shut down nodes 1 and 2:

```
cshutdown -G -N 1 2
```

Step 4: Install the new adapter

Note: Your IBM Customer Engineer (CE) performs this step.

While the Customer Engineer installs the adapter, the system administrator should proceed with the logical configuration of the adapter. **DO NOT REBOOT THE NODE** until “Step 11: Reboot the node” on page 228, when both physical and logical changes are complete.

Step 5: Define the adapter to the SDR

This step creates adapter objects in the SDR for each new adapter being added to a node. The data in the adapter objects is used during the customization or installation steps to configure the adapters on the nodes. You can configure the following adapters with this procedure:

- Ethernet (en)
- FDDI (fi)
- Token ring (tr)

To configure additional adapters, for example Ethernet (en), token ring (tr), or FDDI (fi), use **smit node_data** and select the **Additional Adapter Information** or use the **spadaptrs** command. For these adapters, you can select the Start Frame, Start Slot, Node Count, and Node List fields, or issue the **spadaptrs** command.

Notes:

1. When using the token ring (tr) adapter, you must select the Token Ring Data Rate (4 MB, 16 MB, or autosense).
2. For best results with SMIT, exit and get back into the needed SMIT panel for each different type of adapter. This clears any extraneous values left behind in the panel.
3. Enter the correct value for Ethernet speed (10, 100, 1000, or auto), Duplex (full, half, or auto), and Ethernet Adapter Type (bnc, dix, tp, fiber, or NA), for each new Ethernet adapter on the target node.

The distribution of your IP addresses determines how many times you issue the **spadaptrs** command. You may have to issue it more than once if:

- There are gaps in your IP addresses not caused by wide nodes, high nodes, SP-attached servers, or clustered enterprise servers.
- You want to set up alternate default routes or network masks for certain IP address ranges.

The following example adds SDR information for a tr0 (token ring adapter) for node 1 with IP address 9.114.12.36 and a netmask of 255.255.255.192, and references the Node List field:

```
spadaptrs -l 1 -r 16 tr0 9.114.12.36 255.255.255.192
```

The next example adds SDR information for a second Ethernet adapter on node 1 with a type of bnc, duplexing of half, a speed of 10MBS, an IP address of 9.114.84.11, and a network mask of 255.255.255.192:

```
spadaptrs -l 1 -t bnc -d half -f 10 en1 9.114.84.11 255.255.255.192
```

To validate successful creation of adapter objects, issue the command:

```
splstdata -G -a
```

and find the new objects.

Step 6: Update DCE information for a node (required for DCE)

You need to perform this step only if you have a new adapter on a node that is running DCE (that is, the adapter was added after DCE was configured). This step will create an **ftp** and a **host** account in the DCE Registry database for each new adapter.

1. Login to the control workstation, and then DCE login as a DCE cell administrator.

   Note: This step can be performed from another host configured with a DCE client in the same cell as the node's DCE client.

2. Run the DCE command:

   ```
   kerberos.dce -type admin -ip_name hostname_of_adapter
   ```

   Note: This command does not require AIX root authority, but does require DCE cell administrator authority.

To verify that the entries were created properly in the DCE Registry database, issue the DCE **account show** command, as shown in the following example:

```
dcecp -c account show /.:/host/hostname_of_adapter -all
dcecp -c account show /.:/ftp/hostname_of_adapter -all
```

Step 7: Set node to customize

Each node with a new adapter must be customized. Use the **spbootins** commands to update the boot/install server's response to the **bootp** request from the nodes. If you are customizing nodes in more than one system partition, you must issue these commands in each system partition. For a complete description of the flags associated with these commands, refer to *PSSP: Command and Technical Reference*.

For example, to set the **bootp** response to `customize` for nodes 1 and 2, issue the following command:

```
spbootins -s no -r customize -l 1,2
```

Step 8: Verify installation settings

Make sure that the SDR has the appropriate installation values specified for each of the nodes. Issue the following command to display the values:

```
splstdata -G -b
```

The expected response is that nodes with new adapters should be set to **customize**. All other attributes should contain the values last used to install or customize your nodes.

Step 9: Refresh RSCT subsystems

The SDR has been updated to reflect the new adapters. You must now refresh the RSCT subsystems on the control workstation and all nodes to pick up these changes. Run **syspar_ctrl** on the control workstation to refresh the subsystems on both the control workstation and on the nodes:

```
syspar_ctrl -r -G
```

Step 10: Run setup_server to configure the changes

The **setup_server** command must be run to properly set up NIM on the control workstation. Use the following command:

```
setup_server 2>&1 | tee /tmp/setup_server.out
```

The output is saved in a log file called **/tmp/setup_server.out**. If you have a node defined as a boot/install server (BIS), you must also run **setup_server** on that server node:

```
dsh -w boot_install_node "/usr/lpp/ssp/bin/setup_server \
    2>&1" | tee /tmp/setup_server.boot_install_node.out
```

Step 11: Reboot the node

Once the adapter has been physically installed in the node, the node can be rebooted to complete the logical configuration. If you have any boot/install servers in your system, you must customize them before customizing their clients. Reboot the nodes to be customized by using the **cstartup** command:

```
cstartup -G -N 1 2
```

The node has finished customization when the LEDs become blank, and its **host_responds** value is **yes**.

Verify that the **bootp_response** has been set to **disk** by issuing the following command:

```
splstdata -G -b
```

Step 12: Complete the DCE configuration of the new adapter (required for DCE)

Complete the DCE configuration of the new adapter on the node. This step can be done only after the customize has completed. Use the following command:

```
dsh -w node_name "kerberos.dce -type local"
```

Note: In order for the DCE **kerberos.dce** command to complete successfully, the command must be run as root on the target host, and DCE must be running on that host.

To verify that the entries were created properly in the DCE self-host (machine) key file on *node_name*, issue the DCE **keytab show** command, as shown in the following example:

```
dsh -w node_name "dcecp -c keytab show self"
```

Note: This DCE command has the same requirements as the previous Note.

Using the *hostname_of_adapter* data from “Step 6: Update DCE information for a node (required for DCE)” on page 227, scan the displayed output for **ftp** and **host** fields that contain *hostname_of_adapter*.

Step 13: Rejoin the nodes to the switch network

If you do not have a switch in your SP system, or if the node you are adding the adapter to does not contain a switch, skip the remainder of this step. In SP Switch2 systems, you can have some nodes on the switch and some nodes off of the switch.

If you disabled all nodes from the switch using the **Equiesce** command, you must now issue the **Estart** command in each system partition to rejoin the nodes to the switch network. If you disabled only a few nodes using the **Efence** command, you must now issue the **Eunfence** command to bring those nodes back to the current switch network.

If you stopped the Switch Admin Daemon in "Step 2: Disable nodes from the switch" on page 225, start it now by issuing the command:

```
startsrc -s swtadmd
```

or

```
startsrc -s swtadmd2
```

where **swtadmd** is used for an SP Switch and **swtadmd2** for an SP Switch2.

Step 14: Validate new adapter

Verify that the new adapter is operational. Use **ping** or other network commands.

Deleting a frame, node, SP-attached server, or clustered enterprise server

When you delete a frame, node, SP-attached server, or clustered enterprise server, from your system, you should first plan how the change will affect the remainder of your system. Consider the workload and applications currently running on the hardware to be deleted and plan how to transfer the workload and applications to equivalent SP resources.

- If the node you are deleting is a server node (for example, a switch primary node, primary backup node, NTP server, or boot/install server), reassign the server to another node that is not being deleted. If you reassign a primary or primary backup node, you will need to run the **Estart** command to have it take effect.

  Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd** on SP Switch systems or **startsrc -s swtadmd2** on SP Switch2 systems before issuing the **Estart** command.
- IBM suggests that you delete frames only at the end of your system otherwise you will have to reconfigure the SDR.
- Do these steps in the system partition from which you delete the nodes (SP Switch only).
- If you are deleting an extension node, you need to issue the **enrmnode** command followed by the **enrmadapter** command. For specific instructions, refer to the *PSSP: Administration Guide*.
- If you are deleting a frame with a switch in a single-frame environment, you must quiesce the switch using the **Equiesce** command before deleting the frame.

Note: If you are using the Switch Admin daemon for node recovery, stop it by issuing **stopsrc -s swtadmd** on SP Switch systems or **stopsrc -s swtadmd2** on SP Switch2 systems before issuing the **Equiesce** command.

- If you are deleting a frame that contains a switch that has an SP-attached server connected to it, you must first delete the SP-attached server.

  If you want to move your SP-attached server to another switch, you must add it back to your system with the new configuration information.
- When deleting nodes that were configured for DCE, you will need to unconfigure the nodes from DCE.
- If you a deleting a node that is a logical partition on a pSeries 690 server, follow the instructions in "Deleting an existing LPAR" on page 246 to also delete the partition information on the HMC.

Note: When you delete a node, the attached SP expansion I/O units are also deleted. When you delete a frame, the I/O units attached to nodes on that frame are also deleted regardless if they physically reside on that same frame or on a different frame.

Step 1: Archive the SDR

Before reconfiguring your system, you should back up the SDR by issuing:

```
SDRArchive
```

Note the location and the name of the file created after you issue this command.

Step 2: Unpartition your system (SP Switch only)

If your existing system has multiple partitions defined and you want to delete a frame, you need to bring the system down to one partition before you can delete the additional frame.

Note: The remaining system partition must contain all of the security settings from all of the system partitions before you can power off any nodes for security reasons.

Step 2.1: Repartition your system to a single system partition (SP Switch only)

See the "Managing system partitions" chapter in the *PSSP: Administration Guide* for instructions on partitioning your SP system.

Step 3: Reconfigure LoadLeveler to remove the frame data from the LoadLeveler cluster

For more information on this step, see the "Administration tasks for parallel jobs" chapter in *IBM LoadLeveler for AIX 5L: Using and Administering*.

Step 4: Reconfigure IBM Virtual Shared Disk to remove the frame data from the IBM Virtual Shared Disk cluster

For more information on this step, see *PSSP: Managing Shared Disks*.

Step 5: Shut down SP-attached servers, clustered enterprise servers, or the nodes in the frame

Use the **cshutdown** command to shut down the nodes, SP-attached servers, or clustered enterprise servers that you are deleting from your system. For example, to shut down node number 48 and nodes 16-31 in your system, enter:

```
cshutdown -G -N 48 16-31
```

Step 6: Disconnect the node from the frame (optional)

Perform this step only if you are deleting nodes from a frame that will continue to be part of the SP system. Your IBM Customer Engineer (CE) performs this step. See *RS/6000 SP: Installation and Relocation* for instructions.

To delete an existing LPAR from a pSeries 690 server attached to your SP or clustered enterprise server system, you will need to use the HMC interface to remove the logical partition from the server. Perform the following steps:

1. Start the HMC interface. This can be done in one of the following ways:
   - Login directly to the HMC subsystem. This will automatically start the GUI on the HMC display.
   - Manage the HMC subsystem through a remote WebSM session running on your control workstation or other UNIX workstation that has connectivity to the HMC.
   - Launch a remote WebSM session to the HMC from the Hardware Perspective running on your control workstation.
2. Select the pSeries 690 system object and follow HMC procedures for deleting a logical partition. Refer to the *Hardware Management Console for pSeries Operations Guide* for details on performing this operation.

Step 7: Unconfigure DCE-related information for the node (required for DCE)

After the node has been deleted from the frame, you must remove any DCE-related principles and objects from the DCE registry. You must also unconfigure DCE (Admin portion).

On the control workstation, use the **rm_spsec -t admin** *node_dce_hostname* command first, then do a DCE Admin unconfigure for the node (**smit rmdce**).

Notes:

1. You must have cell administrator authority to perform this step.
2. To remove any additional principals related to the node using the SMIT panels, enter the host name of the adapter to be deleted. For example, on the “Admin only unconfiguration for another machine” panel in the “Machine’s name or TCP/IP address” field, enter the host name for the additional adapters.
3. For the nodes being removed, verify that all DCE principals have been deleted from the DCE registry. Issue:

   ```
   dcecp -c principal catalog -simplename
   ```
4. To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being reconfigured. Refer to the **rm_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.

Step 8: Disconnect SP expansion I/O units from the frame (optional)

If you are deleting a node that has attached SP expansion I/O units, also disconnect all I/O units from the frames in which they reside. Your IBM Customer Engineer (CE) performs this step. See *RS/6000 SP: Installation and Relocation* for instructions.

Step 9: Disconnect the hardware to be deleted

Your IBM Customer Engineer (CE) performs this step.

Step 10: Delete information from the SDR

Perform one of the following steps depending on what you are deleting.

Step 10.1: Frame, SP-attached server, or clustered enterprise server information

Perform this step if you are deleting a frame, SP-attached server, or clustered enterprise server from your system. The following example deletes the last two frames and all of the nodes contained in those frames:

```
spdelfram 3 2
```

Step 10.2: Clustered enterprise servers migration servers

Perform these steps if you are deleting all of your SP frames to convert your SP-attached servers to clustered enterprise servers.

Step 10.2.1: Delete all SP-attached server switch adapters: Perform this step if you are deleting all of your SP frames from a switched SP system to convert your SP-attached servers to clustered enterprise servers. Delete all switch adapter information from the SDR for your SP-attached servers before deleting your SP frames in "Step 10.2.2: Delete all SP frames".

```
spdeladap 5 1 0 css0
```

Step 10.2.2: Delete all SP frames: All of the SP frames in your system must be deleted in a single operation. For example, issue:

```
spdelfram -c -l 1,4
```

Step 10.3: Node information

Perform this step only if you are deleting a node from a frame that will continue to be part of your configuration. Use the **spdelnode** command to delete the node information from the SDR. For example, to delete node 17, issue:

```
spdelnode 2 1 1
```

Step 11: Set up system partitions (SP Switch only)

If you want to partition your system, you can select an alternate configuration from a predefined set of system partitions to implement before booting the nodes or you can use the System Partitioning Aid to generate and save a new layout. Follow the procedure described in the "Managing system partitions" chapter in the *PSSP: Administration Guide* and refer to information in "The System Partitioning Aid" section of the "Planning SP system partitions" chapter in the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*. You do not have to partition your system now as part of this reconfiguration. You can partition it later.

Step 12. Refresh authorization files in the system or in a system partition

Remote commands depend on authorization files on the control workstation and on the nodes for root access. When a node is deleted, these authorization files must be updated to remove the entry for the deleted node on all other nodes and on the control workstation that had knowledge of the deleted node. This could be all nodes in a system partition, all nodes in system partitions with the same authentication method, or all nodes in the entire SP system.

To update the authorization files, issue:

```
dsh -w node_list updauthfiles
```

Step 13: Refresh RSCT subsystems

To refresh the subsystems, issue the following command on the control workstation:

```
syspar_ctrl -r -G
```

This command refreshes all the subsystems in each partition such as **hats** so that it no longer recognizes the deleted hardware.

Step 14: Additional switch configuration

Frames with Switches

If you have deleted a frame with a switch, you will need to perform "Step 14.1: Select a topology file" through "Step 14.4: Storing the switch topology file in the SDR" on page 234 and "Step 14.6: Start the switch" on page 235. If you have an SP Switch, you will need to perform "Step 14.1: Select a topology file" through "Step 14.6: Start the switch" on page 235.

Nodes or SP-attached servers (SP Switch only)

If you have only deleted nodes or an SP-attached server, you will need to perform only "Step 14.3: Annotating a switch topology file" on page 234 and "Step 14.4: Storing the switch topology file in the SDR" on page 234.

Step 14.1: Select a topology file

Select the correct switch topology file by counting the number of node switch boards (NSBs) and intermediate switch boards (ISBs) in your system, then apply these numbers to the naming convention. The switch topology files are in the **/etc/SP** directory on the control workstation.

NSBs are switches mounted in slot 17 of frames containing nodes or SP Switch2 switches mounted in slots 2 through 16 of frames designed as multiple NSB frames. Multiple NSBs are used in systems that require a large number of switch connections for SP-attached servers or clustered enterprise server configurations. ISBs are switches mounted in the switch frame. ISBs are used in large systems, where more than four switch boards exist, to connect many processor frames together. SP-attached servers never contain a node switch board, therefore, never include non-SP frames when determining your topology files.

The topology file naming convention is as follows:

```
expected.top.NSBnumnsb.ISBnumisb.type
```

where:

- *NSBnum* is the number of NSBs in the configuration

- *ISBnum* is the number of ISBs in the configuration
- *type* is the type of topology. The default type is 0.

For example, **expected.top.2nsb.0isb** is a file for a two frame and two switch system with no ISB switches.

The exception to this naming convention is the topology file for the SP Switch-8 configuration, which is **expected.top.1nsb_8.0isb.1**.

See the **Etopology** command in *PSSP: Command and Technical Reference* for additional information on topology file names.

Step 14.2: Managing the switch topology files

The switch topology file must be stored in the SDR. The switch initialization code uses the topology file stored in the SDR when starting the switch (**Estart**). When the switch topology file is selected for your system's switch configuration, it must be annotated with **Eannotator**, then stored in the SDR with **Etopology**. The switch topology file stored in the SDR can be overridden by having an **expected.top** file in **/etc/SP** on the primary node. **Estart** always checks for an **expected.top** file in **/etc/SP** before using the one stored in the SDR. The **expected.top** file is used when debugging or servicing the switch.

Notes:

1. Be aware that **Estart** distributes the topology file to all the nodes in the system partition on the switch. In the case of **expected.top**, this is significant because if the topology file is left on a node and the primary is changed to that node, the topology file will be used. If you have an **expected.top** file in **/etc/SP** on any of the nodes, make sure that you remove it when it is no longer needed.
2. Depending upon your configuration, the first **Estart** of the switch may take longer than subsequent **Estart**s.
3. In a two plane SP Switch2 system, the function of **/etc/SP/expected.top** is taken over by **/etc/SP/expected.top.p0** for plane 0 and by **/etc/SP/expected.top.p1** for plane 1.

Step 14.3: Annotating a switch topology file

Use the **Eannotator** command to update the switch topology file's connection labels with their correct physical locations. Use the **-O yes** flag to store the switch topology file in the SDR. Using **Eannotator** makes the switch hardware easier to debug because the switch diagnostics information is based on physical locations.

For example, to annotate a two-switch or maximum 32-node system, enter:

```
Eannotator -F /etc/SP/expected.top.2nsb.0isb.0 \
           -f /etc/SP/expected.top.annotated -O yes
```

Step 14.4: Storing the switch topology file in the SDR

If you entered **Eannotator -O yes** or **yes** on the Topology File Annotator menu in "Step 14.3: Annotating a switch topology file", skip this step.

Use the **Etopology** command to store the switch topology file in the SDR and make sure that it has been annotated. For example, to store a two-switch or maximum 32-node configuration, enter:

```
Etopology /etc/SP/expected.top.2nsb.0isb.0.annotated
```

Step 14.5: Set the switch clock source for all switches (SP Switch only)

Use SMIT or the **Eclock** command to initialize the switch's clock source. The SMIT and **Eclock** interfaces require that you know the number of Node Switch Boards (NSBs) and Intermediate Switch Boards (ISBs) in your RS/6000 SP system.

Select the **Eclock** topology file from the control workstation's **/etc/SP** subdirectory, based on these numbers. For example, if your RS/6000 SP system has six node switch boards and four intermediate switch boards, you would select **/etc/SP/Eclock.top.6nsb.4isb.0** as an **Eclock** topology file.

See *PSSP: Command and Technical Reference* for the **Eclock** topology file names.

Use the **Eclock** command to set the switch's clock source for all switches.

For example, if your RS/6000 SP system has six node switch boards and four intermediate switch boards, select **/etc/SP/Eclock.top.6nsb.4isb.0** as an **Eclock** topology file. Enter:

```
Eclock -f /etc/SP/Eclock.top.6nsb.4isb.0
```

This command sets the proper clock source settings on all switches within a 96-way (6 nsb, 4 isb) RS/6000 SP system.

To verify the switch configuration information, enter:

```
splstdata -s
```

Step 14.6: Start the switch

Issue the **Estart** command to start the switch.

Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd** on SP Switch systems or **startsrc -s swtadmd2** on SP Switch2 systems before issuing the **Estart** command.

Deleting an adapter from a node

This procedure can be used to delete an adapter from a node.

Step 1: Archive the SDR

To save the current SDR attributes, run the command:

```
SDRArchive
```

This command saves SDR information for all system partitions.

Step 2: Disable nodes from the switch

If you do not have a switch in your SP system, or if you have already logically removed the SP Switch from your nodes, or if the node you are deleting the adapter from does not contain a switch, skip this step.

If you want to bring the switch down for all nodes, issue the **Equiesce** command for each system partition. If you use the **Equiesce** command, you must later restart the switch by issuing the **Estart** command. Issue the **Estart** command prior to "Step 8: Verify installation settings" on page 239.

The Switch Admin Daemon will issue the **Estart** command under certain circumstances, thus re-establishing switch recovery and primary node takeover. To see if you are using the Switch Admin Daemon, issue the following command:

```
lssrc -a | grep swtadmd
```

or

```
lssrc -a | grep swtadmd2
```

where **swtadmd** is used for an SP Switch and **swtadmd2** for an SP Switch2.

If the response returned shows that the subsystem is active, you may want to turn off the Switch Admin Daemon before issuing the **Equiesce** command. To turn off the daemon, issue the following command:

```
stopsrc -s swtadmd
```

or

```
stopsrc -s swtadmd2
```

where **swtadmd** is used for an SP Switch and **swtadmd2** for an SP Switch2.

If you are working with only a few nodes, you can use the **Efence** command to disable these nodes from the switch (if appropriate, first reassign the primary node or primary backup node). To determine if one of the nodes to be fenced is a primary or primary backup node, issue the **Eprimary** command. If you need to reassign the primary or primary backup node, issue the **Eprimary** command with appropriate options. Then issue the **Estart** command to make your choices take effect. You must then issue the **Efence** command to disable the nodes from the switch.

```
Efence -G node_number node_number ... node_number
```

Step 3: Reconfigure subsystems

If an IBM Virtual Shared Disk has been configured to use the adapters being removed, the IBM Virtual Shared Disk configuration must be modified. Refer to *PSSP: Managing Shared Disks* for more information on IBM Virtual Shared Disk configuration.

If the adapter being removed is defined to LoadLeveler in an adapter stanza in the LoadLeveler administration file, you will need to modify your LoadLeveler configuration. Refer to *IBM LoadLeveler for AIX 5L: Using and Administering* for more information.

Step 4: Shutdown the node

The node must be powered off before the adapter can be deleted. Stop the applications, then use SP Perspectives or the **cshutdown** command to shut down the node. For example, to shut down nodes 1 and 2:

```
cshutdown -G -N 1 2
```

Step 5: Update DCE information for a node (required for DCE)

This step is only required for DCE. If you ever plan to reuse the host name or IP address of the adapter being removed for another adapter in another machine, this step must be performed. If you do not plan to reuse the host name or IP address in another machine, this step is optional.

The following procedure assumes that the switch adapter is being removed and that the host names for the adapters follow a standard naming convention. You will need to tailor this procedure for your particular adapter and naming conventions.

If DCE is not configured in your SP system, this step can be skipped.

When DCE is installed on an SP system with a switch adapter, DCE sees the switch adapter as just another network adapter, and adds DCE principal and account entries for the switch adapter as it does for any other network adapter.

This step will delete those entries in the DCE registry database and in the DCE machine self-host (machine) key file on the node for which switch adapters were removed. Specifically, the `host/css_adapter_name` and `ftp/css_adapter_name` entries will be deleted.

If DCE switch entries were not created, or not created in both the DCE registry and in the DCE self-host key files on each host, one or both of the removal commands will fail.

If your site uses a common naming convention for switch adapter hostnames, you can determine if `host/` and `ftp/` entries for switch adapters exist in your DCE registry by using a combination of DCE and AIX commands.

For example, the following command string searches for the occurrence of the string `sn` in all of the principal names in the DCE database. `sn` is part of the naming convention that identifies switch adapter host names:

```
dcecp -c principal catalog | grep "sn"
```

The DCE principal catalog does not require any special privileges.

Using the switch adapter hostname naming convention in the previous example, the same type of check can be performed against the DCE self-host key files on each node.

Note: This check requires that the DCE command run as root on the target nodes.

```
dsh -avG "/usr/bin/dcecp -c keytab show self | grep sn" | dshbak -c
```

Note: Deleting switch `host/` and `ftp/` entries from the DCE registry and the DCE self-host key file on each node is optional. The existence of these entries in both or either location does not impact DCE functionality, nor does it create a security exposure, unless the principals associated with these entries were added to DCE access control lists in your DCE cell.

This procedure deletes entries from the DCE registry. The following examples assume that switch adapter host names must contain `sn`.

1. Login to the control workstation and then DCE login as a DCE cell administrator.

   Note: This step may be performed from another host configured with a DCE client in the same cell as the nodes DCE client.

2. Delete the switch-related entries.

   Delete the accounts first, then the principles, using either DCE **dcecp** commands or through DCE's SMIT interface.

   An example is:

```
dcecp -c account delete host/c186sn01.ppd.pok.ibm.com
dcecp -c account delete ftp/c186sn01.ppd.pok.ibm.com
dcecp -c principal delete host/c186sn01.ppd.pok.ibm.com
dcecp -c principal delete ftp/c186sn01.ppd.pok.ibm.com
```

The following examples assume that switch adapter hostnames must contain `sn`. The DCE commands to follow require that DCE be running on the node, that the invoker of the commands have the DCE self-host (machine) principal credentials, and that the commands be run on the node.

If you have the self-host principal credentials, the output of a **klist** command will contain a line starting with Ticket cache: and a value of:

```
/opt/dcelocal/var/security/creds/dcecred_ffffffff
```

The following examples show the sequence of steps required when a root user logs into a node:

1. Login to each node as a root user, but you may not need to issue a DCE login on the node because root has default access to the self-host principal credentials.
2. Obtain the complete name of the self-host key file object.

   ```
   dcecp -c keytab cat | grep self
   /.../c186dcecell/hosts/c186n01.ppd.pok.ibm.com/config/keytab/self
   ```

   Note: The unique portion of the name begins after the .../keytab/ string. The unique portion of the keytab object is a valid shorthand form of the object.
3. Display the switch entries in the self keytab object (key file).

   ```
   dcecp -c keytab show self | grep sn
   {/.../c186dcecell/host/c186sn01.ppd.pok.ibm.com des 1}
   {/.../c186dcecell/host/c186sn01.ppd.pok.ibm.com des 2}
   {/.../c186dcecell/ftp/c186sn01.ppd.pok.ibm.com des 1}
   {/.../c186dcecell/ftp/c186sn01.ppd.pok.ibm.com des 2}
   ```
4. Delete the switch entries using the data from Step 2 and Step 3.

   You do not need to issue a delete (**remove**) command for each version of the `ftp/` and `host/` entries that exist. However, you must delete both the `ftp/` and `host/` entries.

   ```
   dcecp -c keytab remove self -member ftp/c186sn01.ppd.pok.ibm.com
   dcecp -c keytab show self | grep sn
   {/.../c186dcecell/host/c186sn01.ppd.pok.ibm.com des 1}
   {/.../c186dcecell/host/c186sn01.ppd.pok.ibm.com des 2}
   dcecp -c keytab remove self -member host/c186sn01.ppd.pok.ibm.com
   ```

This completes the switch removal process.

Step 6: Delete switch adapters that connect nodes to the switch planes

Perform this step to delete switch adapters that connect the node from each switch plane. You must delete all switch adapters at this point to disconnect the node from every switch plane in the system.

```
spdeladap 5 1 0 css0
```

To validate the successful deletion of the adapter objects, issue:

```
splstdata -s
```

Step 7: Set node to customize

Each node that has an adapter being deleted must be customized. Use the **spbootins** commands to update the boot/install server's response to the **bootp** request from the nodes. If you are customizing nodes in more than one system partition, you must issue these commands in each system partition. For a complete description of the flags associated with these commands, refer to *PSSP: Command and Technical Reference*.

For example, to set the **bootp** response to `customize` for nodes 1 and 2, issue the following command:

```
spbootins -s no -r customize -l 1,2
```

Step 8: Verify installation settings

Make sure that the SDR has the appropriate values specified for each of the nodes being deleted. Issue the following command to display the values:

```
splstdata -G -b
```

The expected response is that nodes with adapters being deleted should be set to **customize**. All other attributes should contain the values last used to customize your nodes.

Step 9: Run setup_server to configure the changes

The **setup_server** command must be run to properly set up NIM on the control workstation by issuing the following command:

```
setup_server 2>&1 | tee /tmp/setup_server.out
```

The output will be saved in a log file called `setup_server.out`.

If you have a node defined as a boot/install server you must also run **setup_server** out on that node.

```
dsh -w boot/install_node "/usr/lpp/ssp/bin/setup_server \
    2>&1" | tee /tmp/setup_server.boot/install_node.out
```

Step 10: Refresh RSCT subsystems

The SDR has been updated to reflect the deleted adapters. You must now refresh the RSCT subsystems on the control workstation and all nodes to pick up these changes. Run **syspar_ctrl** on the control workstation to refresh the subsystems on both the control workstation and on the nodes:

```
syspar_ctrl -r -G
```

Step 11: Reboot the node

Once the adapter has been physically deleted in the node, the node can be rebooted to complete the logical configuration. If you have any boot/install servers in your system, you must customize them before customizing their clients. Reboot the nodes to be customized by using the **cstartup** command:

```
cstartup -G -N 1 2
```

The node has finished customization when the LEDs become blank, and its **host_responds** value is **yes**.

Verify that the **bootp_response** has been set to **disk** by issuing the following command:

```
splstdata -G -b
```

Step 12: Rejoin the nodes to the switch network

If you do not have a switch in your SP system, or if the node you are deleting the adapter from does not contain a switch, skip the remainder of this step. In SP Switch2 systems, you can have some nodes on the switch and some nodes off of the switch.

If you disabled all nodes from the switch using the **Equiesce** command, you must now issue the **Estart** command in each system partition to rejoin the nodes to the switch network. If you disabled only a few nodes using the **Efence** command, you must now issue the **Eunfence** command to bring those nodes back to the current switch network.

If you stopped the Switch Admin Daemon in “Step 2: Disable nodes from the switch” on page 235, start it now by issuing the command:

```
startsrc -s swtadmd
```

or

```
startsrc -s swtadmd2
```

where **swtadmd** is used for an SP Switch and **swtadmd2** for an SP Switch2.

Replacing a node with an equivalent node

Follow the steps in this section to replace a node in one of your frames with an equivalent node. Equivalent nodes have the same adapters and configuration data, and have the same node type. Note that if you are replacing a node with another node, you must also ensure that the microcode level of the node is correct. Refer to “Step 34: Update the state of the supervisor microcode” on page 60 for more information.

Replacing a node with an equivalent node means removing one node and adding another in its place. Before you do this, you need to determine what applications are currently running on the node. You may need to plan to transfer these applications to another node on your system. Also, if the node you want to delete serves as the boot/install server or as a primary node, you may need to transfer the responsibilities to one of the remaining nodes.

- Be sure to back up any local data you want to preserve.
- Do these steps in the system partition in which you replace the nodes (SP Switch only).

Step 1: Shut down the node and power it off

Prior to shutting down the node, make a list of all the applications and programs running on the node (such as LoadLeveler) and stop them from running.

Use Perspectives or the **cshutdown** command to shut down the nodes. For example, issue:

```
cshutdown -G -N 48
```

Step 2: Security considerations

Step 2.1: Secure remote commands

If running with a secure remote command method enabled, changing the host name or IP address might require a regeneration of public keys and *known_hosts* files, depending on your secure remote command configuration.

Step 2.2: Unconfigure DCE-related information for the node (required for DCE)

After the node has been deleted from the frame, you must remove any DCE-related principles and objects from the DCE registry. You must also unconfigure DCE (Admin portion). You must have cell administrator authority to perform this step.

On the control workstation, use the **rm_spsec -t admin** *node_dce_hostname* command first, then do a DCE Admin unconfigure for the node (**smit rmdce**).

To remove any additional principals related to the node using the SMIT panels, enter the host name of the principals to be deleted. For example, on the “Admin only unconfiguration for another machine” panel in the “Machine’s name or TCP/IP address” field, enter the host name for the additional adapters.

For the nodes being removed, verify that all DCE principals have been deleted from the DCE registry. Issue:

```
dcecp -c principal catalog -simplename
```

You must now issue:

```
setupdce -v

config_spsec -v
```

Note: To run **rm_spsec**, **setupdce**, and **config_spsec** off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being reconfigured. In addition, **rm_spsec** and **config_spsec** may require the **-r** (remote) flag. Refer to the **rm_spsec** and **config_spsec** commands in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.

Step 3: Replace the old node with the new one and power it on

Your IBM Customer Engineer (CE) performs this step.

When the node is powered on, the node supervisor contacts the frame supervisor and informs it that it is running. The frame supervisor updates the **hardmon** daemon, running on the control workstation, that a new node has been added to the system.

Step 4: Update the state of the supervisor microcode

To ensure that you have the latest level of microcode required by the hardware on your SP system, issue the **spsvrmgr** command. For example, to get the status in report form of all of your frames, nodes, and switches, enter:

```
spsvrmgr -G -r status all
```

To update the microcode of the frame supervisor of frame 3, enter:

```
spsvrmgr -G -u 3:0
```

Note: When using the **spled** application, a node in the process of having the microcode on its supervisor card updated will not be displayed in the window.

Refer to the *PSSP: Command and Technical Reference* for more information on using the **spsvrmgr** command.

Step 5: Unallocate NIM resources

Unallocate NIM Resources on the boot/install server for the node you are replacing by issuing the **unallnimres** command. For example:

```
unallnimres -l 33
```

Step 6: Delete the NIM client

Delete the NIM client on the boot/install server for the node you are replacing by issuing the **delnimclient** command. For example, to delete the NIM client for node number 3, issue:

```
delnimclient -l 3
```

Step 7: Acquire the hardware Ethernet address

This step takes the hardware Ethernet address for your new node (either from the node itself or from a file), puts it in the Node Object of the SDR, and sets up the **/etc/bootptab** file on the boot/install server.

Note: If you have information for the node being replaced in **/etc/bootptab.info**, be sure to remove the old address and then add the new Ethernet address.

Use the **sphrdwrad** command to write the hardware Ethernet address to the SDR. For example:

```
sphrdwrad -l 3
```

Step 8: Set up nodes to be installed

Use SMIT or issue the **spbootins** command to change the default boot/install information in the Node Objects in the SDR so that you can indicate a different bootp response for nodes to be installed.

For example:

```
spbootins -s yes -r install -l 3
```

This step also runs the **setup_server** command on all affected server nodes so that the nodes served are installed accordingly at power-up.

Note: In order to run the **spbootins -s yes** command and be authorized to perform a remote command to the target nodes, you must have SDR write authority and be authorized to perform an **rsh** to the target nodes. Therefore, your user ID must be in the appropriate authorization file (**.k5login**, **.klogin**, or **.rhosts**) on the target nodes.

Step 9: Network boot the new SP node

You can network boot the nodes using the following command:

```
nodecond frame_id slot_id &
```

Note: For MCA nodes, the **nodecond** command remotely processes information from the initial AIX firmware menus. You should not change the language option on these menus. The language must be set to English in order for the **nodecond** command to run properly.

Step 10: Run post-installation procedures

Set up the new nodes for any additional tools you use to manage your RS/6000 SP system environment.

See Table 4 on page 107 for a listing of information on the procedures for setting up these facilities.

Replacing a node with a different type of node

In order to replace a node with a different type of node, follow the steps in “Deleting a frame, node, SP-attached server, or clustered enterprise server” on page 229 to delete the node and then follow the steps in “Adding nodes” on page 202 to add a node.

Adding an SP expansion I/O unit to an existing node

Perform the following steps to add an SP expansion I/O unit to an existing node.

Step 1: Archive the SDR

Before adding an SP expansion I/O unit to an existing node, you should back up the SDR by issuing:

```
SDRArchive
```

Note the location and the name of the file created after you issue this command.

Step 2: Shut down the node and power it off

Prior to shutting down the node that you are adding the SP expansion I/O unit to, make a list of all the applications and programs running on the node (such as LoadLeveler) and stop them from running.

Use Perspectives or the **cshutdown** command to shut down the nodes. For example, issue:

```
cshutdown -G -N 45
```

Step 3: Install, cable, and power on the new SP expansion I/O unit

Your IBM Customer Engineer (CE) performs this step. See *RS/6000 SP: Installation and Relocation* for instructions.

Step 4: Power on the node

Power on the node using the **cstartup** command. For example:

```
cstartup -G -N 45
```

Step 5: Verify the SP expansion I/O unit configuration

To verify that the SP expansion I/O unit is properly configured in the SDR, issue:

```
splstdata -x
```

Removing an SP expansion I/O unit from an existing node

Perform the following steps to remove an SP expansion I/O unit from an existing node.

Step 1: Archive the SDR

Before removing an SP expansion I/O unit from an existing node, you should back up the SDR by issuing:

```
SDRArchive
```

Note the location and the name of the file created after you issue this command.

Step 2: Shut down the node and power it off

Prior to shutting down the node that you are deleting the SP expansion I/O unit from, make a list of all the applications and programs running on the node (such as LoadLeveler) and stop them from running.

Use Perspectives or the **cshutdown** command to shut down the nodes. For example, issue:

```
cshutdown -G -N 45
```

Step 3: Disconnect and remove the SP expansion I/O unit

Your IBM Customer Engineer (CE) performs this step. See *RS/6000 SP: Installation and Relocation* for instructions.

Step 4: Remove the SP expansion I/O unit definition from the SDR

Use the **spdelexp** command on the control workstation to remove an SP expansion I/O unit definition from the SDR. For example, issue:

```
spdelexp -x 4,5,7
```

Step 5: Power on the node

Power on the node using the **cstartup** command. For example:

```
cstartup -G -N 45
```

Moving an SP expansion I/O unit

Perform the following steps move an SP expansion I/O unit from one node to a different node.

Step 1: Archive the SDR

Before moving an SP expansion I/O unit from one existing node to another, you should back up the SDR by issuing:

```
SDRArchive
```

Note the location and the name of the file created after you issue this command.

Step 2: Shut down the nodes and power them off

Prior to shutting down the node that you are removing the SP expansion I/O unit from and the node that you are moving it to, make a list of all the applications and programs running on the node such as LoadLeveler and stop them from running.

Use Perspectives or the **cshutdown** command to shut down the nodes. For example, issue:

```
cshutdown -G -N 45 61
```

Step 3: Disconnect the SP expansion I/O unit from the old node and install it on the new node

Your IBM Customer Engineer (CE) performs this step. See *RS/6000 SP: Installation and Relocation* for instructions.

Step 4: Remove the SP expansion I/O unit definition from the SDR

Use the **spdelexp** command on the control workstation to remove the old SP expansion I/O unit definition from the SDR. For example, issue:

```
spdelexp -x 4,5,7
```

Step 5: Power on the node

Power on the nodes using the **cstartup** command. For example:

```
cstartup -G -N 45 61
```

Step 6: Verify the SP expansion I/O unit configuration

To verify that the SP expansion I/O unit is properly configured in the SDR, issue:

```
splstdata -x
```

Reconfiguring IBM @server pSeries 690 logical partitions (LPARs)

The IBM @server pSeries 690 is a multiprocessor server whose processors, memory, and I/O slots can be logically partitioned into separate machine images each running its own version of the operating system and SP software. The system resources are assigned to logical partitions (LPARs) through the Hardware Management Console (HMC). The server is attached to the SP control workstation through an SP administrative Ethernet LAN connection to the HMC. This connection is used by the PSSP Hardware Monitor (**hardmon**) to control and monitor the server hardware and to provide serial terminal access to each LPAR.

As an SP-attached server or clustered enterprise server, the pSeries 690 is represented as a frame. Each LPAR is represented in PSSP as a node. The frame slot number of the LPAR is equivalent to the partition ID assigned to that LPAR by the HMC. This slot number in turn is used to derive the node number following standard PSSP frame-slot number to node number conversion algorithms.

Adding a new LPAR

Adding a new LPAR to an existing pSeries 690 server attached to your SP or clustered enterprise server system is similar to adding a new node to your SP system. However, instead of having your CE physically connect the node hardware to the system, you will need to use the HMC interface to create the logical partition and assign resources to that partition. Once the partition has been created, **hardmon** will automatically find the new LPAR and a new node will be added to the SDR for that pSeries 690 frame. Follow the instructions in “Adding nodes” on page 202 noting the special instructions for pSeries 690 servers in “Step 2: Connect new nodes to the frame” on page 203.

Deleting an existing LPAR

Deleting an existing LPAR from a pSeries 690 server attached to your SP or clustered enterprise server system is similar to deleting a node from your SP system. However, instead of having the CE physically disconnect the node hardware from the system, you will need to use the HMC interface to remove the LPAR from the server. Once the partition has been removed, you will need to delete the node and related information from the SDR on the control workstation as you do for all other nodes. Follow the instructions in “Deleting a frame, node, SP-attached server, or clustered enterprise server” on page 229 noting the special instructions for pSeries 690 servers in “Step 6: Disconnect the node from the frame (optional)” on page 231.

Adding Resources to an LPAR

Adding resources to an existing LPAR on a pSeries 690 server attached to your SP or clustered enterprise server system is similar to adding additional hardware resources to a node on your SP system. When adding resources to a previously-defined partition, PSSP will be impacted as follows:

processors Adding processors to an LPAR will be detected by PSSP the next time the LPAR is booted. The PSSP node boot configuration code will detect the new number of processors and update the node information in the SDR.

memory Adding memory to an LPAR will not be detected by PSSP. This information is not stored in the SDR.

I/O Adding I/O slots to the LPAR may require additional PSSP node configuration steps. If you are adding adapters which will be supported and configured by PSSP, you will need to follow the steps for adding an adapter to your SP node in “Adding an adapter to a node” on page 225.

To add resources to an LPAR, you will need to use the HMC interface. The HMC interface can be started in one of the following ways:

- Login directly to the HMC subsystem. This will automatically start the GUI on the local HMC display.
- Manage the HMC subsystem through the remote WebSM session running on your control workstation or other UNIX workstation that has connectivity to the HMC.
- Launch a remote WebSM session to the HMC from the Hardware Perspective running on your control workstation.

Refer to the *Hardware Management Console for pSeries Operations Guide* for details on adding resources to an LPAR.

Deleting Resources from an LPAR

Deleting resources from an existing LPAR on a pSeries 690 server attached to your SP or clustered enterprise server system is similar to removing hardware resources from a node on your SP system. When deleting resources from a previously-defined partition, PSSP will be impacted as follows:

processors Removing processors from an LPAR will be detected by PSSP the next time the LPAR is booted. The PSSP node boot configuration code will detect the new number of processors and update the node information in the SDR. A message will be written to the

system log indicating that the number of processors detected is less than the number of processors configured for the node during the previous boot of the system.

memory Reducing memory on an LPAR will not be detected by PSSP. This information is not stored in the SDR.

I/O Removing I/O slots from an LPAR may require additional PSSP node configuration steps. If you are deleting adapters which were supported and configured by PSSP, you will need to follow the steps for deleting an adapter from your SP node in “Deleting an adapter from a node” on page 235.

To delete resources from an LPAR, you will need to use the HMC interface. The HMC interface can be started in one of the following ways:

- Login directly to the HMC subsystem. This will automatically start the GUI on the local HMC display.
- Manage the HMC subsystem through the remote WebSM session running on your control workstation or other UNIX workstation that has connectivity to the HMC.
- Launch a remote WebSM session to the HMC from the Hardware Perspective running on your control workstation.

Refer to the *Hardware Management Console for pSeries Operations Guide* for details on deleting resources from an LPAR.

Changing modes of operation

A pSeries 690 server has three separate modes of operation:

Full System Partition mode
Also known as SMP mode. The pSeries 690 server is configured as a single server with no logical partitioning. PSSP will recognize this server as a single frame with a single node, logically located in frame slot 1.

Logical Partition mode
This system is logically partitioned with processor, memory, and I/O resources distributed across from 1 to 16 logical machine images. PSSP recognizes this as a single frame with multiple nodes, one node for each LPAR, logically located in the frame slot corresponding to the partition ID for that LPAR.

Physical Partition mode
This system is logically partitioned across physical hardware boundaries to provide processor and memory affinity. To PSSP, this is identical to Logical Partition mode.

To change the mode of operation between Full Partition mode and either Logical Partition mode or Physical Partition mode on your pSeries 690 server, you will need to reconfigure your PSSP system. Perform the following steps:

1. Delete the original attached server frame and nodes. Follow the instructions in “Deleting a frame, node, SP-attached server, or clustered enterprise server” on page 229.
2. Change the mode of operation. This must be done using the HMC interface. Since the frame has been deleted from PSSP, you will not be able to use the Hardware Perspective to launch a remote WebSM session to the HMC. You can start the HMC interface in one of the following ways:

- Login directly to the HMC subsystem. This will automatically start the GUI on the local HMC display.
- Manage the HMC subsystem through a remote WebSM session running on your control workstation or other UNIX workstation that has connectivity to the HMC.

Refer to *Hardware Management Console for pSeries Operations Guide* for details on changing the mode of operation for your pSeries 690 server.

3. Add the new attached server frame and nodes. Follow the instructions in "Adding a frame, SP-attached server, or clustered enterprise server" on page 195 and "Adding nodes" on page 202.

Adding a switch

The RS/6000 SP system supports three types of switches. They are listed in the following table:

Switch Type	Description	Feature Code
Scalable POWERparallel Switch2 (SP Switch2)	The SP Switch2 has 16 ports for node connections and 16 ports for switch-to-switch connections. The SP Switch2 can interconnect all of the processor nodes that are currently available from IBM. When the control workstation and all of the nodes are running PSSP 3.4, you have optional switch connectivity. You can connect a selection of nodes to the SP Switch2 and leave some nodes off the switch. The SP Switch2 supports two different switch configurations so that nodes can communicate over one or two switch planes. A two plane SP Switch2 system has two sets of switches and two switch adapters per node. The switch planes are disjoint - each is cabled exactly like a single plane, both with the same topology, and communication across the pair of planes is achieved by a data striping algorithm in the software. You can only use this option with nodes that can have two adapters. You cannot install one adapter in some nodes and two adapters in other nodes. The SP Switch2 can be configured with one switch plane, with one adapter per node attached to the one switch plane network. The SP Switch2 supports multiple node switch boards in one frame. With PSSP 3.4, you can have up to eight SP Switch2 node switch boards in a single SP frame. Multiple node switch boards are used in systems that require a large number of switch connections for SP-attached servers or clustered enterprise server configurations. The PSSP 3.4 software supports the removable and hot-pluggable interposer cards that provide the bulkhead connections to the switch cable, the concurrent replacement of any failed power supplies or fans, and replacement of the supervisor while the switch is operating. The SP Switch2 does not support SP system partitioning or the SP Switch Router.	4012
Scalable POWERparallel Switch (SP Switch)	This switch can interconnect all the currently-orderable SP processor nodes, SP Switch Routers, and SP-attached servers. It has 16 ports for node connections and 16 ports for switch-to-switch connections.	4011
Scalable POWERparallel Switch-8 (SP Switch-8)	This switch offers 8 internal connections to provide enhanced functions for small systems with up to 8 total nodes. It does not support POWER3 SMP high nodes, SP-attached servers, or scaling to larger systems.	4008

If you are adding any of these switches, your IBM Customer Engineer (CE) installs the switch hardware on your RS/6000 SP system. You must plan your switch configuration and record your information on the Switch Configuration Worksheet, located in the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*. Both the *RS/6000 SP: Planning, Volume 1, Hardware and Physical*

Environment and the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* offer information to help you with the worksheet. After you complete the worksheet, follow these steps to install and configure the switch according to your configuration worksheet.

Note: The extension node only communicates with the SP Switch and the SP Switch-8. If you have multiple partitions, and you want the nodes in each partition to use the extension node, you must add an extension node in each partition.

Adding a switch to a switchless system

Note: If your system contains SP-attached servers, you must ensure that the switch port numbers assigned to them will be valid after the switch is configured. The system's switch port numbers may change during configuration. Refer to the "Understanding node numbering and switch port numbering" section of the "Defining the configuration that fits your needs" chapter in *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*

Step 1: Redefine the system to a single partition

Refer to the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for more information.

Step 2: Install the Communication Subsystem software

Install the level of Communication Subsystem software (**ssp.css**) on the control workstation that corresponds with the PSSP code version on the control workstation. You will also need the **ssp.st** file set installed for Job Switch Resource Table services.

Use SMIT or **installp** to install **ssp.css** on the control workstation. For more information, see "Step 19: Install PSSP on the control workstation" on page 31.

Step 3: Install the new switch

Your IBM Customer Engineer (CE) performs this step. This step may include installing the switch adapters and installing a new frame supervisor card.

Step 4: Initialize switch information

Issue the following command:

```
Eprimary -init
```

Step 5: Configure the switch adapters for each node

Use SMIT or the **spadaptrs** command to create **css0** adapter objects in the SDR for each new node. If you are adding two switch planes, you will need to create **css1** adapter objects also. Refer to your Switch Configuration Worksheet. Use this data to configure the switch on the nodes.

See "Step 40: Configure additional adapters for nodes" on page 68 for more information.

DCE notes:

1. All adapters must have an **ftp** and a **host** account defined in the DCE database.
2. If you are adding an additional adapter to a previously DCE-configured node, perform the following steps:
   a. Login to the control workstation as a cell administrator

b. Run **kerberos.dce -type admin -ip_name** *hostname_of_adapter*

Refer to *IBM Distributed Computing Environment for AIX: Administration Commands Reference* for more information on the **kerberos.dce** command.

Step 6: Configure the aggregate IP interface for nodes (SP Switch2 only)

Refer to “Step 41: Configure the aggregate IP interface for nodes (SP Switch2 only)” on page 73 for more information.

Step 7: Update the state of the supervisor microcode

Refer to “Step 34: Update the state of the supervisor microcode” on page 60 for more information.

Step 8: Update the SDR

To update the SDR switch information, issue the following command:

```
/usr/lpp/ssp/install/bin/hmreinit
```

Step 9: Set up the switch

Do installation steps “Step 62: Set up the switch” on page 92 through “Step 63: Verify the switch primary and primary backup nodes” on page 95. If you have an SP Switch, you must do “Step 64: Set the switch clock source for all switches (SP Switch only)” on page 96 and you may do “Step 65: Set up system partitions (SP Switch or switchless systems only)” on page 97. If you have set up system partitions, you need to reconfigure them.

Step 10: Refresh RSCT subsystems

To refresh the subsystems, issue the following command on the control workstation after adding the switch:

```
syspar_ctrl -r -G
```

This command refreshes all the subsystems in each partition such as **hats** so that it recognizes the new switch.

Step 11: Set up nodes to customize

Previously, you installed the Communication Subsystem Software (CSS) on the control workstation. Now, you need to install this software on all the nodes in the system. Use SMIT or issue the **spbootins** command to do this by changing the boot/install information for all the Node Objects in the SDR to specify a bootp response of **customize**.

Set the nodes to **customize** with the following command:

```
spbootins -r customize -l node_list
```

Step 12: Reboot all nodes

1. Reboot all the nodes for node customization.
2. When the nodes have booted and initialized, validate that **hostResponds** is active for the nodes.

Step 13: Start up and verify the switch

1. To complete this step, follow the installation steps starting at “Step 76: Start the optional switch” on page 103 and proceed through the rest of the chapter.
2. To restart Perspectives, enter the following command:

   ```
   perspectives &
   ```

Adding a switch to a system with existing switches

Step 1: Redefine the system to a single partition (SP Switch or switchless systems only)

Refer to the *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for information on setting up partitions.

Step 2: Install the new switch

Your IBM Customer Engineer (CE) performs this step. This step includes installing the switch adapters and installing new frame supervisors.

Step 3: Update the state of the supervisor microcode

Refer to “Step 34: Update the state of the supervisor microcode” on page 60 for more information.

Step 4: Configure the adapters for each node

Use SMIT or the **spadaptrs** command to create **css0** adapter objects in the SDR for each new node. If you are adding two switch planes, you will need to create **css1** adapter objects also. Refer to your Switch Configuration Worksheet. Use this data to configure the switch on the nodes.

See “Step 40: Configure additional adapters for nodes” on page 68 for more information.

DCE notes:

1. All adapters must have an **ftp** and a **host** account defined in the DCE database.
2. If you are adding an additional adapter to a previously DCE-configured node, perform the following steps:
   a. Login to the DCE cell as a cell administrator
   b. Run **kerberos.dce -type admin -ip_name** *hostname_of_adapter*

Refer to *IBM Distributed Computing Environment for AIX: Administration Commands Reference* for more information on the **kerberos.dce** command.

Step 5: Configure the aggregate IP interface for nodes (SP Switch2 only)

Refer to “Step 41: Configure the aggregate IP interface for nodes (SP Switch2 only)” on page 73 for more information.

Step 6: Set up the switch

Do installation steps “Step 62: Set up the switch” on page 92 through “Step 63: Verify the switch primary and primary backup nodes” on page 95. If you have an SP Switch, you must do “Step 64: Set the switch clock source for all switches (SP Switch only)” on page 96 and you may do “Step 65: Set up system partitions (SP Switch or switchless systems only)” on page 97. If you have set up system partitions, you need to reconfigure them.

Step 7: Refresh RSCT subsystems

To refresh the subsystems, issue the following command on the control workstation:

```
syspar_ctrl -r -G
```

This command refreshes all the subsystems in each partition such as **hats** so that it recognizes the new switch.

Step 8: Set up nodes to customize

All nodes in the frames supported by the switch being added must be recustomized. Use SMIT or issue the **spbootins** command to do this by changing the boot/install information for all the Node Objects in the SDR to specify a bootp response of **customize**.

Set the nodes to **customize** by issuing the following command:

```
spbootins -r customize -l node_list
```

Step 9: Reboot all nodes

1. Reboot all of the nodes in the frames supported by the switch being added.
2. Verify the production partition. When the nodes have booted and initialized, the production partition is ready for use.

Step 10: Start up and verify the switch

1. To complete this step, follow the installation steps starting at “Step 76: Start the optional switch” on page 103 and proceed through the rest of the chapter.
   - If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd2**.
2. To restart Perspectives, enter the following command:

   ```
   perspectives &
   ```

Upgrading the switches in your system

This section discusses how to reconfigure your system to remove a switch and replace it with a new switch. It does not discuss how to install a switch initially into your system. For this information, refer to “Chapter 2. Installing and configuring a new RS/6000 SP system” on page 11 or “Adding a switch” on page 248.

Note: To successfully run this procedure, you **must** replace **all** SP Switch switches in your system with SP Switch2 switches at the same time. In addition, each SP Switch2 switch **must** be placed in the locations that each SP Switch previously occupied. You cannot remove any SP Switch and replace it with an SP Switch2 placed in an alternate location.

If you need to move the switches in your system to locations other than where the SP Switch switches were located, you will first need to perform this procedure and place the SP Switch2 switches into the locations previously occupied by the SP Switch switches. After this procedure is complete, you can then delete and add frames as appropriate to change the locations of your switches.

Prerequisites to transferring switches

If you are removing a High Performance Switch and replacing it with the SP Switch, you need to do the following:

- Ensure your entire SP system has been migrated to at least PSSP 2.4. For instructions, refer to “Chapter 4. Migrating the software on your RS/6000 SP system” on page 127.
- Ensure that you have the PDU or SEPBU twin tailed frame supervisor card with the appropriate supervisor card level installed.

Prior to removing an SP Switch and replacing it with the SP Switch2, you need to do the following:

- If your nodes do not support the SP Switch2, you must first upgrade them before continuing.
- Ensure that all nodes have been migrated to at least PSSP 3.2 or later.

Step 1: Prepare for the switch transfer

IBM suggests that your system be configured as a single partition. If your system is a single partition, skip this step and proceed to the next step. If your system has multiple partitions, continue with this step.

Step 1.1: Archive the SDR

Always archive the SDR before partitioning (or repartitioning) your system. If you change your mind once you have committed a system partition configuration, or if applying a system partition configuration fails, you should use the archived SDR to recreate the previous system partition configuration. (To do this, select the Restore System Partition Configuration option from the SMIT menu.)

If using:	Do this:
SMIT	At the System Partition Configuration menu: **SELECT** Archive System Data Repository
SDRArchive	Enter `SDRArchive`
Note: When you archive the SDR from either SMIT or by using the **SDRArchive** command, the archive produced is in tar format. Issuing **tar -x** for this archive does not restore the system partition configuration. To restore the system partition configuration, always use the SMIT Restore System Partition Configuration option.	

Step 1.2: Return to a single partition environment

Either restore a previous single partition or set up a single partition.

To Restore a Single Partition:		To Set Up a Single Partition:	
Using SMIT		**Using SMIT**	
TYPE	smit syspar	**TYPE**	smit syspar
SELECT	Restore System Partition Configuration	**SELECT**	Select System Partition Configuration
ENTER	The name of a previously archived system partition configuration or press **List** for a list of choices.		• A list of system partition configurations appears.
PRESS	Enter to apply the new configuration.	**SELECT**	A single system partition
Using the command line		**Using the command line**	
The full path name of the **sprestore_config** command is: `/usr/lpp/ssp/bin/sprestore_config` *`archived_file`*		Issue the **spapply_config** command.	

Note: IBM suggests that your system be in one partition when transferring switches. If this is the case, the IBM Customer Engineer (CE) can use the control workstation to diagnose the switch and determine that it is functioning correctly. If there are multiple partitions, however, the IBM Customer Engineer (CE) cannot diagnose the switch completely using the control workstation. In this event, ask your IBM Customer Engineer (CE) to request that the laptop tool be shipped with the SP Switch Miscellaneous Equipment Specifications (MES).

Step 2: Remove existing css0 device entries from the ODM

To remove existing css0 device entries from the ODM, issue the following command:

```
dsh -a /usr/sbin/rmdev -l css0 -d
```

Step 3: Shut down all nodes and clean up the control workstation

1. Check for any unique switch topology files on nodes and save them.
   a. Issue the following command:
      ```
      dsh -a ls -al /etc/SP/expected.top
      ```
   b. Move any files you find to a new location. For example, (**/tmp/expected.top**).
2. Shut down the nodes using the **cshutdown** command. For example:
   ```
   cshutdown -G -N 3
   ```

Step 4: Replace a switch and switch adapters

This step may involve replacing the switch, switch adapters, switch cables, frame supervisors, and for some node types, the OCS modules. For more information regarding the instances when OCS module replacement is necessary, refer to the switch Miscellaneous Equipment Specifications (MES) that you received when you ordered the switch.

Your IBM System Engineer will perform this step for you using the instructions found in the Miscellaneous Equipment Specifications (MES) that you received when you ordered the switch.

Step 5: Update the SDR

To update the SDR switch information, issue the following command:

```
/usr/lpp/ssp/install/bin/hmreinit
```

Step 6: Update the state of the supervisor microcode

Refer to “Step 34: Update the state of the supervisor microcode” on page 60 for more information.

Step 7: Initialize switch information (SP Switch2 only)

Issue the following command:

```
Eprimary -init
```

Step 8: Annotate a switch topology file

Use **Eannotator** to update the switch topology file’s connection labels with their correct physical locations. Use the **-O yes** flag to store the switch topology file in the SDR. Using **Eannotator** makes the switch hardware easier to debug because the switch diagnostics information is based on physical locations.

For example, to annotate a two-switch or maximum 32-node system, enter:

```
Eannotator -F /etc/SP/expected.top.2nsb.0isb.0 \
           -f /etc/SP/expected.top.annotated -O yes
```

SP Switch note:

If you are replacing a High Performance Switch with an SP Switch and your system has multiple partitions, you must run the **Eannotator** command against all switch topology files that are in use for any system partition.

Storing the switch topology file in the SDR

If you entered **Eannotator -O yes** or **yes** previously when issuing the **Eannotator** command, skip this step.

Use **Etopology** to store the switch topology file in the SDR and make sure that it has been annotated. For example, to store a two-switch or maximum 32-node configuration, enter:

```
Etopology /etc/SP/expected.top.2nsb.0isb.0.annotated
```

SP Switch note:

If you are replacing a High Performance Switch with an SP Switch and your system has multiple partitions, you must run the **Etopology** command against all switch topology files that are in use for any system partition.

Step 9: Set the switch primary and primary backup nodes

Frame 1, node 1 is the default primary node.

If you have an SP Switch or SP Switch2 system, the primary backup node takes over for the primary node when it detects that the primary node is no longer functional. By default, a node is selected from a frame that is different from the primary node. If no other frame exists (for example, a single frame system), a node is selected from a switch chip that is different from the primary node. If no other switch chip is available, any available node on the switch is selected. From the command line, use the **Eprimary** command to verify this node or change the primary or primary backup to another node. For example:

```
Eprimary 1 -backup 16
```

This command, without any parameters, returns the node number of the current primary node, the primary backup node, the oncoming primary node, and the oncoming primary backup node.

Step 10: Set the switch clock source for all switches (SP Switch only)

Use the **Eclock** command to initialize the switch's Clock Source. The **Eclock** command requires that you know the number of Node Switch Boards (NSBs) and Intermediate Switch Boards (ISBs) in your RS/6000 SP system.

Select the **Eclock** topology file from the control workstation's **/etc/SP** subdirectory, based on these numbers. For example, if your RS/6000 SP system has six node switch boards and four intermediate switch boards, you would select **/etc/SP/Eclock.top.6nsb.4isb.0** as an **Eclock** topology file.

```
Eclock -f /etc/SP/Eclock.top.6nsb.4isb.0
```

This command sets the proper clock source settings on all switches within a 96-way (6 nsb, 4 isb) RS/6000 SP system.

Note: Be careful when using **Eclock** after the switch is initialized. **Eclock** modifies the switch clocking and is disruptive to the entire system (all system partitions).

See the *PSSP: Command and Technical Reference* for the **Eclock** topology file names.

Step 11: Set up nodes to customize

Since there was a switch topology change, you need to install this software on all the nodes in the system. Use SMIT or issue the **spbootins** command to do this by changing the boot/install information for all the Node Objects in the SDR to specify a bootp response of **customize**.

Set the nodes to **customize** with the following command:

```
spbootins -r customize -l node_list
```

Step 12: Power on the node

Power on the node using the **cstartup** command. For example:

```
cstartup -G -N 48
```

Step 13: Verify SP switch adapters

To verify that the switch adapters are functioning correctly, issue either of the following commands.

For the SP Switch adapter, issue:

```
SDRGetObjects switch_responds
```

For the SP Switch2 adapter, issue:

```
SDRGetObjects Adapter adapter_type==css0
```

If *adapter_config_status* shows other than **css_ready**, such as **diag_fail**, contact your IBM service representative.

Step 14: Start the switch

This step initializes the optional switch. Perform this step from the command line by issuing:

```
Estart
```

Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd** on SP Switch systems or **startsrc -s swtadmd2** on SP Switch2 systems before issuing the **Estart** command.

Check the switchResponds (SP Switch) or the switchResponds0 (SP Switch2) indicator for each node.

Step 15: Run a verification test on the switch

Run a verification test to ensure the installation of the switch is complete. You can do this using the command line. For example:

```
CSS_test
```

If the verification test fails, see the section on “Diagnosing switch problems” in the *PSSP: Diagnosis Guide*.

Step 16: Reapply your system partition configuration (SP Switch only)

This is an optional step to perform only if you have multiple partitions. To reapply your original system partition configuration, follow Steps 2 through 6 of "Partitioning the SP system" in the "Managing system partitions" chapter of the *PSSP: Administration Guide*.

Adding a switch plane (SP Switch2 only)

This section discusses how to reconfigure a one plane system to a two plane system.

Step 1: Quiesce your system

Before adding a switch plane, be sure to quiesce your system as follows:

- All users are logged off nodes
- No jobs are running
- Batch submission queues, such as LoadLeveler queues, should be shut down or quiesced and the node should be removed from the queue.
- If you are using the Switch Admin daemon for node recovery, stop it by issuing **stopsrc -s swtadmd2**.
- Quiesce the switch using the **Equiesce** command.

Step 2: Install the new switch plane hardware

Your IBM Customer Engineer (CE) performs this step. This step includes installing new switches, switch adapters, and switch cables.

Step 3: Update the state of the supervisor microcode

Refer to "Step 34: Update the state of the supervisor microcode" on page 60 for more information.

Step 4: Configure the adapters for each node

Use SMIT or the **spadaptrs** command to create css1 switch adapter objects in the SDR for each node. Refer to your Switch Configuration Worksheet. Use this data to configure the plane of switches on the nodes.

See "Step 40: Configure additional adapters for nodes" on page 68 for more information.

DCE notes:

1. All adapters must have an **ftp** and a **host** account defined in the DCE database.
2. If you are adding an additional adapter to a previously DCE-configured node, perform the following steps:
   a. Login to the control workstation as a cell administrator
   b. Run **kerberos.dce -type admin -ip_name** *hostname_of_adapter*

Refer to *IBM Distributed Computing Environment for AIX: Administration Commands Reference* for more information on the **kerberos.dce** command.

Step 5: Configure the aggregate IP interface for nodes (optional)

Refer to "Step 41: Configure the aggregate IP interface for nodes (SP Switch2 only)" on page 73 for more information.

Step 6: Reconfigure the number of switch planes on your system

To change the number of switch planes defined on your system, issue:

```
spswplane -p 2
```

Step 7: Set up the new switch plane

Do installation steps “Step 62: Set up the switch” on page 92 through “Step 63: Verify the switch primary and primary backup nodes” on page 95.

Step 8: Refresh RSCT subsystems

To refresh the subsystems, issue the following command on the control workstation:

```
syspar_ctrl -r -G
```

This command refreshes all the subsystems in each partition such as **hats** so that it recognizes the new switch.

Step 9: Set up nodes to customize

Perform this step to configure the new adapters on the nodes and to set up the nodes to use the new switch plane. Use SMIT or issue the **spbootins** command to do this by changing the boot/install information for all the Node Objects in the SDR to specify a bootp response of **customize**.

Set the nodes to **customize** by issuing the following command:

```
spbootins -r customize -l node_list
```

Step 10: Reboot all nodes

Reboot all the nodes.

Step 11: Start up and verify the switch

Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd2** before issuing the **Estart** command.

To complete this step, perform the following installation steps:

“Step 76: Start the optional switch” on page 103

“Step 77: Verify that the switch was installed correctly” on page 104

“Step 79: Tune the network adapters” on page 105

Chapter 7. Performing software maintenance

This chapter provides information on updating installation images, LPs, and support programs.

Updating and maintaining installation images

This section addresses maintaining mksysb installation images.

Adding mksysb images to the control workstation

Changes in system configurations may require you to create additional mksysbs or replace existing ones. To create additional mksysbs, use one of the following commands:

- mksysb — used with AIX 4.x and AIX 5L 5.1 systems
- smit mksysb — supported SMIT interface

Store the resulting mksysb on the control workstation using a different name from the mksysbs that already exists on the control workstation. If you do not use a different name, you will overwrite the existing mksysb and your changes will not get propagated to boot/install servers or nodes.

To replace an existing mksysb, create a new one and store in on the control workstation using the same name as the existing mksysb.

Note: In order to reinstall your nodes, the mksysb image and the lppsource that you use must both contain the same version, release, modification, and fix levels of AIX. If you do not have a mksysb image at the same level as your lppsource, you may do one of the following:

1. Make your own updated mksysb image. In order to do this, you will need to:
   a. Update an existing lppsource to the most recent maintenance level of AIX.
   b. Perform a BOS node upgrade on a single node as described in “BOS node upgrade” on page 153 **or** follow the steps in “Installing updates on a per node basis” on page 264.
   c. Make a mksysb image of that node as described in “Installing updates through reinstallation” on page 266.
   d. Use the mksysb created in Step 1c along with your updated lppsource to install your remaining nodes.
2. Contact IBM Level 1 service to obtain an updated mksysb image.

Restoring the control workstation from a mksysb image

When you need to restore the control workstation from a mksysb image, for example in the case of a disk in the rootvg volume group being replaced, follow the procedure documented in the “Installing BOS for a system backup” chapter of the *AIX Installation Guide*. Make sure that the mksysb image you are using is the most recent copy from the control workstation. Otherwise, you may have a version that does not match with the version on the nodes.

Restoring the node from a mksysb image

To restore the node from a mksysb image, refer to “Test your image on a single node” on page 269. Make sure that the mksysb you want to restore on the node is the correct one for that node.

Installing program updates

This section provides instructions for applying software updates (PTFs) to the SP system.

Before you begin

You need to be aware of the following items before you apply PTFs to the SP system.

PSSP READMEs

Make sure you read the *README* document that comes with any updates to PSSP. This information can also be viewed by running **installp -i** on the installation images. This document may convey important information that you need to know prior to installing the PTF. It may also contain instructions for activating particular fixes. This and additional information can also be printed to the screen during PTF installation.

Working with nodes that are down

Nodes that were down when service was applied must be updated when they become available. Simply follow the same procedure you used when updating the rest of the nodes.

Updating the css file set

When reinstalling or updating the **ssp.css** file set of PSSP, you must reboot all affected nodes to load changes that affect the kernel extensions.

Choosing an approach

Note: Performing this task requires that your identity be authenticated as an authorized user of the system management commands and the Perspectives interface shown in the following steps.

For DCE, you should **dce_login** to the SP administrative principal created in “Step 22.3: Create SP administrative principals” on page 45.

For Kerberos V4, you should use the principal created in “Step 21: Initialize RS/6000 SP Kerberos V4 (optional)” on page 38.

There are two approaches to installing program updates.

Approach	Description
Per-node	Apply the maintenance on each node individually. For the per-node approach, you can apply service on all the nodes in one of the following ways: • Logging in to each node and using SMIT for installation • Running the installation command on each node using **dsh**
Reinstall	Install the maintenance on a single node, building an image on that node, and then propagate the changes to all other nodes by reinstalling that image on each node.

The reinstall approach requires you install the programs and updates on your maintenance test node using SMIT or **installp** in order to generate an installation image and then place the new image on the control workstation. SMIT or the **spbootins** command enables you to specify all the nodes to be reinstalled with the new image. All that remains is to boot the nodes, which causes the nodes to be reinstalled.

Regardless of which approach you choose, do the following:

1. Apply the desired maintenance to a single test node. This allows you to gauge how long the service takes for a single node and enables you to verify the success of the maintenance before applying the service to the rest of the nodes.
2. Generate a mksysb image of your updated system. In the event of a required reinstall, you can use the mksysb image instead of reapplying the maintenance to any nodes that need to be reinstalled.

Which approach is right for you?

How do you know which approach to take? Consider these factors:

- How many nodes you have in your system.
- How long it takes to apply the desired maintenance.
- Whether you have user data on a node in its root volume group. (This data is destroyed on a reinstall.)

If you have fewer than 16 nodes in your system or the maintenance is minor, it may be faster to apply the maintenance directly on each node rather than to reinstall. On the other hand, if you have a large amount of maintenance to do and you have no user data to preserve in the root volume group, it may be faster to install the maintenance once, generate a new installation image, and reinstall all your nodes.

Preparing the control workstation

Regardless of your approach, you must install the maintenance of the control workstation first.

If you are using an HACWS configuration, before beginning, make sure that HACMP is running on both control workstations and that the primary control workstation is the active control workstation. Then perform Steps 1 through 4 on both the primary and backup control workstations and continue with Step 5. Once you start this procedure, you should not perform a control workstation failover at any time before Step 7.

1. Create a backup mksysb image of the control workstation.
2. Copy the PTFs into an appropriate directory on the control workstations, for example:

   `/spdata/sys1/install/pssplpp/`*`code_version`*`/ptf2`

   Note: Beware that the AIX update CD may contain PTFs for several levels of AIX. You need only to copy the PTFs pertaining to the levels of AIX installed on your system.
3. Run the **inutoc .** command in that directory to build the new **.toc** file.
4. To apply PTFs, you need write access to the SDR. Obtain DCE or Kerberos V4 credentials for the SP administrative principal as described in “Step 24: Obtain credentials” on page 46.
5. Apply the PTFs to the control workstation.

6. You may need to reboot the control workstation. Check the *README* in the PTF to see if this is required. If you have an HACWS and rebooting the control workstations is required, perform the preceding steps on both control workstations, then follow the instructions for rebooting in your HACMP documentation.

   If you have an HACWS configuration and rebooting was **not** required, you may want to recycle the control workstation applications to ensure that any fixes that affect HACMP/HACWS are enabled. Note that control workstation services will become unavailable during this procedure. To recycle the applications, first stop them on the backup control workstation using the following command. When the command completes, repeat it on the primary control workstation.

   ```
   /usr/sbin/hacws/spcw_apps -d
   ```

   Now restart the applications, first on the primary control workstation and then on the backup control workstation (note the order is the opposite of stopping them):

   ```
   /usr/sbin/hacws/spcw_apps -a
   ```

7. Verify that **/spdata/sys1/install/pssplpp** is exported to all nodes. (In an HACWS environment, it needs only to be exported on the active control workstation.)
8. Verify the correct operation of all SP and AIX control workstation functions.

Installing updates on a per node basis

This section outlines a procedure you can follow to install PTFs on your system.

Task A: Apply PTFs on one SP node and verify correct operation

Some PTFs require you apply them to all the nodes in your system. To do this, set up a test AIX 4.3.3 (SP Switch only) or AIX 5L 5.1 partition. See the *PSSP: Administration Guide* for instructions.

1. Select one node for PTF installation.
2. Create a backup mksysb image of the test node.
3. NFS mount **/spdata/sys1/install/pssplpp** from the control workstation onto that node:

   ```
   /usr/sbin/mount cw:/spdata/sys1/install/pssplpp /mnt
   ```

4. Apply PTFs to that node.
5. Reboot the node and verify correct operation.
6. Verify correct installation and operation of the node.
7. Create a mksysb image of this node and store it in an appropriate directory on the control workstation, for example:

   ```
   /spdata/sys1/install/images/bos.obj.ssp.43.ptf2
   ```

   Notes:

   a. You may want to install one node using the mksysb image you just saved to make sure the image is correct.

   b. If DCE is running on the host that the mksysb image is made from, you must first turn autostart off for the DCE daemons. To do this, issue:

      ```
      config.dce -autostart no
      ```

      then create the mksysb image.

Task B: Apply PTFs to all nodes

You can follow these instructions to install the PTF code from the directory onto each node. You can also install the PTFs by using the mksysb you just saved.

1. Using the **dsh** command, NFS mount **/spdata/sys1/install/pssplpp** onto each node. You can exclude the node you used for testing since it is now at the correct level.

   ```
   dsh -a /usr/sbin/mount cw:/spdata/sys1/install/pssplpp /mnt
   ```

2. Use **dsh** to run the appropriate **installp** command to apply PTFs to all nodes.

   ```
   dsh -a /usr/sbin/installp -Xagd /mnt/code_version ssp.st
   ```

   Note: The **dsh -f** option allows the **dsh** commands to be fanned out to multiple nodes.

3. Reboot the nodes and verify their correct operation.

Task C: Commit PTFs on the nodes and the control workstation

Committing the PTF will save file system space, but once the PTF is committed, you can never reject it. If you are not required to commit, you can skip this task.

1. Using **dsh**, commit the PTFs onto all nodes.
2. Commit the PTFs on the control workstation.

Task D: Update the state of the supervisor microcode

Refer to “Step 34: Update the state of the supervisor microcode” on page 60 for more information.

Task E: Update the SPOT when installing AIX BOS service updates

Perform the following steps on the control workstation and on all of the boot/install servers:

1. Deallocate the SPOT from all clients using the **unallnimres** command.
2. On the control workstation only, copy the install images from the AIX BOS Service Updates to the **lppsource** directory that corresponds to the appropriate SPOT. For example, the directory could be:

   ```
   /spdata/sys1/install/aix433/lppsource
   ```

3. For Boot Install Server (BIS) nodes, you must ensure that the BIS host name is in the **/.rhosts** file on the control workstation.
4. On the control workstation only, run **nim -o check** *lpp_source* to create a new **.toc** file.
5. Issue **smit nim_res_op**
   a. Select the appropriate SPOT.
   b. Select the **update_all** function.
   c. Hit **F4** in the “Source of Install Images” field and select the appropriate **lppsource**.
   d. Hit **Enter** twice to initiate the update.
   e. After the update completes, run **setup_server** to reallocate the SPOT to the necessary clients.
6. If you added **.rhosts** entries in Step 3, you can now delete them.

Note: In a multiple Boot/Install Server (BIS) environment, the following actions can only be performed on one BIS at a time due to an AIX constraint regarding the **inutoc** command and the **.toc** file:

1. Installing NIM master file sets
2. Creating the SPOT

For more information, see the “NIM errors in a multiple boot/install server (BIS) environment” section of the “Diagnosing NIM problems” chapter of the *PSSP: Diagnosis Guide*.

Task F: Create new mksysb images

Create a new backup mksysb image of the control workstation. The mksysb image you created earlier is now the mksysb image for all nodes. Store any earlier mksysb images you created before you installed the PTFs in case you need to restore your system to its previous maintenance level.

Installing updates through reinstallation

Performing this task requires that your identity be authenticated as an authorized user of the system management commands and the Perspectives interface shown in the following steps.

For DCE, you should **dce_login** to the SP administrative principal created in “Step 22.3: Create SP administrative principals” on page 45.

For Kerberos V4, you should use the principal created in “Step 21: Initialize RS/6000 SP Kerberos V4 (optional)” on page 38.

See the “Security features on the SP system” chapter in *PSSP: Administration Guide* for more information.

If you choose the reinstall approach, you may want to target particular nodes for maintenance images. For instance, if you have groups of nodes with distinct identities such as:

- Boot/install servers
- Compute nodes
- Interactive nodes

You may want to select one representative node in each group from which to apply maintenance and generate new installation images. The other nodes in the group should always then be reinstalled with the image generated from that node.

Note: For system partitions (SP Switch only):

1. All nodes using the switch must be at the same level of base AIX.
2. The communications subsystem software (**ssp.css**) must be at the same level on each node using the switch.

Build an installation image

Prior to creating a mksysb, you must do the following:

Notes:

1. Make sure that no files named **/etc/niminfo** or **/etc/niminfo.prev** exist. If they do exist, rename them. These files should be saved for possible debugging later on.
2. Verify the host name resolution for the control workstation and any nodes where the mksysb may be installed, if they are listed in the **/etc/hosts** file.
3. If you want your machine to be a NIM master, make sure that the image you are building from had not executed the **inurid -r** command. Check to see if this command was run by issuing the following:

   ```
   /usr/lib/instl/inurid -q ; echo $?
   ```

   If the return code=1, the **inurid -r** command was executed. IBM suggests that you not execute the **inurid -r** command on any machine in case you want to use the machine as a boot/install server in the future.

Before building an installation image, install all the LPs and service to the node on which you intend to create your installation image. Careful planning in selecting LPs and required service spares you from repeating this process.

Note: If DCE is running on the host that the mksysb image is made from, you must first turn autostart off for the DCE daemons. To do this, issue:

```
config.dce -autostart no
```

then create the mksysb image.

After you have installed the LPs and service required for your nodes, you are ready to make an installation image. Login to the node where you want to create an installation image and enter:

```
smit -C mksysb
```

Enter the file name of the image that you want to create and press Enter.

Note: The installation tools require that the name of this image begin with **bos.obj**. A suggested naming convention for these images is:

bos.obj.*level*.*date*

For example:

```
bos.obj.433.20000428
```

You can generate a mksysb image to install on your nodes. You can do this on a node after you have installed at least that node, or you can use a standalone workstation (either the control workstation or a different standalone workstation). However, after you have done “Step 21: Initialize RS/6000 SP Kerberos V4 (optional)” on page 38, you cannot use the control workstation for this purpose. Running the **setup_authent** command on any workstation makes the workstation unsuitable for generating a mksysb image for a node.

The control workstation (or a different standalone workstation) must be at the right level of the AIX RS/6000 operating system and must have all required PTFs applied. One way to ensure this is to install the mksysb image that resides in the **spimg installp** image on its product tape. You can also install any of the RS/6000 SP software options on that machine. After installing extra LPs or maintenance, you can generate a mksysb image of the system and copy it back to **/spdata/sys1/install/images** on the control workstation. You can then use that image to install your nodes.

When using a mksysb, we strongly suggest that any AIX corrective service applied to the mksysb should also be placed in the **lppsource** directory. The Shared Product Object Tree (SPOT) should also be updated. Refer to the procedure documented in “Task E: Update the SPOT when installing AIX BOS service updates” on page 265.

After you install a node, you can install any RS/6000 SP software options or LPs, other LPs, or any maintenance that you need. You can then test your node and generate a mksysb image of that node. Before doing this, you may want to remove certain files from the node that you do not want in the image. These include any mksysb images in that node’s **/spdata/sys1/install/images** directory (if the node is a boot/install server itself).

Do not remove **/home** from the node. If you do so, when you use the mksysb image to install a boot/install server node, you cannot create the **netinst** user ID that is required for network install after you install the image. After you create the mksysb image, copy it back to **/spdata/sys1/install/images** on the control workstation and install the rest of your nodes with that image.

Consider the following items prior to creating a mksysb backup:

- Excluding Mounted Filesystems:

  Are there any mounted filesystems from other systems? The mksysb command automatically backs up everything unless you specify what to exclude. Some example of items you might want to exclude are AFS, DFS, or NFS filesystems. You can exclude files by specifying the "-e" flag and creating an **/etc/exclude.rootvg** file with a list of the files to exclude. The file might contain something like this:

  ```
  /afs
  /usr/public
  ```

  where the **/afs** is an AFS filesystem and the **/usr/public** is a shared NFS filesystem that is mounted on this system.

  If you specify the **-e** flag with the **/etc/exclude.rootvg** file containing this information, mksysb image will not back up any files in these directories or create the directories.

- Generate a New **/image.data** File

  The **/image.data** file in the mksysb image contains information about the operating system level, the logical volumes, the filesystems, and other information. This information may be out of date. Therefore, you should always create the file during the mksysb creation using the **-i** option.

- Create MAP files

  Note: Do not use the **-m** flag on the mksysb command. This flag generates a layout mapping of the logical-to-physical partitions for each logical volume in the volume group. This mapping is used to allocate the same logical-to-physical partition mapping when the image is restored. Do not use this option because it requires that you have the exact same physical disk configurations when restoring the image. This option causes an installation failure if this requirement is not met. The installation process will stop with an LED code of C48 and the installation will automatically switch from a noprompt to a prompt mode.

- If you have additional node customizations that cannot be included in the SDR, you should update **/tftpboot/script.cust**, **/tftpboot/firstboot.cust**, and **/tftpboot/tuning.cust**. For example, customized parameters such as maxuproc, maxmbuff, rspoolsize, spoolsize, and asynchronous I/O should be preserved by adding the appropriate **chdev** or **chgcss** commands to these files. Otherwise, reinstallation of the mksysb image will restore the system default values.

  For more information on the node customization scripts, see "Appendix E. User-supplied node customization scripts" on page 297.

To create the mksysb image, use the following command to create a **/image.data** file. It will also expand **/tmp** if needed:

```
mksysb -i -X
```

The following command creates a **/image.data** file, expands **/tmp**, and excludes all files listed in **/etc/exclude.rootvg**.

```
mksysb -e  -i -X
```

Unconfigure DCE-related information for the node (required for DCE)

Note: You must have cell administrator authority to perform this step.

If a node was previously configured for DCE, you must remove any DCE-related principles and objects from the DCE registry before issuing the **nodecond** command.

1. On the control workstation, use the **rm_spsec -t admin** *node_dce_hostname* command.

   Note: To run this command remotely off of the SP, you must set the SP_NAME environment variable to point to the SDR you want to access. Refer to the **rm_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.

2. Do a DCE Admin unconfigure for the node (**smit rmdce**).

   Note: To remove any additional principals related to the node using the SMIT panels, enter the host name of the adapter to be deleted. For example, on the “Admin unconfiguration for another machine” panel in the “Machine’s name or TCP/IP address” field, enter the host name for the additional adapters.

3. For the nodes being removed, verify that all DCE principals have been deleted from the DCE registry. Issue:

   ```
   dcecp -c principal catalog -simplename
   ```

You must now create new DCE information for the node by performing the following steps:

1. Run the **setupdce** command.

   Notes:

   a. You will be prompted for the cell administrator’s password when you issue this command.

   b. To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being configured. The value must be a resolvable address. For example:

      ```
      export SP_NAME=spcws.abc.com
      ```

2. As an ID with cell administrator authority, run the **config_spsec -v** command.

   Note: To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being configured. Refer to the **config_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.

Test your image on a single node

After you create your netinstall image, you must test the image. Use **ftp** in binary mode or **rcp** to transfer the image to the control workstation. The installation tools require all installation images to reside on the control workstation. Remember that the image must be placed in the **/spdata/sys1/install/images** directory and that the permissions must allow it to be read by **other**.

To test the image, you must reinstall a node with this new image and run your applications to see if it meets your requirements.

Propagate your installation image

After you create the netinstall image and copy it to the control workstation in **/spdata/sys1/install/images**, you are ready to reinstall.

On the control workstation, issue the following:

```
spchvgobj -r selected_vg -i install_image_name \
          -l node_list

spbootins -r install -l node_list
```

In the following example:

```
spchvgobj -r selected_vg -i bos.obj.433.20000428 -l 23

spbootins -r install -l 23
```

would change the SDR information to install the image **bos.obj.433.20000428** on node 23 from its boot/install server.

This command automatically issues **setup_server** on node 23's boot/install server to update its install files with the new information. Note that **setup_server** copies the installation image from the control workstation to the appropriate boot/install servers if the boot/install server is not the control workstation. If the image has to be copied to a boot/install server, this may take some time. You can now use Perspectives to reset the node, which causes the netinstall of the test image on the node.

The boot list of a node is set by **/etc/rc.sp** to boot from **hdisk0** only. When the **bootp_response** of a node is changed from **disk** to be some other response, such as "install," the node is sent a boot list command to cause it to boot from **ent0** before **hdisk0**. If a node is down when its **bootp_response** is changed, its boot list is not changed. From the Hardware Perspective, select one or more nodes and then select Actions → Network Boot.

After the netinstall is complete, the node reboots and you can verify the image runs your applications. When you are satisfied that this image meets your requirements, you can use the **spbootins** and the **spchvgobj** commands to change the **install_image** attribute and **bootp_response** and reinstall the rest of your nodes.

You may need to override the physical partition size (PPSIZE) of the root volume group in the mksysb. See "Appendix F. Overriding the PPSIZE in a mksysb image" on page 301 for more information.

Chapter 8. Performing hardware maintenance

The information in this chapter contains steps that you must perform for hardware maintenance of your system.

Replacing the frame supervisor

To replace the frame supervisor with an updated frame supervisor, perform the following steps. As with many PSSP commands, you must have the appropriate authority or credentials to use them. See “Step 24: Obtain credentials” on page 46.

Step 1: Install the new frame supervisor card

Your IBM Customer Engineer (CE) performs this step.

Step 2: Reconfigure the hardware monitor to recognize the new hardware

Issue the following command to update the frame supervisors with the frame id:

```
hmcmds setid frame_number:0
```

For example:

```
hmcmds setid 1:0
```

Step 3: Update the state of the supervisor microcode

To ensure that you have the latest level of microcode required by the hardware on your SP system, issue the **spsvrmgr** command. For example, to get the status in report form of all of your frames, nodes, and switches, enter:

```
spsvrmgr -G -r status all
```

To update the microcode of the frame supervisor of frame 3, enter:

```
spsvrmgr -G -u 3:0
```

Refer to the *PSSP: Command and Technical Reference* for more information on using the **spsvrmgr** command.

Step 4: Issue Eannotator

Use **Eannotator** to update the switch topology file’s connection labels with their correct physical locations. Use the **-O yes** flag to store the switch topology file in the SDR.

For example, to annotate a two-switch or maximum 32-node system, enter:

```
Eannotator -F /etc/SP/expected.top.2nsb.01sb.0 \
           -f /etc/SP/expected.top.annotated -O yes
```

Step 5: Issue Eclock (SP Switch only)

Use the **Eclock** command to set the switch’s clock source for all switches. For example, if your RS/6000 SP system has six node switch boards and four intermediate switch boards, select **/etc/SP/Eclock.top.6nsb.4isb.0** as an **Eclock** topology file.

```
Eclock -f /etc/SP/Eclock.top.6nsb.4isb.0
```

This command sets the proper clock source settings on all switches within a 96-way (6nsb, 4isb) RS/6000 system.

Use the following command to verify the status in the supervisor cards:

```
hmmon -G -Q -v mux 1:17
```

Step 6: Issue Estart

Run **Estart** to restart the switch.

Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd** on SP Switch systems or **startsrc -s swtadmd2** on SP Switch2 systems before issuing the **Estart** command.

Replacing a fixed disk

This procedure applies only to disks that are part of the **rootvg** volume group.

- Performing this task requires that your identity be authenticated as an authorized user of the system management commands and Perspectives. See “Step 24: Obtain credentials” on page 46 for more information.
- If possible, back up any required data on the disk before replacing it.
- Do these steps in the system partition in which you replace the disk.

Step 1: Shut down the node

You can use Perspectives to power off the node.

If using:	Do this:	
Perspectives	**SELECT**	The **Hardware Perspective** icon by double clicking
	SELECT	The Nodes Pane • The Nodes Pane receives focus.
	SELECT	The node to power off
	PRESS	Power Off icon
	SELECT	The power off options
	PRESS	Apply
cshutdown	Issue: `cshutdown -G -N 48`	

Step 2: Replace the disk

Your IBM Customer Engineer (CE) performs this step. Make sure that the new disk has equivalent or greater disk space available.

Step 2a: Unconfigure DCE-related information for the node (required for DCE)

Note: You must have cell administrator authority to perform this step.

If a node was previously configured for DCE, you must remove any DCE-related principles and objects from the DCE registry before issuing the **nodecond** command.

1. On the control workstation, use the **rm_spsec -t admin** *dce_hostname* command.

   Note: To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being configured. Refer to the **rm_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.
2. Do a DCE Admin unconfigure for the node (**smit rmdce**).

   Note: To remove any additional principals related to the node using the SMIT panels, enter the host name of the adapter to be deleted. For example, on the “Admin unconfiguration for another machine” panel in the “Machine’s name or TCP/IP address” field, enter the host name for the additional adapters.
3. For the nodes being removed, verify that all DCE principals have been deleted from the DCE registry. Issue:

   ```
   dcecp -c principal catalog -simplename
   ```

You must now create new DCE information for the node by performing the following steps:

1. Run the **setupdce** command.

   Notes:

   a. You will be prompted for the cell administrator’s password when you issue this command.

   b. To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being configured. The value must be a resolvable address. For example:

      ```
      export SP_NAME=spcws.abc.com
      ```
2. As an ID with cell administrator authority, run the **config_spsec -v** command.

   Note: To run this command off of the SP, you must set the SP_NAME environment variable on the remote workstation to point to the SDR of the SP system being configured. Refer to the **config_spsec** command in *PSSP: Command and Technical Reference* for a description of the **-r** (remote) flag.

Step 3: Install the system image on the disk

Use the **spbootins** command to:

1. Set *bootp_response* to **install** using the **spbootins** command. For example:

   ```
   spbootins -s no -r install -l 33
   ```
2. Ensure that the installation information is set up properly by running the **splstdata** command. For example:

   ```
   splstdata -b -l 33
   ```
3. Run **setup_server** on the boot/install server.

Step 4: Network boot the node

Use Perspectives to network boot the node.

If using:	Do this:	
Perspectives	**SELECT**	The **Hardware Perspective** icon by double clicking
	SELECT	The Nodes Pane • The Nodes Pane receives focus.
	Select	Nodes to network boot
	PRESS	Actions → Nodes → Netboot
	PRESS	Apply
nodecond	Issue: `nodecond` *`frame_id slot_id`* `&`	

Replacing a disk not in rootvg

To replace disks that are *not* in the **rootvg** volume group:

Step 1: Shut down the node

You can use Perspectives to power off the node.

If using:	Do this:	
Perspectives	**SELECT**	The **Hardware Perspective** icon
	SELECT	The Nodes Pane • The Nodes Pane receives focus.
	SELECT	The node to power off
	PRESS	Power Off icon
	SELECT	The power off options
	PRESS	Apply
cshutdown	Issue: `cshutdown -G -N 48`	

Step 2: Replace the failing fixed disk

Your IBM Customer Engineer (CE) performs this step. Make sure that the new disk has equivalent or greater disk space available.

Step 3: Power on the node

If using:	Do this:
Perspectives	From **Hardware Perspective** menu bar **SELECT** The **Hardware Perspective** icon **SELECT** The Nodes Pane • The Nodes Pane receives focus. **SELECT** Node to be powered on **SELECT** Actions → Nodes → Power On... **SELECT** power on options **PRESS** Apply or from **Hardware Perspective** tool bar **SELECT** The **Hardware Perspective** icon **SELECT** The Nodes Pane • The Nodes Pane receives focus. **SELECT** Node to be powered on **PRESS** Power on icon from Tool bar **SELECT** Power on options **PRESS** Apply
cstartup	Issue: `cstartup -G -N 45`

Step 4: Restore any data previously backed-up

You will need to re-create the volume group and restore any files back into the user-based file system.

Replacing an I/O Planar card or an Ethernet adapter

Note: SDR adapter information includes the physical location code of the adapter. If you are placing the replacement adapter in a different location, you must run the **spadaptrs** command to update the physical location code in the SDR.

The I/O Planar card for nodes with an integrated Ethernet or the SP Ethernet administrative LAN adapter for other nodes has a hardware address used to network boot the nodes. This procedure updates the SDR hardware Ethernet address.

- A ticket is required to perform some of the commands in this procedure. Run the **k4init** or **dce_login** command.
- Do these steps in the system partition in which you replace the card.

Step 1: Shut down the node

You can use Perspectives to power off the node.

If using:	Do this:	
Perspectives	**SELECT**	The **Hardware Perspective** icon
	SELECT	The Nodes Pane • The Nodes Pane receives focus.
	SELECT	The node to power off
	PRESS	Power Off icon
	SELECT	The power off options
	PRESS	Apply
cshutdown	Issue: `cshutdown -G -N 48`	

Step 2: Replace the card

Your IBM Customer Engineer (CE) performs this step.

Step 3: Unallocate NIM resources

Unallocate NIM Resources on the boot/install server for the node you are replacing by issuing the **unallnimres** command. For example:

```
unallnimres -l 33
```

Step 4: Delete the NIM client

Delete the NIM client on the boot/install server for the node you are replacing by issuing the **delnimclient** command from the control workstation. For example, to delete node 3 from the boot/install server, issue:

```
delnimclient -l 3
```

Step 5: Obtain the new hardware Ethernet address

If you have Ethernet information for the node being replaced in **/etc/bootptab.info**, be sure to remove the old address.

Use the **sphrdwrad** command to get the new hardware Ethernet address.

Step 6: Re-create the NIM client definition

To re-create the NIM client definition, issue the **mknimclient** command. For example, to re-create the NIM client for node number 3, issue:

```
mknimclient –l 3
```

Changing Ethernet cable types

Note: If the adapter was originally specified using a physical location code and you are placing the replacement adapter in a different location, you must run the **spadaptrs** command to update the physical location code in the SDR.

To change an Ethernet cable type, perform the following steps. As with many PSSP commands, you must have the appropriate authority or credentials to use them. See “Step 24: Obtain credentials” on page 46.

Step 1: Power off the nodes of the cable being replaced

You can use Perspectives to power off the nodes.

If using:	Do this:
Perspectives	From the **Hardware Perspective** menu bar **SELECT** The **Hardware Perspective** icon **SELECT** The Nodes Pane • The Nodes Pane receives focus. **SELECT** The nodes to power off **PRESS** The Power Off icon **SELECT** The Power Off options **PRESS** Apply
cshutdown	Issue: `cshutdown -G -N 48`

Step 2: Replace the cable and recable your network

Your IBM Customer Engineer (CE) performs this step.

Step 3: Deallocate NIM resources

If any resources are allocated to the affected nodes, you must deallocate the NIM resources. For example, to deallocate resources for nodes 1, 2, and 3 from the boot/install server, issue:

```
unallnimres -l 1,2,3
```

Step 4: Delete the NIM clients

Delete the NIM clients on the boot/install server for the nodes you are replacing by issuing the **delnimclient** command from the control workstation. For example, to delete nodes 1, 2, and 3 from the boot/install server, issue:

```
delnimclient -l 1,2,3
```

Step 5: Change the Ethernet cable type entry in the SDR

To change the Ethernet cable type entry in the SDR, perform the following steps.

Note: You must reenter data that is not changing. If you do not know these values, you can obtain them by issuing the following two commands.

For the default route, issue:

```
splstdata -n
```

For the netmask, issue:

```
splstdata -a
```

If using:	Do this:	
SMIT	**TYPE**	smit enter_data
	SELECT	Node Database Information
	SELECT	SP Ethernet Information
	ENTER	Start frame, start slot, and node count
		OR
		Node group
		OR
		Node list
		AND
		Starting node's SP Ethernet administrative LAN adapter host name or IP address (for models other than the IBM @server pSeries 690, you must specify the en0 adapter name), netmask, default route host name or IP address.
	SELECT	Ethernet Adapter Type and type F4 to generate the list
	CHOOSE	The appropriate type from the list
	SELECT	**Ok**
spadaptrs	This example configures an SP Ethernet administrataive LAN adapter network of 16 nodes with IP addresses ranging from 129.33.32.1 to 129.33.32.16, a netmask of 255.255.255.192, and a default route of 129.33.32.200. The adapter is a twisted pair 10/100 Ethernet adapter using auto-negotiate for the communication transfer and rate. `spadaptrs -e 129.33.32.200 -t tp -d auto -f auto 1 1 16 en0 129.33.32.1 255.255.255.192`	

Note: If you are replacing your Ethernet adapter, you need to acquire the new hardware Ethernet address using the **sphrdwrad** command.

Step 6: Set the affected nodes to customize and run setup_server

Use the **spbootins** command to set the nodes to customize and run the **setup_server** command. For example, to set nodes 1, 2, and 3 to customize, issue:

```
spbootins -r customize -s yes -l 1,2,3
```

Note: In order to run the **spbootins -s yes** command, you must have SDR write authority and be authorized to perform an **rsh** to the target nodes. Therefore, your user ID must be in the appropriate authorization file (**.k5login**, **.klogin**, or **.rhosts**) on the target nodes.

Step 7: Power on the nodes

Power on the nodes. This causes the customization to take place that updates the adapters on the nodes.

Chapter 9. Installing the optional PSSP T/EC adapter

The information in this chapter is optional unless you plan on installing a PSSP Tivoli Enterprise Console (T/EC) adapter.

The PSSP T/EC adapter forwards events generated by the Event Management subsystem to a Tivoli Enterprise Console. The adapter consists of the **tecad_pssp** command and the **rvclasses.cfg** and **pssp_classes.baroc** configuration files. This product is offered as an optional file set, **ssp.tecad**, and is installed using installp. The PSSP T/EC adapter can be installed on any node or in the control workstation. IBM suggests that the PSSP T/EC adapter be installed in each node that will be generating events to be forwarded to the Tivoli Enterprise Console. By doing this, it is then possible to run the **tecad_pssp** command locally in each node, thus causing the event forwarding overhead to be distributed across the system.

Once the PSSP T/EC adapter is installed, the system administrator must complete the SP installation by running the **install_agent** command which places the event definitions in the SDR. This command can be run on any node in which the PSSP T/EC adapter is installed. Since the SDR is partition sensitive, the **install_agent** command must be run once in every system partition in which the adapter will be used.

The **install_agent** command is included in the **/usr/lpp/ssp/tecad** installation directory. It requires the name of the configuration file to be installed; which in the standard installation case is rvclasses.cfg (also in the installation directory). For example, the following command will install the SDR classes for the PSSP T/EC adapter:

```
/usr/lpp/ssp/tecad/install_agent /usr/lpp/ssp/tecad/rvclasses.cfg
```

The installation directory also contains a README file that contains information on how to add custom classes to the PSSP T/EC adapter. The use of custom classes is not supported and is, therefore, not included in this book.

The **tecad_pssp** command also requires a configuration file with the location of the T/EC server. Refer to *PSSP: Command and Technical Reference* for the format of this file and a description of the command.

The **ssp.tecad** file set also contains the **pssp_classes.baroc** class definition file. This file needs to be installed using the standard procedure for installing event adapters in T/EC. Refer to *Tivoli Enterprise Console Event Integration Facility Guide* for details.

Chapter 10. Installing extension nodes

The information in this chapter is optional unless you plan on installing an extension node.

Note: Extension nodes are only supported with SP Switch or switchless systems, not with the SP Switch2.

Extension nodes are non-standard nodes that extend your system's capabilities or scope, but cannot be used in all the same ways as standard SP nodes. Extension nodes are attached to the system via an *extension node adapter*.

A specific type of extension node is a *dependent node*. A dependent node depends on SP nodes for certain functions, but implements much of the switch-related protocol that standard nodes use.

You can use the command line interface or SMIT to add, delete, or modify the extension node information in the SDR. This information consists of both an extension node and an extension node adapter definition.

To add an extension node, there are several steps to follow. You perform some of these steps on the extension node. Most of the steps, however, must be performed on the control workstation. You will have to login to the extension node, either directly or via a telnet session to perform the steps required on the extension node.

As with many PSSP commands, you must have the appropriate authority or credentials to use them. See "Step 24: Obtain credentials" on page 46.

This chapter provides a high-level discussion of the steps you must perform on the control workstation and on the extension node. For more specific information, refer to the following documents:

- *PSSP: Administration Guide*
- *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment*
- *SP Switch Router Adapter Guide* that comes with the SP Switch Router Adapter
- All of the IBM SP Switch Router documentation

Control workstation steps

Perform the following steps on the control workstation:

1. Install the **ssp.spmgr** file set

   If the file set **ssp.spmgr** is not already installed on the control workstation, install it now by issuing the following commands:

If using:	Do this:
SMIT	**TYPE** **smit install_latest** • The Install New Software Products at Latest Level window appears. **ENTER** **/spdata/sys1/install/pssplpp/***code_version* for Input Device **PRESS** **Do** to display the default install parameters. **PRESS** **List** to show options. **SELECT** **ssp.spmgr** **PRESS** **Do** to complete option selection and to begin installation. When the installation is complete, check the SMIT log file for the installation status. If errors occur, see the *IBM AIX Problem Solving Guide and Reference*.

installp	You can use **installp** to install multiple file sets. For example, to install all of the file sets, enter:
	`installp -a -d /spdata/sys1/install/pssplpp/code_version -X ssp.spmgr`
	Note: installp automatically commits the packaging file set when you specify the **-a** option.
	To list all of the options for **ssp**, enter:
	`installp -l -d /spdata/sys1/install/pssplpp/code_version/pssp.installp`

2. Verify Port Information

   Once **ssp.spmgr** is installed on the control workstation, UDP port 162 becomes reserved for SNMP traffic between the SP Extension Node SNMP Manager and the SNMP Agent on the extension node. This port is the default SNMP trap port for any SNMP Manager. If another LP which includes an SNMP Manager is configured on the control workstation, there is a conflict which you would need to resolve by modifying the entry for **spmgrd-trap** in **/etc/services** on the control workstation to specify an unused UDP port number.
3. Define the Extension Node

   Issue the **endefnode** command to add an extension node definition in the SDR. You need to issue this command for each extension node you want to configure. For more information, refer to the *PSSP: Command and Technical Reference*.
4. Define the Extension Node Adapter

   Issue the **endefadapter** command to add an extension node adapter definition in the SDR. You need to issue this command for each extension node adapter you want to configure.

Extension node steps

Perform the following steps on the extension node:

1. Optionally, Define the Extension Node and Extension Node Adapter

   IBM requires that you define the extension node and the extension node adapter on the control workstation, as discussed in the previous section. In addition, you can optionally define these items on the extension node.

   Once the items have been defined, the extension node is configured automatically. The SNMP Agent on the extension node polls the SPMGR Manager running on the control workstation for the information at a designated interval. The SPMGR Manager looks for the information for the extension node adapter specified and passes back the configuration information. Once the

extension node and extension node adapter information has been added to the SDR, the SPMGR Manager is able to respond with the configuration information.

For more information, refer to the *SP Switch Router Adapter Guide*.

2. Configure the SNMP Agent to communicate with the extension node SNMP Manager on the control workstation. See the *SP Switch Router Adapter Guide* for further information.
3. Install the SP Switch Router and all IP interfaces connected to it.

   For more information, refer to the *SP Switch Router Adapter Guide*.

Activate extension node step

Issue the **Estart** command on the control workstation to reconfigure the switch to recognize the new extension nodes.

Note: If you are using the Switch Admin daemon for node recovery, start it by issuing **startsrc -s swtadmd** on SP Switch systems or **startsrc -s swtadmd2** on SP Switch2 systems before issuing the **Estart** command.

Extension node verification steps

There are several ways to verify that the extension node has been successfully installed and is operating correctly. Perform steps 1, 2, and 3. You can optionally perform step 4:

1. Issue the following command to verify the definition of the extension node:

   ```
   splstnodes -t dependent node_number reliable_host_name \
   management_agent_hostname extension_node_identifier snmp_community_name
   ```

2. Issue the following command to verify the definition of the extension node adapter:

   ```
   splstadapters -t dependent node_number netaddr netmask
   ```

3. Issue the following command to verify that the extension node is connected to the switch:

   ```
   SDRGetObjects switch_responds
   ```

   You should receive output similar to the following:

   ```
   node_number switch_responds autojoin  isolated     adapter_config_status
   1           0               0         0            css_ready
   3           0               0         0            css_ready
   4           0               0         0            css_ready
   5           0               0         0            css_ready
   6           0               0         0            css_ready
   7           0               0         0            css_ready
   ```

4. Use the Perspectives GUI to verify the definitions of the extension node and extension node adapter, and to verify that the extension node is connected to the switch.

   Issue **perspectives &** to bring up Perspectives as a background process.

   SELECT The **Hardware Perspective** icon by double clicking

   SELECT The IP node icon in the Nodes Pane

   SELECT The Notebook icon on the tool bar

   The notebook associated with the node appears. Review the node information to ensure it is operating correctly.

Appendix A. SAGE job descriptions

This appendix describes the job descriptions for the system administration levels discussed in “Who should use this book” on page xvii.

The System Administrators Guide of USENIX (SAGE), a professional organization for system administrators, has developed a set of job descriptions for system administrators. The Core templates described in this appendix describe the core set of attributes for system administrators. The “Additional skill areas” on page 286 provides details on additional skill areas that may be important at your site.

For more information, contact the USENIX Association or access the the following Web site:

http://www.sage.usenix.org/sage/

Core templates defined by SAGE

The following are templates for the core skills required for Junior and Intermediate/Advanced System Administrators.

Junior system administrator

Required skills

- Strong interpersonal and communication skills; capable of training users in applications and UNIX fundamentals, and writing basic documentation.
- Highly skilled in using most UNIX commands/utilities.
- Familiarity with most basic system administration tools and processes; for example, can boot/shutdown a machine, add and remove user accounts, use backup programs and fsck, maintain system database files (groups, hosts, aliases).
- Fundamental understanding of a UNIX-based operating system; for example, understands job control, soft and hard links, distinctions between the kernel and the shell.

Required background

One to three years of system administration experience.

Desirable background

- A degree in computer science or a related field.
- Familiarity with networked/distributed computing environment concepts; for example, can use the route command, add a workstation to a network, and mount remote filesystems.
- Ability to write scripts in some administrative language (Tk, Perl, and shell).
- Programming experience in any applicable language.

Appropriate responsibilities

Administers a small site alone or assists in the administration of a larger system. Works under the general supervision of a system administrator or computer systems manager.

Intermediate/advanced system administrator

Required skills

- Strong interpersonal and communication skills; capable of writing purchase justifications, training users in complex topics, making presentations to an internal audience, and interacting positively with upper management.
- Independent problem solving; self-direction.
- Is comfortable with most aspects of UNIX system administration; for example, configuration of mail systems, system installation and configuration, printing systems, fundamentals of security, installing third-party software.
- A solid understanding of a UNIX-based operating system; understands paging and swapping, interprocess communication, devices and what device drivers do, file system concepts (″inode″, ″superblock″).
- Familiarity with fundamental networking/distributed computing environment concepts; can configure NFS and NIS, can use nslookup or dig to check information in the DNS, understands basic routing concepts.
- Ability to write scripts in some administrative language (Tk, Perl, and shell).
- Ability to do minimal debugging and modification of C programs.

Required background

Three to five years system administration experience.

Desirable background

- A degree in computer science or a related field.
- Significant programming background in any applicable language.

Appropriate responsibilities

- Receives general instructions for new responsibilities from supervisor.
- Administers a midsized site alone or assists in the administration of a larger site.
- Initiates some new responsibilities and helps to plan for the future of the site/network.
- Manages novice system administrators or operators.
- Evaluates and/or recommends purchases; has strong influence on purchasing process.

Additional skill areas

Use this list to identify additional skill areas that may be important at your site.

Local Environment Experience

Experience with the specific operating systems, applications, or programming languages in use at the site (for example SunOS, AIX, CAE/CAD software, FrameMaker, Mathematica, Fortran, Ada). Experience with the work done by the users at the site.

Heterogeneity Experience

Experience with more than one UNIX-based operating system. Experience with sites running more than one UNIX-based operating system. Familiarity with both System V and BSD-based UNIX operating systems. Experience with non-UNIX operating systems (for example, MS-DOS, Macintosh OS, or VMS). Experience with internetworking UNIX and other operating systems (MS-DOS, Macintosh OS, VMS).

Programming Skills

Extensive programming experience in an administrative language (Tk, Perl, and shell). Extensive programming experience in any applicable language.

Networking Skills

Experience configuring network file systems (for example, NFS, RFS, or AFS). Experience with network file synchronization schemes (for example, rdist and track). Experience configuring automounters. Experience configuring license managers. Experience configuring NIS/NIS+. Experience with TCP/IP networking protocols (ability to debug and program at the network level). Experience with non-TCP/IP networking protocols (for example, OSI, Chaosnet, DECnet, Appletalk, Novell Netware, Banyan Vines). Experience with high-speed networking (for example, FDDI, ATM, or SONET). Experience with complex TCP/IP networks (networks that contain routers). Experience with highly complex TCP/IP networks (networks that contain multiple routers and multiple media). Experience configuring and maintaining routers. Experience maintaining a site-wide modem pool/terminal servers. Experience with X/X terminals. Experience with dial-up networking (for example, SLIP, PPP, or UUCP). Experience at a site that is connected to the Internet. Experience installing/configuring DNS/BIND. Experience installing/administering Usenet news. Experience as postmaster of a site with external connections.

Security

Experience with network security (for example, building firewalls, deploying authentication systems, or applying cryptography to network applications). Experience with classified computing. Experience with multi-level classified environments. Experience with host security (for example, passwords, uids/gids, file permissions, file system integrity, use of security packages).

Site Specialities

Experience at sites with over 1000 computers, over 1000 users, or over a terabyte of disk space. Experience with supercomputers. Experience coordinating multiple independent computer facilities (for example, working for the central group at a large company or university). Experience with a site with 100% uptime requirement. Experience developing/implementing a site disaster recovery plan. Experience with a site requiring charge-back accounting.

Documentation

Background in technical publications, documentation, or desktop publishing.

Databases

Experience using relational databases. Experience using a database query language. Experience programming in a database query language. Previous experience as a database administrator.

Hardware

Experience installing and maintaining the network cabling in use at the site. Experience installing boards and memory into systems. Experience with SCSI device setup and installation. Experience installing/configuring peripherals (for example, disks, modems, printers, or data acquisition devices). Experience with board-level diagnosis and repair of computer systems. Experience with component-level diagnosis and repair of computer system.

Management

Budget responsibility. Experience in writing personnel reviews, and ranking processes. Experience in interviewing and hiring.

Appendix B. Directory information

The following table lists the directories created when you install the RS/6000 SP LP image.

Table 6. RS/6000 SP LP Images Directories

Directory	Contents
/etc/amd	amd files
/etc/auto	automounter files
/etc/ssp/css	css user space library
/spdata/sys1/err_methods	Contains scripts for error handling
/spdata/sys1/ha	Contains files related to RSCT
/spdata/sys1/k4srvtabs	Contains node's Kerberos V4 srvtab files
/spdata/sys1/keyfiles	Contains node's keyfiles
/spdata/sys1/logtables	Sample service collection tables
/spdata/sys1/pman	Contains problem management configuration files
/spdata/sys1/sdr	SDR database files
/spdata/sys1/spmon	hardmon internal security and threshold tables
/spdata/sys1/spsec	Security Services files
/spdata/sys1/st	Job Switch Resource Table Services data files
/spdata/sys1/syspar_configs	System partition configuration files
/spdata/sys1/ucode	Contains supervisor microcode levels
/spdata/sys1/vsd	Partitioning data with regard to vsds
/usr/lpp/ssp	Package files
/usr/lpp/ssp/amd	amd/amq files
/usr/lpp/ssp/bin	ssp command files
/usr/lpp/ssp/bin/spd	Advanced Diagnostics commands
/usr/lpp/ssp/config	Setup file for switch, and so on
/usr/lpp/ssp/config/admin	User mgmt SMIT stanzas
/usr/lpp/ssp/config/cmi	Cluster setup/management SMIT stanzas
/usr/lpp/ssp/config/spmgrd	MIB files for SNMP manager
/usr/lpp/ssp/css	Communication subsystem (css) files
/usr/lpp/ssp/css/diags	css diagnosis files
/usr/lpp/ssp/css/spd	Advanced Diagnostics files
/usr/lpp/ssp/filec	File collection installation files
/usr/lpp/ssp/info	InfoExplorer setup files
/usr/lpp/ssp/inst_root	Root package files
/usr/lpp/ssp/install/bin	Configuration scripts
/usr/lpp/ssp/install/config	Install/config data and templates
/usr/lpp/ssp/kerberos/bin	Kerberos V4 user commands

Table 6. RS/6000 SP LP Images Directories (continued)

Directory	Contents
/usr/lpp/ssp/kerberos/etc	Kerberos V4 daemons and privileged administrative commands
/usr/lpp/ssp/lib	ssp libraries
/usr/lpp/ssp/man	man page files
/usr/lpp/ssp/perl/bin	Perl translation commands
/usr/lpp/ssp/perl/lib	Perl functions and headers
/usr/lpp/ssp/perl/lib/sys	C headers converted to Perl
/usr/lpp/ssp/perl/man	Perl man page files
/usr/lpp/ssp/perspectives	Perspectives files
/usr/lpp/ssp/public	Public tar files
/usr/lpp/ssp/rcmd/bin	Kerberos-authenticated remote commands
/usr/lpp/ssp/rcmd/etc	RS/6000 SP command daemon
/usr/lpp/ssp/resctr	Resource Center HTML files and programs
/usr/lpp/ssp/samples	Sample scripts
/usr/lpp/ssp/sysctl/bin	Supported client code
/var/adm/acct	Created when accounting is configured on nodes.
/var/adm/cacct	Cluster consolidated accounting data
/var/adm/SPlogs	Logs and other run time data collectors
/var/adm/SPlogs/auth_install	Security services logs
/var/adm/SPlogs/auto	automount logs
/var/adm/SPlogs/cs	Cluster startup/shutdown logs
/var/adm/SPlogs/css	css logs
/var/adm/SPlogs/css0	css logs (SP Switch2 only)
/var/adm/SPlogs/css1	css logs (SP Switch2 only)
/var/adm/SPlogs/css0/p0	css logs (SP Switch2 only)
/var/adm/SPlogs/css1/p0	css logs (SP Switch2 only)
/var/adm/SPlogs/filec	File collection logs
/var/adm/SPlogs/kerberos	Kerberos server logs
/var/adm/SPlogs/kfserver	Keyfile server logs
/var/adm/SPlogs/sdr	System Data Repository logs
/var/adm/SPlogs/SPconfig	Vital product data information
/var/adm/SPlogs/spmgr	Trace file for SNMP manager
/var/adm/SPlogs/spmon	SP System Monitor logs
/var/adm/SPlogs/st	Job Switch Resource Table Services logs
/var/adm/SPlogs/sysctl	Sysctl server log file
/var/adm/SPlogs/sysman	Installation, configuration, and console logs
/var/kerberos/database	Kerberos server database and other files

Appendix C. SP Perspectives tasks

This appendix discusses how to start and use Perspectives. In addition, a table appears describing how to perform some of the most common administrative tasks.

Starting and using Perspectives

Perspectives is a graphical user interface you can use to manage and monitor the SP system. Complete instructions for using Perspectives appear in the online help information.

You can use Perspectives during the installation and migration process to perform a variety of tasks including installing software and verifying executed steps.

To start Perspectives by bringing up the Launch Pad, first export your display and then type **perspectives &** to run the process in the background. The Launch Pad gives you graphical access to Perspectives applications. To launch an application, either double click on the icon associated with an application or select the icon using the Launch action found under the Actions menu.

In each of the Perspectives, to perform an action, all you need to do is to click on the icon of one or more objects and then click on an action in either the menu bar or the tool bar. If you need more information to complete the action, a dialog box will appear. In addition, each Perspective has online help available from the Help menu selection.

Note: Many of the steps require you to use SMIT. You can use SMIT directly by issuing SMIT commands or you can access SMIT using Perspectives. Often used SMIT fast-paths such as **smit_verify** are available from the Perspectives Launch Pad.

The following table discusses the most common tasks that you can perform using Perspectives.

Table 7. Performing Common Tasks Using SP Perspectives

Task	**Perspectives**
Bringing up the GUI	**TYPE perspectives &** to bring up the Launch Pad
Launching applications	Double click on application icon in Launch Pad or **SELECT** application icon **SELECT** Actions → Launch
Closing the GUI	**SELECT** Window → Exit to close the Launch Pad or any Perspective
Closing a window	**SELECT** Cancel button or select Close from the top left of window.

Table 7. Performing Common Tasks Using SP Perspectives (continued)

Task	Perspectives
Powering off one node	From **Hardware Perspective** menu bar **SELECT** Node to be powered off **SELECT** Actions → Power Off, Reset, or Shutdown... **SELECT** Power off options **PRESS** Apply or from **Hardware Perspective** tool bar **SELECT** Node to be powered off **PRESS** Power off icon from Tool bar **SELECT** power off options **PRESS** Apply
Powering off more than one node	From **Hardware Perspective** menu bar **SELECT** Nodes to be powered off **SELECT** Actions → Power Off, Reset, or Shutdown... **SELECT** Power off options **PRESS** Apply or from **Hardware Perspective** menu bar with Node Groups **SELECT** Node group to be powered off **SELECT** Actions → Power Off, Reset, or Shutdown... **SELECT** Power off options **PRESS** Apply or from **Hardware Perspective** tool bar **SELECT** Nodes to be powered off **PRESS** Power off icon from Tool bar **SELECT** Power off options **PRESS** Apply
Powering on one node	From **Hardware Perspective** menu bar **SELECT** Node to be powered on **SELECT** Actions → Power On or Cluster Power On... **SELECT** Power on options **PRESS** Apply or From **Hardware Perspective** tool bar **SELECT** Node to be powered on **PRESS** Power on icon from Tool bar **SELECT** Power on options **PRESS** Apply

Table 7. Performing Common Tasks Using SP Perspectives (continued)

Task	Perspectives
Powering on more than one node	From **Hardware Perspective** menu bar **SELECT** Nodes to be powered on **SELECT** Actions → Power On or Cluster Power On... **SELECT** Power on options **PRESS** Apply or From **Hardware Perspective** menu bar with Node Groups **SELECT** Node groups to be powered on **SELECT** Actions → Power On or Cluster Power On... **SELECT** Power on options **PRESS** Apply or From **Hardware Perspective** tool bar **SELECT** Nodes to be powered on **PRESS** Power on icon from Tool bar **SELECT** Power on options **PRESS** Apply
Verifying node status for many nodes: hostResponds, powerLED, switchResponds, LCD and LED	From **Hardware Perspective** menu bar **Click** On the Nodes pane **SELECT** View → Set Monitoring **SELECT** hostResponds **PRESS** Apply Repeat these steps for individual viewing of node powerLED and switchResponds status, or select all three. You can create multiple Nodes panes and monitor one condition in each pane. Check to see if any node icons do not turn green. If a red X is showing, there is a problem with that node. If a “?” appears in the node, there is a communication problem with the underlying subsystem. The LCD and LED are always displayed in its own window. Bring up the LCD and LED: **SELECT** The Nodes pane **SELECT** Actions → LCD and LED Display

Table 7. Performing Common Tasks Using SP Perspectives (continued)

Task	Perspectives
Verifying node status for one node: hostResponds, powerLED, and switchResponds.	From **Hardware Perspective** node notebook **SELECT** A node by double clicking on it. **SELECT** Node Status tab You can also now perform actions from the Node Status page, such as power on and off, and so on. To verify the status of an extension node: **SELECT** The extension node by double clicking on it. Once the notebook appears, check that the switchResponds variable is running. The Node Status page also displays the value of the LCD and LED display for the node.
Viewing node information	**SELECT** A node by double clicking on it. **SELECT** Tab corresponding to desired page
Changing to the system (global) view	From **Hardware Perspective** menu bar **SELECT** System icon **SELECT** Actions → Set Current System Partition If a Syspar pane is not present: **SELECT** View → Change System Partition and select Global
Changing to a partition (syspar) view	From **Hardware Perspective** menu bar **SELECT** Syspar icon **SELECT** Actions → Set Current System Partition If a Syspar pane is not present: **SELECT** View → Change System Partition and select the syspar name you want

Appendix D. Boot/install server configuration commands

This appendix lists the names of the boot/install server configuration commands which make up setup_server. For more information on these commands, refer to the *PSSP: Command and Technical Reference*.

Table 8. Boot/Install Server Configuration Commands

Command	**Description.**
allnimres	Allocates NIM resources from NIM server to client and prepares the client node for boot/installation.
unallnimres	Unallocates NIM resources from NIM server to client.
mknimclient	Makes a node a NIM client of its boot server as specified in the SDR.
delnimclient	Deletes the NIM client definition of a node on it's boot server.
mknimmast	Configures a node as a NIM master.
delnimmast	Unconfigures a node as a NIM master.
mknimint	Creates the necessary interfaces on a NIM master.
mknimres	Creates the necessary NIM resources on a NIM master.
mkconfig	Creates the config.info file on the server (must run on the server).
mkinstall	Creates the install.info file on the server (must run on the server).
setup_CWS	Set up the necessary files and other conditions on the control workstation.
create_krb_files	Creates the necessary Kerberos V4 and tftp access files on a boot server.
export_clients	Exports required file systems from a NIM master to its clients.
setup_server	Sets up a boot/install server (same function, different internals).

Appendix E. User-supplied node customization scripts

IBM provides the opportunity to run two different customer-supplied scripts during node installation:

script.cust

This script is run from the PSSP NIM customization script (**pssp_script**) after the node's AIX and PSSP software have been installed, but before the node has been rebooted. This script is run in a limited environment where not all services are fully configured:

- The installation RAM file system is in place (and not the normal disk file system).
- The node has communications connectivity and routing to its boot/install server only. For example, there is no general connectivity to the control workstation or any internal or external network.
- None of the node's network adapters, with the exception of the installation adapter, are configured and available.
- There is no access to authentication services.

As a result of these restrictions, you should restrict your use of **script.cust** to function that must be performed prior to the post-installation reboot of the node. An example of appropriate function in **script.cust** is the installation of an LP that requires a reboot of the node to render it fully functional. See the sample **script.cust** that follows for additional examples.

Note: **script.cust** is also run during the boot of a node that has been set to "customize."

A sample of the **script.cust** node customization script is located at **/usr/lpp/ssp/samples/script.cust**.

firstboot.cust

This script is run during the first boot of the node immediately after it has been installed. This script runs in a more "normal" environment where most all services have been fully configured:

- The node's disk file system is in place.
- The node has established its defined communications connections including all routing.
- The **css0** adapter has been configured, but is not available.
- The other network adapters configured by PSSP have been started.
- Authentication services are available.

This script is a preferred location for node customization functions that do not require a reboot of the node to become fully enabled.

Note: **firstboot.cust** is also run during the boot of a node that has been set to "migrate."

A sample of the **firstboot.cust** node customization script is located at **/usr/lpp/ssp/samples/firstboot.cust**.

Note: Your security environment is not set up during **script.cust** processing. If you require security function, perform your customization during **firstboot.cust** processing.

Be aware that the root password is not set if you are installing from a minimal mksysb image. If you want to have the root password set, you can do one of the following:

- Modify the **firstboot.cust** file by uncommenting the portion of the code to copy the user information from the server to the node.
- If you did not copy the user management files over, then immediately after the node is up without a password, login to the node as root and set the password either through the SMIT panels under Security & Users or with the **passwd** command.

Name resolution of the control workstation host name is required on the nodes. **firstboot.cust** shows how to copy **/etc/resolv.conf** or **/etc/hosts**, and how to define the node to NIS, if that is being used. If you have created a **mksysb** image to install on the nodes, and it already has name resolution defined in it, you do not need to do anything additional.

You should use any customizations that were created by using the **chdev** or **chgcss** commands, except for those that are already stored in the SDR. For example, customized parameters such as maxuproc, maxmbufs, rpoolsize, spoolsize, and asynchronous I/O should be preserved by using the **chdev** or **chgcss** command in these files.

How to use the node customization scripts

Both **script.cust** and **firstboot.cust** are run during node installation if the corresponding file is located in the **/tftpboot** directory of the node's boot/install server. If you use a common version of **script.cust** and **firstboot.cust** across all the nodes in your SP system, then placing one or both of these files on the control workstation prior to installing any nodes will cause them to be copied and run on the boot/install server nodes (if you have configured boot/install servers) and, subsequently, to all nodes in your SP as they are installed.

To use the **script.cust** customization script, do the following tasks:

- Examine the sample **script.cust** file located in **/usr/lpp/ssp/samples** to determine whether any lines are appropriate to your installation.
- Copy the file to **/tftpboot/script.cust** on the control workstation and uncomment, edit, or add any necessary lines. Make sure you change the file's permissions to be "world-readable." If you forget to set the permissions properly, the script will not run.

To use the **firstboot.cust** customization script, do the following tasks:

- Examine the sample **firstboot.cust** file located in **/usr/lpp/ssp/samples** to determine whether any lines are appropriate to your installation.
- Copy the file to **/tftpboot/firstboot.cust** on the control workstation and uncomment, edit, or add any necessary lines. Make sure you change the file's permissions to be "world-readable." If you forget to set the permissions properly, the script will not run.

The **firstboot.cmds** file contains a sample sysctl procedure that can be called from **firstboot.cust** in restricted root rcmd mode and secure remote command mode. It

copies files from the server (typically the control workstation) to the node. If you are using AIX remote commands, the server must be authorized to issue remote commands to the requesting node.

To use this procedure, perform the following steps on the server:

- Edit the file to remove the comments from the remote command calls below what you want to have executed during the node's **firstboot.cust** processing. Add additional remote command calls as required for your node customization.
- Ensure that the **firstboot.cmds** file is loaded by sysctl by editing **/etc/sysctl.conf** to include the **/usr/lpp/ssp/samples/sysctl/firstboot.cmds** file.
- Start and restart the **sysctl** daemon to reload this changed procedure by issuing:

  ```
  stopsrc -s sysctld

  startsrc -s sysctld
  ```

- Edit the **firstboot.cust** file that will be copied to the node during installation. Invoke this procedure by adding the following:

  /bin/sysctl -L -h $SERVER copy_env_files

Once the node is installed, or as part of **firstboot.cust**, the remote command authorization files on the node serviced by the non-control workstation boot/install server need to be updated with the following changes:

standard An entry for the boot/install server node host name in **/.rhosts**

k4 An entry for the boot/install server node rcmd principal in **/.klogin**

dce An entry for the **self-host** and the **spbgroot** principal for the boot/install server node

Migration and coexistence issues related to the node customization scripts

Prior to PSSP 2.4, the **script.cust** user customization script ran immediately after the installation of the node, but prior to the reboot of that node. As of PSSP 2.4, **script.cust** still runs at this time, but the environment established for this script has changed somewhat (see the preceding discussion). In particular, **script.cust** now runs before full network connectivity has been established. For instance, this means that access to the authentication server is not available when this script runs.

The **firstboot.cust** customization script is now available for user customization tasks. This script runs after the initial post-installation reboot of the node. Full network connectivity (through the I/O adapters automatically configured by PSSP) is established and the authentication server is accessible.

Because of this change, the version of **script.cust** that you may have written to use during the installation of nodes at PSSP levels prior to Version 2.4, may no longer function correctly. To rectify this situation, you should review the contents of your **script.cust** file during the migration of your nodes to the latest level of PSSP. You may need to move some function from your existing **script.cust** file to a new **firstboot.cust** file. See the preceding discussion and the new sample files for more specific information.

If you will be installing nodes at different PSSP levels where one or more nodes are at an earlier PSSP level and one or more are at PSSP 2.4 or later, you may need to structure your **script.cust** file to contain logic that performs certain functions on earlier PSSP level nodes and different functions on PSSP 2.4 or later level nodes.

The sample versions of **script.cust** and **firstboot.cust** contain sample code to implement the coincident use of these scripts on nodes at differing PSSP and AIX levels. (Note that **firstboot.cust** is run on a node being installed at an earlier PSSP level if that node's boot/install server is running at the PSSP 2.4 or later level.)

tuning.cust file

A sample of the **tuning.cust** file is located at **/usr/lpp/ssp/samples**. Additional IBM-supplied tuning files are available to set your performance parameters. One file is for the commercial environment, one is for the development environment, and one is for the scientific environment. For more information on these files, refer to *PSSP: Command and Technical Reference*.

Appendix F. Overriding the PPSIZE in a mksysb image

When installing a new node in an existing configuration, the mksysb **image.data** file will be used to describe the node's rootvg volume group. If the new node has a different size disk than the existing nodes, the parameters in the **image.data** file may have to be altered.

Refer to the AIX **mkvg** command for related hardware limitations. A 4.5 GB disk requires an 8 MB partition size and a 9 GB drive requires 16 MB.

If the PPSIZE entry in the **image.data** file is too small for the disk that is being installed, **no** action is needed. The AIX install program will automatically increase the PPSIZE to the appropriate number and adjust the LPs to match (cutting them in half or to one-quarter of their original size).

If you want to change the PPSIZE to a value other than the default or to modify other **image.data** parameters, follow this procedure. For more information about the **image.data** file, see the *AIX Files Reference*.

To alter the PPSIZE of the root volume group, do the following:

1. Extract the **image.data** file from mksysb:

   ```
   cd /spdata/sys1/install/images
   restore -xvf MKSYSB ./image.data
   ```

   where MKSYSB is the file name of your mksysb image.

2. Use a text editor to modify the PPSIZE in the **vg_data** stanza of the **image.data** file. The **vg_data** stanza looks like the following:

   ```
   vg_data:
      VGNAME= rootvg
      PPSIZE= 4
      VARYON= yes
      VG_SOURCE_DISK_LIST= hdisk0
      QUORUM= 2
   ```

   Each logical volume in the group has a stanza named **lv_data**. They also must be modified to be consistent with the new PPSIZE.

3. Create an **image.data** resource from the SMIT panel:

   ```
   smitty nim_res

   Select: Define a Resource
   Select: image_data
   ```

   On the SMIT panel, enter these values:

   ```
     Resource Name              [imagedata]
     Resource Type              image_data
     Server of Resource         [master]
     Location of Resource       [DIR/image.data]
     comments                   []
   ```

4. Set the node to install by using the **spbootins** command:

   ```
   spbootins -r install -l node_number
   ```

5. Allocate the **image.data** NIM resource to the node:

   ```
   nim -Fo reset node_hostname
   ```

   This will remove the boot and nim_script entries which will change the NIM state so that the **image.data** NIM resources can be allocated to the node.

Bring up the SMIT panel:

```
smitty nim_mac

Select: Manage Network Install Resource Allocation
Select: Allocate Network Install Resources
Select: node_hostname

Select: imagedata            image_data
```

When the command completes, hit cancel twice to return to the “Manage Machines” panel.

```
Select: Perform Operations on Machines
Select: node_hostname
Select: bos_inst = perform a BOS installation
```

The “Perform a Network Install” panel will come up. Enter these values:

```
Target Name                                  node_hostname
Source for BOS Runtime Files                        mksysb
installp Flags                                      [-agX]
Fileset Names                                       []
Remain NIM client after install?                    yes
Initiate Boot Operation on Client?                  no
Set Boot List if Boot not Initiated on Client?      no
Force Unattended Installation Enablement?           no
```

Press Enter after making all desired changes. Exit SMIT.

6. Verify that the imagedata resource and all NIM resources associated with an “install” were allocated.

   ```
   lsnim -c resources node_hostname
   ```

   The following output will be generated:

   ```
   lppsource_AIX421   lpp_source
   spot_AIX421        spot
   noprompt           bosinst_data
   psspscript         script
   mksysb_1           mksysb
   imagedata          image_data
   boot               boot represents the network boot resource
   nim_script         nim_script - directory containing customization script
   ```

7. Perform a netboot for all the nodes in question.

Appendix G. Reserving ports

Components of the high availability infrastructure (Topology Services, Group Services, and Event Management) allocate port numbers from the range 10000 to 10100, inclusive. Topology Services, Group Services, and Event Management allocate port numbers when they are configured on the control workstation.

- Topology Services and Group Services allocate one port number per system partition.
- Event Management allocates one port number per system partition plus one.

If a port number from this range is already listed in **/etc/services** on the control workstation, then these subsystems try another port number. Therefore, if any customer subsystem uses port numbers in the range 10000-10100, such port numbers must be reserved in **/etc/services** on the control workstation before Topology Services, Group Services, and Event Management allocate port numbers.

Procedures

Topology Services, Group Services, and Event Management allocate their port numbers when their control scripts, respectively **hatsctrl**, **hagsctrl**, and **haemctrl** are run with the **−a** flag. Allocated port numbers are saved in the System Data Repository (SDR). If port numbers are found in the SDR from a previous invocation of a control script, the saved port numbers are used and new values are not allocated.

When the **syspar_ctrl −A** command is invoked, the **hatsctrl**, **hagsctrl**, and **haemctrl** scripts are invoked with the **−a** flag. As part of installation and migration procedures, **syspar_ctrl −A** is invoked. Specifically, it is invoked as part of the "Start RSCT subsystems" step in "Chapter 2. Installing and configuring a new RS/6000 SP system" on page 11 and in "Chapter 4. Migrating the software on your RS/6000 SP system" on page 127. If port numbers must be reserved in **/etc/services**, it must be done prior to "Start RSCT subsystems" in "Chapter 2. Installing and configuring a new RS/6000 SP system" on page 11 or when preparing to migrate as described in "Chapter 4. Migrating the software on your RS/6000 SP system" on page 127.

To reserve a port number, enter it in the **/etc/services** file on the control workstation using an editor or by invoking the **smit clientnet** command. The service name and protocol should be specified for the application or subsystem that is using the port number. If the application or subsystem does not run on the control workstation, the port numbers can be reserved by using "dummy" service names. In the latter case, reserve the port number for use with both the **tcp** and **udp** protocols.

Considerations for Network Information Service (NIS)

If the control workstation is a NIS client, port numbers must be reserved using standard NIS procedures. This involves updating the NIS services map on the master server and any slave servers that serve the NIS domain to which the control workstation belongs. See the "Network Information Service" chapter of the *AIX System Management Guide: Communications and Networks* for the procedures to update NIS maps.

Resolving port number conflicts

It may be the case when the availability infrastructure subsystems are configured, that you do not know if the subsystem you intend to install later is also using port numbers in the range 10000 through 10100. You do not reserve these port numbers, which leads to a conflict at a later time. If the subsystem cannot be configured to use a different port number, the port numbers used by the availability infrastructure subsystems can be reallocated to avoid any conflict. Use the following procedure:

1. Execute the following commands on each node, in the order listed:

   ```
   haemctrl -c
   hagsctrl -c
   hatsctrl -c
   ```

2. Execute the following commands on the control workstation in each system partition, in the order listed:

   ```
   haemctrl -u
   hagsctrl -u
   hatsctrl -u
   ```

   The **−u** flag is the same as the **−c** flag of these scripts with the addition of removing the subsystem's port numbers from the SDR.

   Caution: Steps 1 and 2 must execute successfully before proceeding.

3. Reserve the port numbers needed by your subsystem.
4. Issue the following commands, first on the control workstation in each system partition and then on each node, in the order listed:

   ```
   hatsctrl -a
   hatsctrl -s

   hagsctrl -a
   hagsctrl -s

   haemctrl -a
   haemctrl -s
   ```

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.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Department LJEB/P905
2455 South Road
Poughkeepsie, NY 12601-5400
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

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

AFS
AIX
AIX 5L
DATABASE 2
DB2
DFS
e (logo)
ES/9000
ESCON
HACMP/6000
IBM
IBM (logo)

IBMLink
LoadLeveler
Micro Channel
Open Class
POWERparallel
POWERserver
pSeries
PTX
Redbooks
RS/6000
RS/6000 Scalable POWERparallel Systems
Scalable POWERparallel Systems
SP
System/370
System/390
TURBOWAYS
VisualAge

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, MS-DOS, Windows, Windows NT, BackOffice, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Publicly Available Software

PSSP includes software that is publicly available:

expect
Programmed dialogue with interactive programs

Perl Practical Extraction and Report Language

SUP Software Update Protocol

Tcl Tool Command Language

TclX Tool Command Language Extended

Tk Tcl-based Tool Kit for X-windows

This book discusses the use of these products only as they apply specifically to the RS/6000 SP system. The distribution for these products includes the source code and associated documentation. **/usr/lpp/ssp/public** contains the compressed **tar** files of the publicly available software. (IBM has made minor modifications to the versions of Tcl and Tk used in the SP system to improve their security characteristics. Therefore, the IBM-supplied versions do not match exactly the versions you may build from the compressed **tar** files.) All copyright notices in the documentation must be respected. You can find version and distribution information for each of these products that are part of your selected install options in the **/usr/lpp/ssp/README/ssp.public.README** file.

Specified Operating Environment

Hardware Specifications

The following list contains the minimum hardware requirements for your RS/6000 SP system:

- A control workstation - refer to *RS/6000 SP: Planning, Volume 2, Control Workstation and Software Environment* for the recommended hardware configuration
- If you have clustered enterprise servers attached to your SP, there are cabling requirements that you must also consider. See *RS/6000 Enterprise Server: Installation and Service Guide*, SA38-0558 for more information. Additional hardware requirements for clustered enterprise servers include:
  - A control workstation
  - One to 16 clustered enterprise servers

Programming Specifications

The following list contains the minimum PSSP software requirements for your RS/6000 SP system:

- AIX 4.3.3, AIX 5L 5.1, or later Base Operating System
- PSSP 3.4
- At least one of the following:
  - C for AIX, Version 5.0.2 or later
  - VisualAge C++ Professional for AIX, Version 5.0.2 or later

Glossary of Terms and Abbreviations

A

ACL. Access Control List. A list that defines who has permission to access certain services; that is, for whom a server may perform certain tasks. This is usually a list of principals with the type of access assigned to each.

adapter. An adapter is a mechanism for attaching parts. For example, an adapter could be a part that electrically or physically connects a device to a computer or to another device. In the SP system, network connectivity is supplied by various adapters, some optional, that can provide connection to I/O devices, networks of workstations, and mainframe networks. Ethernet, FDDI, token-ring, HiPPI, SCSI, FCS, and ATM are examples of adapters that can be used as part of an SP system.

address. A character or group of characters that identifies a register, a device, a particular part of storage, or some other data source or destination.

AFS. A distributed file system that provides authentication services as part of its file system creation.

AIX. Abbreviation for Advanced Interactive Executive, IBM's licensed version of the UNIX operating system. AIX is particularly suited to support technical computing applications, including high function graphics and floating point computations.

API. Application Programming Interface. A set of programming functions and routines that provide access between the Application layer of the OSI seven-layer model and applications that want to use the network. It is a software interface.

application. The use to which a data processing system is put; for example, a payroll application, an airline reservation application.

application data. The data that is produced using an application program.

ARP. Address Resolution Protocol.

ATM. Asynchronous Transfer Mode. (See *TURBOWAYS 100 ATM Adapter*.)

authentication. The process of validating the identity of either a user of a service or the service itself. The process of a principal proving the authenticity of its identity.

authorization. The process of obtaining permission to access resources or perform tasks. In SP security services, authorization is based on the principal identifier. The granting of access rights to a principal.

authorization file. A type of ACL (access control list) used by the IBM AIX remote commands and the IBM PSSP Sysctl and Hardmon components.

B

batch processing. (1) The processing of data or the accomplishment of jobs accumulated in advance in such a manner that each accumulation thus formed is processed or accomplished in the same run. (2) The processing of data accumulating over a period of time. (3) Loosely, the execution of computer programs serially. (4) Computer programs executed in the background.

BOS. The AIX Base Operating System.

C

call home function. The ability of a system to call the IBM support center and open a PMR to have a repair scheduled.

CDE. Common Desktop Environment. A graphical user interface for UNIX.

charge feature. An optional feature for either software or hardware for which there is a charge.

CLI. Command Line Interface.

client. (1) A function that requests services from a server and makes them available to the user. (2) A term used in an environment to identify a machine that uses the resources of the network.

CMI. Centralized Management Interface provides a series of SMIT menus and dialogues used for defining and querying the SP system configuration.

Concurrent Virtual Shared Disk. A virtual shared disk that can be concurrently accessed by more than one server.

connectionless. A communication process that takes place without first establishing a connection.

connectionless network. A network in which the sending logical node must have the address of the receiving logical node before information interchange can begin. The packet is routed through nodes in the network based on the destination address in the packet. The originating source does not receive an acknowledgment that the packet was received at the destination.

control workstation. A single point of control allowing the administrator or operator to monitor and manage the SP system using the IBM AIX Parallel System Support Programs.

credentials. A protocol message, or part thereof, containing a ticket and an authenticator supplied by a client and used by a server to verify the client's identity.

css. Communication subsystem.

D

daemon. A process, not associated with a particular user, that performs system-wide functions such as administration and control of networks, execution of time-dependent activities, line printer spooling and so forth.

DASD. Direct Access Storage Device. Storage for input/output data.

DCE. Distributed Computing Environment.

DFS. distributed file system. A subset of the IBM Distributed Computing Environment.

DNS. Domain Name Service. A hierarchical name service which maps high level machine names to IP addresses.

E

Error Notification Object. An object in the SDR that is matched with an error log entry. When an error log entry occurs that matches the Notification Object, a user-specified action is taken.

ESCON. Enterprise Systems Connection. The ESCON channel connection allows the RS/6000 to communicate directly with a host System/390; the host operating system views the system unit as a control unit.

Ethernet. (1) Ethernet is the standard hardware for TCP/IP local area networks in the UNIX marketplace. It is a 10-megabit per second baseband type LAN that allows multiple stations to access the transmission medium at will without prior coordination, avoids contention by using carrier sense and deference, and resolves contention by collision detection (CSMA/CD). (2) A passive coaxial cable whose interconnections contain devices or components, or both, that are all active. It uses CSMA/CD technology to provide a best-effort delivery system.

Ethernet network. A baseband LAN with a bus topology in which messages are broadcast on a coaxial cabling using the carrier sense multiple access/collision detection (CSMA/CD) transmission method.

event. In Event Management, the notification that an expression evaluated to true. This evaluation occurs each time an instance of a resource variable is observed.

expect. Programmed dialogue with interactive programs.

expression. In Event Management, the relational expression between a resource variable and other elements (such as constants or the previous value of an instance of the variable) that, when true, generates an event. An example of an expression is `X < 10` where X represents the resource variable `IBM.PSSP.aixos.PagSp.%totalfree` (the percentage of total free paging space). When the expression is true, that is, when the total free paging space is observed to be less than 10%, the Event Management subsystem generates an event to notify the appropriate application.

F

failover. Also called fallover, the sequence of events when a primary or server machine fails and a secondary or backup machine assumes the primary workload. This is a disruptive failure with a short recovery time.

fall back. Also called fallback, the sequence of events when a primary or server machine takes back control of its workload from a secondary or backup machine.

FDDI. Fiber Distributed Data Interface.

FFDC. First Failure Data Capture.

Fiber Distributed Data Interface (FDDI). An American National Standards Institute (ANSI) standard for 100-megabit-per-second LAN using optical fiber cables. An FDDI local area network (LAN) can be up to 100 km (62 miles) and can include up to 500 system units. There can be up to 2 km (1.24 miles) between system units and concentrators.

file. A set of related records treated as a unit, for example, in stock control, a file could consist of a set of invoices.

file name. A CMS file identifier in the form of 'filename filetype filemode' (like: TEXT DATA A).

file server. A centrally located computer that acts as a storehouse of data and applications for numerous users of a local area network.

File Transfer Protocol (FTP). The Internet protocol (and program) used to transfer files between hosts. It is an application layer protocol in TCP/IP that uses TELNET and TCP protocols to transfer bulk-data files between machines or hosts.

First Failure Data Capture (FFDC). A set of utilities used for recording persistent records of failures and significant software incidents. It provides a means of

associating failures to one another, thus allowing software to link effects of a failure to their causes and thereby facilitating discovery of the root cause of a failure.

foreign host. Any host on the network other than the local host.

FTP. File transfer protocol.

G

gateway. An intelligent electronic device interconnecting dissimilar networks and providing protocol conversion for network compatibility. A gateway provides transparent access to dissimilar networks for nodes on either network. It operates at the session presentation and application layers.

H

HACMP. High Availability Cluster Multi-Processing for AIX.

HACWS. High Availability Control Workstation function, based on HACMP, provides for a backup control workstation for the SP system.

Hardware Management Console (HMC). The *IBM Hardware Management Console for pSeries* is an installation and service support processor that runs only the HMC software. For an IBM @server pSeries 690 server to run the PSSP software, an HMC is required with a network connection to the PSSP control workstation . The HMC provides the following functions for the p690 server:

- Creating and maintaining a multiple partition environment
- Detecting, reporting, and storing changes in hardware conditions
- Acting as a focal point for service representatives to determine an appropriate service strategy

Hashed Shared Disk (HSD). The data striping device for the IBM Virtual Shared Disk. The device driver lets application programs stripe data across physical disks in multiple IBM Virtual Shared Disks, thus reducing I/O bottlenecks.

help key. In the SP graphical interface, the key that gives you access to the SP graphical interface help facility.

High Availability Cluster Multi-Processing. An IBM facility to cluster nodes or components to provide high availability by eliminating single points of failure.

HiPPI. High Performance Parallel Interface. RS/6000 units can attach to a HiPPI network as defined by the ANSI specifications. The HiPPI channel supports burst rates of 100 Mbps over dual simplex cables; connections can be up to 25 km in length as defined by the standard and can be extended using third-party HiPPI switches and fiber optic extenders.

home directory. The directory associated with an individual user.

host. A computer connected to a network, and providing an access method to that network. A host provides end-user services.

HMC. Hardware Management Console.

I

instance vector. Obsolete term for resource identifier.

Intermediate Switch Board. Switches mounted in the switch expansion frame.

Internet. A specific inter-network consisting of large national backbone networks such as APARANET, MILNET, and NSFnet, and a myriad of regional and campus networks all over the world. The network uses the TCP/IP protocol suite.

Internet Protocol (IP). (1) A protocol that routes data through a network or interconnected networks. IP acts as an interface between the higher logical layers and the physical network. This protocol, however, does not provide error recovery, flow control, or guarantee the reliability of the physical network. IP is a connectionless protocol. (2) A protocol used to route data from its source to it destination in an Internet environment.

IP address. A 32-bit address assigned to devices or hosts in an IP internet that maps to a physical address. The IP address is composed of a network and host portion.

ISB. Intermediate Switch Board.

K

Kerberos. A service for authenticating users in a network environment.

kernel. The core portion of the UNIX operating system which controls the resources of the CPU and allocates them to the users. The kernel is memory-resident, is said to run in “kernel mode” and is protected from user tampering by the hardware.

Kernel Low-Level Application Programming Interface (KLAPI). KLAPI provides transport service for communication using the SP Switch.

L

LAN. (1) Acronym for Local Area Network, a data network located on the user’s premises in which serial

transmission is used for direct data communication among data stations. (2) Physical network technology that transfers data a high speed over short distances. (3) A network in which a set of devices is connected to another for communication and that can be connected to a larger network.

LAPI. Low-level Communication API.

local host. The computer to which a user's terminal is directly connected.

log database. A persistent storage location for the logged information.

log event. The recording of an event.

log event type. A particular kind of log event that has a hierarchy associated with it.

logging. The writing of information to persistent storage for subsequent analysis by humans or programs.

Low-level Communication API (LAPI). A low-level (low overhead) message passing protocol that uses a one-sided active message style interface to transfer messages between processes. LAPI is an IBM proprietary interface designed to exploit the SP switch adapters.

M

mask. To use a pattern of characters to control retention or elimination of portions of another pattern of characters.

menu. A display of a list of available functions for selection by the user.

Message Passing Interface (MPI). An industry standard message passing protocol that typically uses a two-sided send-receive model to transfer messages between processes.

Motif. The graphical user interface for OSF, incorporating the X Window System. Also called OSF/Motif.

MPI. Message Passing Interface.

MTBF. Mean time between failure. This is a measure of reliability.

MTTR. Mean time to repair. This is a measure of serviceability.

N

naive application. An application with no knowledge of a server that fails over to another server. Client to server retry methods are used to reconnect.

network. An interconnected group of nodes, lines, and terminals. A network provides the ability to transmit data to and receive data from other systems and users.

Network Interface Module (NIM). A process used by the Topology Services daemon to monitor each network interface.

NFS. Network File System. NFS allows different systems (UNIX or non-UNIX), different architectures, or vendors connected to the same network, to access remote files in a LAN environment as though they were local files.

NIM. (1) Network Installation Management is provided with AIX to install AIX on the nodes. (2) Network Interface Module is a process used by the Topology Services daemon to monitor each network interface.

NIM client. An AIX system installed and managed by a NIM master. NIM supports three types of clients:
- Standalone
- Diskless
- Dataless

NIM master. An AIX system that can install one or more NIM clients. An AIX system must be defined as a NIM master before defining any NIM clients on that system. A NIM master managers the configuration database containing the information for the NIM clients.

NIM object. A representation of information about the NIM environment. NIM stores this information as objects in the NIM database. The types of objects are:
- Network
- Machine
- Resource

NIS. Network Information System.

node. In a network, the point where one or more functional units interconnect transmission lines. A computer location defined in a network. The SP system can house several different types of nodes for both serial and parallel processing. These node types can include thin nodes, wide nodes, 604 high nodes, as well as other types of nodes both internal and external to the SP frame.

Node Switch Board. Switches mounted on frames that contain nodes.

NSB. Node Switch Board.

NTP. Network Time Protocol.

O

ODM. Object Data Manager. In AIX, a hierarchical object-oriented database for configuration data.

P

parallel environment. A system environment where message passing or SP resource manager services are used by the application.

Parallel Environment. A licensed IBM program used for message passing applications on the SP or RS/6000 platforms.

parallel processing. A multiprocessor architecture which allows processes to be allocated to tightly coupled multiple processors in a cooperative processing environment, allowing concurrent execution of tasks.

parameter. (1) A variable that is given a constant value for a specified application and that may denote the application. (2) An item in a menu for which the operator specifies a value or for which the system provides a value when the menu is interpreted. (3) A name in a procedure that is used to refer to an argument that is passed to the procedure. (4) A particular piece of information that a system or application program needs to process a request.

partition. See system partition.

Perl. Practical Extraction and Report Language.

perspective. The primary window for each SP Perspectives application, so called because it provides a unique view of an SP system.

pipe. A UNIX utility allowing the output of one command to be the input of another. Represented by the | symbol. It is also referred to as filtering output.

PMR. Problem Management Report.

POE. Formerly Parallel Operating Environment, now Parallel Environment for AIX.

port. (1) An end point for communication between devices, generally referring to physical connection. (2) A 16-bit number identifying a particular TCP or UDP resource within a given TCP/IP node.

predicate. Obsolete term for expression.

Primary node or machine. (1) A device that runs a workload and has a standby device ready to assume the primary workload if that primary node fails or is taken out of service. (2) A node on the switch that initializes, provides diagnosis and recovery services, and performs other operations to the switch network. (3) In IBM Virtual Shared Disk function, when physical disks are connected to two nodes (twin-tailed), one node is designated as the primary node for each disk and the other is designated the secondary, or backup, node. The primary node is the server node for IBM Virtual Shared Disks defined on the physical disks under normal conditions. The secondary node can become the server node for the disks if the primary node is unavailable (off-line or down).

Problem Management Report. The number in the IBM support mechanism that represents a service incident with a customer.

process. (1) A unique, finite course of events defined by its purpose or by its effect, achieved under defined conditions. (2) Any operation or combination of operations on data. (3) A function being performed or waiting to be performed. (4) A program in operation. For example, a daemon is a system process that is always running on the system.

protocol. A set of semantic and syntactic rules that defines the behavior of functional units in achieving communication.

R

RAID. Redundant array of independent disks.

rearm expression. In Event Management, an expression used to generate an event that alternates with an original event expression in the following way: the event expression is used until it is true, then the rearm expression is used until it is true, then the event expression is used, and so on. The rearm expression is commonly the inverse of the event expression (for example, a resource variable is on or off). It can also be used with the event expression to define an upper and lower boundary for a condition of interest.

rearm predicate. Obsolete term for rearm expression.

remote host. *See foreign host.*

resource. In Event Management, an entity in the system that provides a set of services. Examples of resources include hardware entities such as processors, disk drives, memory, and adapters, and software entities such as database applications, processes, and file systems. Each resource in the system has one or more attributes that define the state of the resource.

resource identifier. In Event Management, a set of elements, where each element is a name/value pair of the form `name=value`, whose values uniquely identify the copy of the resource (and by extension, the copy of the resource variable) in the system.

resource monitor. A program that supplies information about resources in the system. It can be a command, a daemon, or part of an application or subsystem that manages any type of system resource.

resource variable. In Event Management, the representation of an attribute of a resource. An example of a resource variable is `IBM.AIX.PagSp.%totalfree`, which represents the percentage of total free paging

space. `IBM.AIX.PagSp` specifies the resource name and `%totalfree` specifies the resource attribute.

Restricted Root Access (RRA). Restricted root access (RRA) limits the uses of the **rsh** and **rcp** commands within PSSP software. When RRA is enabled, it restricts root **rsh** and **rcp** authorizations from the nodes to the control workstation, and from one node to another. However, control workstation to node **rsh** and **rcp** access is still permitted.

RISC. Reduced Instruction Set Computing (RISC), the technology for today's high performance personal computers and workstations, was invented in 1975. Uses a small simplified set of frequently used instructions for rapid execution.

rlogin (remote LOGIN). A service offered by Berkeley UNIX systems that allows authorized users of one machine to connect to other UNIX systems across a network and interact as if their terminals were connected directly. The rlogin software passes information about the user's environment (for example, terminal type) to the remote machine.

RPC. Acronym for Remote Procedure Call, a facility that a client uses to have a server execute a procedure call. This facility is composed of a library of procedures plus an XDR.

RRA. Restricted Root Access.

RSH. A variant of RLOGIN command that invokes a command interpreter on a remote UNIX machine and passes the command line arguments to the command interpreter, skipping the LOGIN step completely. See also *rlogin*.

S

SCSI. Small Computer System Interface.

Secondary node. In IBM Virtual Shared Disk function, when physical disks are connected to two nodes (twin-tailed), one node is designated as the primary node for each disk and the other is designated as the secondary, or backup, node. The secondary node acts as the server node for the IBM Virtual Shared disks defined on the physical disks if the primary node is unavailable (off-line or down).

server. (1) A function that provides services for users. A machine may run client and server processes at the same time. (2) A machine that provides resources to the network. It provides a network service, such as disk storage and file transfer, or a program that uses such a service. (3) A device, program, or code module on a network dedicated to providing a specific service to a network. (4) On a LAN, a data station that provides facilities to other data stations. Examples are file server, print server, and mail server.

shell. The shell is the primary user interface for the UNIX operating system. It serves as command language interpreter, programming language, and allows foreground and background processing. There are three different implementations of the shell concept: Bourne, C and Korn.

Small Computer System Interface (SCSI). An input and output bus that provides a standard interface for the attachment of various direct access storage devices (DASD) and tape drives to the RS/6000.

Small Computer Systems Interface Adapter (SCSI Adapter). An adapter that supports the attachment of various direct-access storage devices (DASD) and tape drives to the RS/6000.

SMIT. The System Management Interface Toolkit is a set of menu driven utilities for AIX that provides functions such as transaction login, shell script creation, automatic updates of object database, and so forth.

SNMP. Simple Network Management Protocol. (1) An IP network management protocol that is used to monitor attached networks and routers. (2) A TCP/IP-based protocol for exchanging network management information and outlining the structure for communications among network devices.

socket. (1) An abstraction used by Berkeley UNIX that allows an application to access TCP/IP protocol functions. (2) An IP address and port number pairing. (3) In TCP/IP, the Internet address of the host computer on which the application runs, and the port number it uses. A TCP/IP application is identified by its socket.

standby node or machine. A device that waits for a failure of a primary node in order to assume the identity of the primary node. The standby machine then runs the primary's workload until the primary is back in service.

subnet. Shortened form of subnetwork.

subnet mask. A bit template that identifies to the TCP/IP protocol code the bits of the host address that are to be used for routing for specific subnetworks.

subnetwork. Any group of nodes that have a set of common characteristics, such as the same network ID.

subsystem. A software component that is not usually associated with a user command. It is usually a daemon process. A subsystem will perform work or provide services on behalf of a user request or operating system request.

SUP. Software Update Protocol.

switch capsule. A group of SP frames consisting of a switched frame and its companion non-switched frames.

Sysctl. Secure System Command Execution Tool. An authenticated client/server system for running commands remotely and in parallel.

syslog. A BSD logging system used to collect and manage other subsystem's logging data.

System Administrator. The user who is responsible for setting up, modifying, and maintaining the SP system.

system partition. A group of nonoverlapping nodes on a switch chip boundary that act as a logical SP system.

T

tar. Tape ARchive, is a standard UNIX data archive utility for storing data on tape media.

Tcl. Tool Command Language.

TclX. Tool Command Language Extended.

TCP. Acronym for Transmission Control Protocol, a stream communication protocol that includes error recovery and flow control.

TCP/IP. Acronym for Transmission Control Protocol/Internet Protocol, a suite of protocols designed to allow communication between networks regardless of the technologies implemented in each network. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It assumes that the underlying protocol is the Internet Protocol.

Telnet. Terminal Emulation Protocol, a TCP/IP application protocol that allows interactive access to foreign hosts.

ticket. An encrypted protocol message used to securely pass the identity of a user from a client to a server.

Tk. Tcl-based Tool Kit for X Windows.

TMPCP. Tape Management Program Control Point.

token-ring. (1) Network technology that controls media access by passing a token (special packet or frame) between media-attached machines. (2) A network with a ring topology that passes tokens from one attaching device (node) to another. (3) The IBM Token-Ring LAN connection allows the RS/6000 system unit to participate in a LAN adhering to the IEEE 802.5 Token-Passing Ring standard or the ECMA standard 89 for Token-Ring, baseband LANs.

transaction. An exchange between the user and the system. Each activity the system performs for the user is considered a transaction.

transceiver (transmitter-receiver). A physical device that connects a host interface to a local area network, such as Ethernet. Ethernet transceivers contain electronics that apply signals to the cable and sense collisions.

transfer. To send data from one place and to receive the data at another place. Synonymous with move.

transmission. The sending of data from one place for reception elsewhere.

TURBOWAYS 100 ATM Adapter. An IBM high-performance, high-function intelligent adapter that provides dedicated 100 Mbps ATM (asynchronous transfer mode) connection for high-performance servers and workstations.

U

UDP. User Datagram Protocol.

UNIX operating system. An operating system developed by Bell Laboratories that features multiprogramming in a multiuser environment. The UNIX operating system was originally developed for use on minicomputers, but has been adapted for mainframes and microcomputers. **Note:** The AIX operating system is IBM's implementation of the UNIX operating system.

user. Anyone who requires the services of a computing system.

User Datagram Protocol (UDP). (1) In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. UDP is used for application-to-application programs between TCP/IP host systems. (2) A transport protocol in the Internet suite of protocols that provides unreliable, connectionless datagram service. (3) The Internet Protocol that enables an application programmer on one machine or process to send a datagram to an application program on another machine or process.

user ID. A nonnegative integer, contained in an object of type *uid_t*, that is used to uniquely identify a system user.

V

Virtual Shared Disk, IBM. The function that allows application programs executing at different nodes of a system partition to access a raw logical volume as if it were local at each of the nodes. In actuality, the logical volume is local at only one of the nodes (the server node).

W

workstation. (1) A configuration of input/output equipment at which an operator works. (2) A terminal or

microcomputer, usually one that is connected to a mainframe or to a network, at which a user can perform applications.

X

X Window System. A graphical user interface product.

Bibliography

This bibliography helps you find product documentation related to the RS/6000 SP hardware and software products.

You can find most of the IBM product information for RS/6000 SP products on the World Wide Web. Formats for both viewing and downloading are available.

PSSP documentation is shipped with the PSSP product in a variety of formats and can be installed on your system. The man pages for public code that PSSP includes are also available online.

Finally, this bibliography contains a list of non-IBM publications that discuss parallel computing and other topics related to the RS/6000 SP.

Information formats

Documentation supporting RS/6000 SP software licensed programs is no longer available from IBM in hardcopy format. However, you can view, search, and print documentation in the following ways:

- On the World Wide Web
- Online from the product media or the SP Resource Center

Finding documentation on the World Wide Web

Most of the RS/6000 SP hardware and software books are available from the IBM Web site at:

http://www.ibm.com/servers/eserver/pseries

You can view a book or download a Portable Document Format (PDF) version of it. At the time this manual was published, the Web address of the ″RS/6000 SP Hardware and Software Books″ page was:

http://www.rs6000.ibm.com/resource/aix_resource/sp_books

However, the structure of the RS/6000 Web site can change over time.

Accessing PSSP documentation online

On the same medium as the PSSP product code, IBM ships PSSP man pages, HTML files, and PDF files. In order to use these publications, you must first install the **ssp.docs** file set.

To view the PSSP HTML publications, you need access to an HTML document browser such as Netscape. The HTML files and an index that links to them are installed in the **/usr/lpp/ssp/html** directory. Once installed, you can also view the HTML files from the RS/6000 SP Resource Center.

If you have installed the SP Resource Center on your SP system, you can access it by entering the **/usr/lpp/ssp/bin/resource_center** command. If you have the SP Resource Center on CD-ROM, see the **readme.txt** file for information about how to run it.

To view the PSSP PDF publications, you need access to the Adobe Acrobat Reader. The Acrobat Reader is shipped with the AIX Bonus Pack and is also freely available for downloading from the Adobe Web site at:

http://www.adobe.com

To successfully print a large PDF file (approximately 300 or more pages) from the Adobe Acrobat reader, you may need to select the “Download Fonts Once” button on the Print window.

Manual pages for public code

The following manual pages for public code are available in this product:

SUP /usr/lpp/ssp/man/man1/sup.1

Perl (Version 4.036)

/usr/lpp/ssp/perl/man/perl.man

/usr/lpp/ssp/perl/man/h2ph.man

/usr/lpp/ssp/perl/man/s2p.man

/usr/lpp/ssp/perl/man/a2p.man

Manual pages and other documentation for **Tcl**, **TclX**, **Tk**, and **expect** can be found in the compressed **tar** files located in the **/usr/lpp/ssp/public** directory.

RS/6000 SP planning publications

This section lists the IBM product documentation for planning for the IBM RS/6000 SP hardware and software.

IBM RS/6000 SP:

- *Planning, Volume 1, Hardware and Physical Environment*, GA22-7280
- *Planning, Volume 2, Control Workstation and Software Environment*, GA22-7281

RS/6000 SP hardware publications

This section lists the IBM product documentation for the IBM RS/6000 SP hardware.

IBM RS/6000 SP:

- *Planning, Volume 1, Hardware and Physical Environment*, GA22-7280
- *Planning, Volume 2, Control Workstation and Software Environment*, GA22-7281
- *Installation and Relocation*, GA22-7441
- *System Service Guide*, GA22-7442
- *SP Switch Service Guide*, GA22-7443
- *SP Switch2 Service Guide*, GA22-7444
- *Uniprocessor Node Service Guide*, GA22-7445
- *604 and 604e SMP High Node Service Guide*, GA22-7446
- *SMP Thin and Wide Node Service Guide*, GA22-7447
- *POWER3 SMP High Node Service Guide*, GA22-7448

RS/6000 SP Switch Router publications

The RS/6000 SP Switch Router is based on the Ascend GRF switched IP router product from Lucent Technologies. You can order the SP Switch Router as the IBM 9077.

The following publications are shipped with the SP Switch Router. You can also order these publications from IBM using the order numbers shown.

- *Ascend GRF GateD Manual*, GA22-7327
- *Ascend GRF 400/1600 Getting Started*, GA22-7368
- *Ascend GRF Configuration and Management*, GA22-7366
- *Ascend GRF Reference Guide*, GA22-7367
- *SP Switch Router Adapter Guide*, GA22-7310

Related hardware publications

For publications on the latest IBM @server pSeries and RS/6000 hardware products, see the Web site:

http://www.ibm.com/servers/eserver/pseries/library/hardware_docs/

That site includes links to the following:

- General service documentation
- Guides by system (pSeries and RS/6000)
- Installable options
- IBM Hardware Management Console for pSeries guides

RS/6000 SP software publications

This section lists the IBM product documentation for software products related to the IBM RS/6000 SP. These products include:

- IBM Parallel System Support Programs for AIX (PSSP)
- IBM LoadLeveler for AIX 5L (LoadLeveler)
- IBM Parallel Environment for AIX (Parallel Environment)
- IBM General Parallel File System for AIX (GPFS)
- IBM Engineering and Scientific Subroutine Library (ESSL) for AIX
- IBM Parallel ESSL for AIX
- IBM High Availability Cluster Multi-Processing for AIX (HACMP)

PSSP Publications

IBM RS/6000 SP:

- *Planning, Volume 2, Control Workstation and Software Environment*, GA22-7281

PSSP:

- *Installation and Migration Guide*, GA22-7347
- *Administration Guide*, SA22-7348
- *Managing Shared Disks*, SA22-7349
- *Diagnosis Guide*, GA22-7350
- *Command and Technical Reference*, SA22-7351

- *Messages Reference*, GA22-7352
- *Implementing a Firewalled RS/6000 SP System*, GA22-7874

RS/6000 Cluster Technology (RSCT):
- *Event Management Programming Guide and Reference*, SA22-7354
- *Group Services Programming Guide and Reference*, SA22-7355
- *First Failure Data Capture Programming Guide and Reference*, SA22-7454

LoadLeveler Publications

LoadLeveler:
- *Using and Administering*, SA22-7311
- *Diagnosis and Messages Guide*, GA22-7277

GPFS Publications

GPFS:
- *Problem Determination Guide*, GA22-7434
- *Administration and Programming Reference*, SA22-7452
- *Concepts, Planning, and Installation*, GA22-7453

Parallel Environment Publications

Parallel Environment:
- *Installation Guide*, GA22-7418
- *Messages*, GA22-7419
- *MPI Programming Guide*, SA22-7422
- *MPI Subroutine Reference*, SA22-7423
- *Hitchhiker's Guide*, SA22-7424
- *Operation and Use, Volume 1*, SA22-7425
- *Operation and Use, Volume 2*, SA22-7426

Parallel ESSL and ESSL Publications
- *ESSL Products: General Information*, GC23-0529
- *Parallel ESSL: Guide and Reference*, SA22-7273
- *ESSL: Guide and Reference*, SA22-7272

HACMP Publications

HACMP:
- *Concepts and Facilities*, SC23-4276
- *Planning Guide*, SC23-4277
- *Installation Guide*, SC23-4278
- *Administration Guide*, SC23-4279
- *Troubleshooting Guide*, SC23-4280
- *Programming Locking Applications*, SC23-4281
- *Programming Client Applications*, SC23-4282
- *Master Index and Glossary*, SC23-4285
- *HANFS for AIX Installation and Administration Guide*, SC23-4283

- *Enhanced Scalability Installation and Administration Guide, Volume 1*, SC23-4284
- *Enhanced Scalability Installation and Administration Guide, Volume 2*, SC23-4306

AIX publications

You can find links to the latest AIX publications on the Web site:

http://www.ibm.com/servers/aix/library/techpubs.html

DCE publications

The DCE library consists of the following books:

- *IBM DCE for AIX: Administration Commands Reference*
- *IBM DCE for AIX: Administration Guide—Introduction*
- *IBM DCE for AIX: Administration Guide—Core Components*
- *IBM DCE for AIX: DFS Administration Guide and Reference*
- *IBM DCE for AIX: Application Development Guide—Introduction and Style Guide*
- *IBM DCE for AIX: Application Development Guide—Core Components*
- *IBM DCE for AIX: Application Development Guide—Directory Services*
- *IBM DCE for AIX: Application Development Reference*
- *IBM DCE for AIX: Problem Determination Guide*
- *IBM DCE for AIX: Release Notes*

You can view a DCE book or download a Portable Document Format (PDF) version of it from the IBM DCE Web site at:

http://www.ibm.com/software/network/dce/library

Redbooks

IBM's International Technical Support Organization (ITSO) has published a number of redbooks related to the RS/6000 SP. For a current list, see the ITSO Web site at:

http://www.ibm.com/redbooks

Non-IBM publications

Here are some non-IBM publications that you might find helpful.

- Almasi, G., Gottlieb, A., *Highly Parallel Computing*, Benjamin-Cummings Publishing Company, Inc., 1989.
- Foster, I., *Designing and Building Parallel Programs*, Addison-Wesley, 1995.
- Gropp, W., Lusk, E., Skjellum, A., *Using MPI*, The MIT Press, 1994.
- Message Passing Interface Forum, *MPI: A Message-Passing Interface Standard, Version 1.1*, University of Tennessee, Knoxville, Tennessee, June 6, 1995.
- Message Passing Interface Forum, *MPI-2: Extensions to the Message-Passing Interface, Version 2.0*, University of Tennessee, Knoxville, Tennessee, July 18, 1997.
- Ousterhout, John K., *Tcl and the Tk Toolkit*, Addison-Wesley, Reading, MA, 1994, ISBN 0-201-63337-X.
- Pfister, Gregory, F., *In Search of Clusters*, Prentice Hall, 1998.

- Barrett, D., Silverman, R., *SSH The Secure Shell The Definitive Guide*, O'Reilly, 2001.

Index

Special Characters

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

U

V

W

Readers' Comments — We'd Like to Hear from You

Parallel System Support Programs for AIX
Installation and Migration Guide
Version 3 Release 4

Publication No. GA22-7347-03

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name _____ Address _____

Company or Organization _____

Phone No. _____

Readers' Comments — We'd Like to Hear from You
GA22-7347-03

Cut or Fold
Along Line

Fold and Tape **Please do not staple** Fold and Tape

PLACE
POSTAGE
STAMP
HERE

IBM Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie NY 12601-5400

Fold and Tape **Please do not staple** Fold and Tape

GA22-7347-03

Cut or Fold
Along Line

Program Number: 5765-D51

GA22-7347-03