

SAN Boot configuration for pSeries Linux servers in DS4000 Storage Subsystems



This document describes how to remotely install (SAN boot) a Linux operating system onto a pSeries host server that is connected by Fibre Channel to a DS4000 Storage Subsystem.

Software requirements

You can use this SAN boot method to install one of the following Linux operating systems. The software levels listed are minimum versions.

SUSE Linux Enterprise Server 9 (SLES) Service Pack 1 (or later)

- Kernel: 2.6.5-7.139
- RDAC driver: 09.01.B5.02

Red Hat Enterprise Linux 3, Update 4 (or later)

- Kernel: 2.4.21-27.EL.img (ppc64 Architecture)
- RDAC driver: 09.10.A5.06

Hardware requirements

The following hardware is required:

- IBM @server pSeries server(s)
- Supported IBM pSeries host bus adapters (HBAs)
- DS4000 Storage Subsystem

Assumptions

This procedure assumes the following:

- No installation is present.
- After the installation is complete, the boot and swap devices will run on a DS4000 Storage Subsystem.

Preparing for the SAN boot procedure

Complete the following steps before you begin the installation:

1. Be familiar with the following topics:
 - How to connect to and operate an IBM pSeries server
 - How to create LUNs on the DS4000 Storage Subsystem

- How to set up a SAN network or direct-attached SAN storage devices, so that the configured server can access the DS4000 LUNs
 - How the Linux kernel boot process functions, and what processes and procedures are used to boot a Linux distribution for a local storage device
2. Ensure that you have a copy of your Red Hat or SLES operating system installation file available, either on CD or on an accessible network.
 3. Verify that the BOOT BIOS is enabled on the HBA devices that are configured for the server. (The BOOT BIOS setting enables discovery and use of SAN disk devices during the installation process.)
 4. Ensure that the size of the boot device that you plan to use is at least 4 GB. (This is the minimum size for installing all packages from the installation media to the boot devices.)

In addition, ensure that the size of the swap device is equivalent to or larger than the size of the server's physical memory.

Note: For simplicity, the procedure in this document assumes that the boot, root, and swap devices are all located on the same device. However, this is not a requirement for the installation.

Steps for performing the SAN boot procedure

The installation media (source for installation) that is used in the following steps is an NFS-exported set of CDs. However, you can use any of the following installation sources:

- CD
- NFS
- HTTP
- FTP
- Other source

Note: Not all of the details about installation and partitioning are documented in this procedure. To determine which packages are needed for the type of system you are installing, you must refer to installation procedures.

Complete the following steps.

1. From the System Management Services menu (SMS), select your installation source, and boot from the media. (See your server's installation manual for further instructions.)
2. Verify that the HBA module is loaded. (For Red Hat, the module is **lpfc**; for SLES, the module is **lpfcdd**.)
3. Verify that the SAN devices you are using for the installation have been detected successfully.

Note: If the SAN devices are already configured for multiple path access, Linux discovers the same physical device multiple times -- once for each logical path to the device.

Red Hat only: The Red Hat Enterprise Linux 2.4 kernel does not probe past LUN 0 if the device is not a well-known device. The devices are part of `scsi_scan.c`. Therefore, for newer devices, you need to identify your boot disk as LUN 0.

4. Note which SAN device will be used for the installation (for example, `/dev/sdh`.) Also, note which HBA device is used to discover the SAN device. (You will need this information in step 5.)

5. Select the installation options according to your operating system, as follows.

Note: Installation options include partitioning, which is necessary to ensure that the device that you noted in step 4 is used as the root/boot installation target.

For Red Hat: Select the installation options that are appropriate for your configuration. For partitioning options, select `Autopartition` then select the drive you are using as the installation target. Then, continue to the steps listed in “Completing the Red Hat procedure” on page 4.

For SLES: Select the installation options that are appropriate for your configuration until the Installation Settings dialog displays. Then complete the following steps:

- a. In the Installation Settings dialog, select `Change`.
- b. In the next dialog that displays, select `Partitioning`.
- c. In the Custom Partitioning dialog, select the device, then select `Custom partitioning for experts`.
- d. Ensure that there is a PReP boot partition on the root/boot device, and that it is the first partition.
- e. Continue to partition devices as required to complete this configuration.
- f. Finish the installation. The system reboots.

Note: When attempting to create the yaboot boot loader, an error occurs that states that the 'fcp' device type is unknown. Click **OK**, and then select `No` when asked to retry.

- g. Open the SMS menu to boot from the installation source media. (If you are booting from a CD, continue to a point where you can abort the installation and return to the command-line menu. If you are booting from the network, the command-line menu is already present.)
- h. Select the option to boot from an installed system, select the root device that you just installed, and continue.
- i. When the YaST installation wizard displays, use the wizard to finish the installation.
- j. Continue to “Completing the SLES procedure.”

Completing the SLES procedure

After you finish the previous steps, you must complete the following steps to finish the SAN boot procedure for the SLES system:

1. Use YaST to upgrade the installed system to the latest SLES service pack. (As of this writing, the latest service pack is SP 1.)
2. Run the `cat /etc/lilo.conf` command, and verify that the data in this file looks correct for the boot and root partitions.
3. Run the `lilo` command. This permits the installation of the yaboot loader to the PReP boot partition of the drive (resolving the installation error that occurred in the previous steps).
4. Reboot again and open the SMS menu. This time, the boot device that you set up in the previous steps is ready to be used. Select the option to boot from a hard drive/SAN, and select the HBA that is associated with the SAN disk device on which the installation was completed.
5. Set the installed boot device as the default boot device for the system.

Note: This step is not required. However, it is recommended to enable unattended reboots after this procedure is complete.

6. The installation boot device should now be listed in the "bootable" devices that are discovered on the selected HBA. Select the appropriate device, and boot.

If all of the preceding steps were completed accurately, the SLES system is now booted in single-path mode from the SAN. Complete the following steps to verify the installation:

1. Use the **df** command to check the mounted devices and verify that the root is mounted in the appropriate location.
2. Use the **swapon -s** command to verify that the swap device is correctly mounted.
3. Verify that the other configured partitions are also correctly mounted.

This completes the single-path SAN boot procedure for SLES systems.

For information about modifying the boot/root and other devices for booting using the RDAC driver, see the *IBM TotalStorage DS4000 Storage Manager 9 Installation and Support Guide for AIX, HP-UX, Solaris and Linux on POWER (GC26-7705)*.

Completing the Red Hat procedure

After you finish the previous steps, you must complete the following steps to finish the SAN boot procedure for the Red Hat system:

1. Reboot again and open the SMS menu. This time, the boot device that you set up in the previous steps is ready to be used. Select the option to boot from a hard drive/SAN, and select the HBA that is associated with the SAN disk device on which the installation was completed.
2. Set the installed boot device as the default boot device for the system.

Note: This step is not required. However, it is recommended to enable unattended reboots after this procedure is complete.

3. The installation boot device should now be listed in the "bootable" devices that are discovered on the selected HBA. Select the appropriate device, and boot.

If all of the preceding steps were completed accurately, the Red Hat system is now booted in single-path mode from the SAN. Complete the following steps to verify the installation:

1. Check the mounted devices, and verify that the root is mounted in the appropriate location.
2. Verify that the swap device and other configured partitions are correctly mounted.

This completes the single-path SAN boot procedure for Red Hat systems.

For information about modifying the boot/root and other devices for booting using the RDAC driver, see the *IBM TotalStorage DS4000 Storage Manager 9 Installation and Support Guide for AIX, HP-UX, Solaris and Linux on POWER (GC26-7705)*.

First edition

This document describes how to remotely install certain versions of Red Hat Enterprise Linux or SUSE Enterprise Linux onto a pSeries host server that is connected by Fibre Channel to a DS4000 Storage Subsystem.

© Copyright International Business Machines Corporation 2005. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

GC26-7773-00

