

HP StorageWorks

Booting Linux x86 and x86_64 systems from a storage area network application notes

Legal and notice information

© Copyright 2003–2005 Hewlett-Packard Development Company, L.P.

© Copyright 2005 Emulex Corporation

© Copyright 2005 QLogic Corp.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Linux is a U.S. registered trademark of Linus Torvalds.

Booting Linux x86 and x86_64 systems from a storage area network application notes

About this document

These application notes describe boot setup procedures from a storage area network (SAN) with X86 and X86_64 Linux® systems.

For the latest version of this document, access the HP storage web site:
<http://h18006.www1.hp.com/storage/saninfrastructure.html>.

This section describes the content reflected in this document, including:

- [Application notes information](#), page 3
- [Intended audience](#), page 3

Application notes information

These application notes contain the following major topics:

- [Prerequisites](#), page 4
- [Introduction](#), page 5
- [Required hardware and software](#), page 6
- [Boot from SAN procedures](#), page 9
- [Maintenance](#), page 16

Intended audience

This document is intended for storage system administrators who are familiar with:

- Red Hat Enterprise Linux or SUSE LINUX Enterprise Server operating systems
- Host bus adapters (HBAs)
- Fibre Channel switch configuration, management, and zoning
- Storage arrays:
 - Enterprise Modular Array (EMA)
 - Enterprise Virtual Array (EVA)
 - Modular Smart Array (MSA)
 - Disk Array XP (XP)



NOTE: For information about booting to a SAN on Modular Smart Arrays, go to the HP web site <http://h18006.www1.hp.com/storage/arrayssystems.html>.

- Rapid Deployment Pack



NOTE: The Rapid Deployment Pack Support Matrix contains detailed support information. You can access this information on the HP web site <http://www.hp.com/servers/rdp>.

Accessing future product updates

HP strongly recommends that customers sign up online using the Subscriber's choice web site: <http://www.hp.com/go/e-updates>.

Subscribing to this service provides you with e-mail updates on the latest product enhancements, newest versions of drivers, and firmware documentation updates as well as instant access to numerous other product resources.

After signing up, you can quickly locate your products by selecting **Business support** and then **Storage** under Product Category.

Prerequisites

To set up booting from a SAN, you should be familiar with:

- HSG80, EVA, MSA, or XP Fibre Channel RAID arrays. This information is in the RAID array and platform software documentation.
- RA8000, MA8000, EMA12000, EMA16000, ESA 12000
 - Set and/or verify the HSG80 controller port topology for your configuration.
 - Verify that all existing storage units have the proper access.

 **NOTE:** References to the HSG80 are intended to cover all of the RA8000/MA8000 RAID array products.

The following procedures are not described in this document. However, they must be performed before performing boot from SAN procedures.

- EVA—Verify that all existing storage units (virtual disks) have the proper access settings.
- XP—Verify that all existing storage units (virtual disks) have the proper access settings.
- Linux operating systems—Verify that one of the operating systems listed in [Table 1](#) on page 6 is to be installed.
- Switch configurations and zoning—Check the switch to ensure there are no zoning conflicts.
- If you are using the Rapid Deployment Pack, verify that the Rapid Deployment Pack and NFS servers reside in the same network.

Introduction

Traditionally, HP ProLiant servers boot operating systems from internal SCSI and IDE storage devices. With the support of external booting in HP StorageWorks HBAs and RAID arrays, customers can optionally eliminate server-based internal boot devices. Booting from an external device decreases downtime in the event of server failure through faster server replacement.

Why boot to a SAN?

Booting to a SAN provides:

- Improved disaster recovery
- Reduced backup window through the SAN
- Additional SAN-managed features

Additional information

For detailed information about HBAs, switches, storage arrays, and storage software, access the following HP web sites:

Topic	HP web site
HBAs and switches	http://h18006.www1.hp.com/storage/saninfrastructure.html
Storage array systems	http://h18006.www1.hp.com/storage/arraysystems.html
Platform software	http://h18006.www1.hp.com/storage/software.html
ProLiant servers	http://h18004.www1.hp.com/products/servers/platforms/storage.html

Supported configurations

This document describes single-server configurations for EMAs, EVA, MSA, and XP.

Required hardware and software

Table 1 lists the *minimum* required hardware and software for booting from a SAN. HP recommends using the *latest* supported versions for operating systems, firmware, BIOS, and drivers. Check for the latest supported versions on the following web site:

<http://h18006.www1.hp.com/storage/saninfrastructure.html>.

Table 1 Required hardware and software

Supported feature	Requirement
Linux operating systems	<p>QLogic HBAs</p> <ul style="list-style-type: none">• Red Hat Enterprise Linux 3.0 (Update 1 minimum)—Note that booting from EVA8000, EVA6000, or EVA4000 requires Red Hat Enterprise Linux 3.0 (Update 5 minimum).• Red Hat Enterprise Linux 4 (base release minimum)• SUSE Linux Enterprise Server 8/United Linux (SP3 minimum)• SUSE Linux Enterprise Server 9 (SP1 minimum) <p>Emulex HBAs</p> <ul style="list-style-type: none">• Red Hat Enterprise Linux 3.0 (Update 4 minimum)—Note that booting from EVA8000, EVA6000, or EVA4000 requires Red Hat Enterprise Linux 3.0 (Update 5 minimum).• Red Hat Enterprise Linux 4 (base release minimum)• SUSE Linux Enterprise Server 9 (SP1 minimum)
RAID arrays	<p>HP supports booting from SAN using non-shared LUNs on the following arrays:</p> <p>EVA8000, EVA6000, EVA5000, EVA4000, EVA3000 MSA 500 G2, MSA1000, MSA1500 cs XP12000, XP512/48, XP128/1024</p> <p>The following arrays are only supported with Qlogic HBAs: RA8000, MA8000, EMA12000, EMA16000, ESA 12000 Fibre Channel</p>
Fibre Channel switches	<p>B-Series, C-Series, and M-Series switch products. For a list of supported switch models, see the <i>HP StorageWorks SAN design reference guide</i> available on the following web site:</p> <p>http://h18006.www1.hp.caom/products/storageworks/san/documentation.html.</p>
Servers	<p>All HP ProLiant Servers, including BL and ML/DL series, are supported. For detailed information, see: http://h18004.www1.hp.com/products/servers/platforms/storage.html.</p> <p>Be sure that the servers have the latest ROM BIOS.</p>
HBA models	<p>See the HBA release notes for the latest firmware, driver, BIOS versions, and LUN limitations. The release notes are available on the following web site:</p> <p>http://h18006.www1.hp.com/storage/saninfrastructure.html.</p>

Table 1 Required hardware and software (continued)

Supported feature	Requirement
2Gb HBA models	FCA2214 2Gb PCI-X HBA, part number 281541-B21 FCA2214DC 2Gb PCI-X HBA, part number 321835-B21 QLogic Dual Port Fibre Channel Mezzanine HBA 2-Gb for BL20Gp2 G2, part number 300874-B21 QLogic Dual Port Fibre Channel Balcony HBA 2-Gb for BL30Gp, part number 354054-B21 QLogic-based HP BL20p G2, Fibre Channel Mezzanine HBA, part number 300874-B21 QLogic-based HP BL25p/BL45p FCA, part number 381881-B21 Emulex-based Fibre Channel Mezzanine HBA for HP p-class systems for BL30pG1, BL25p, BL35p, and BL45p, part number 394588-B21 Emulex-based Fibre Channel Mezzanine HBA for HP p-class systems for BL20pG3, part number 394757-B21
4Gb HBA models	FC1142SR 4Gb PCI-Express HBA, part number AE311A FC1242SR DC 4Gb PCI-Express HBA, part number AE312A
HBA BIOS and drivers	You must load the latest supported HP drivers after you complete the boot from SAN installation. The drivers are available on the SAN infrastructure web site: http://h18006.www1.hp.com/storage/saninfrastructure.html

Table 1 Required hardware and software (continued)

Supported feature	Requirement
QLogic	2Gb HBAs HBA BIOS: 1.34 minimum HBA drivers: 7.01.01 minimum for 2.4 kernels 8.00.00 minimum for 2.6 kernels 4Gb HBAs HBA BIOS: 1.06 minimum HBA drivers: 7.07.03 minimum for 2.4 kernels 8.01.03 minimum for 2.6 kernels Note the following: <ul style="list-style-type: none">• The HBA firmware loads at driver load time.• HSG80-based storage is supported only with 2Gb HBAs with the 6.06.50 driver in single-path mode with Red Hat AS 2.1 and HBA BIOS 1.34
Emulex	HBA BIOS: 1.70a3 minimum HBA drivers: 7.3.2 minimum for 2.4 kernels 8.0.16.17 minimum for 2.6 kernels HBA firmware: 1.91a2 You must manually update the HBA firmware.

Switch-based zoning

You must be familiar with switch-based zoning. For specific information about zoning for your environment, see your switch's installation guide.

Boot from SAN procedures

Perform the following procedures for setting up your Linux systems with HSG80, EVA, and MSA-based storage to boot from a SAN:

- [Step 1: Set up the storage arrays](#), page 9
- [Step 2: Set up the host systems](#), page 10
 - [Step 2A: Configure the host boot order](#), page 10
 - [Step 2B: Set up switch zoning](#), page 10
 - [Step 2C: Create a LUN](#), page 11
 - [Step 2D: Configure the HBAs](#), page 12
- [Step 3: Select the operating system installation method](#), page 14
 - [Option 1: Use CD-ROM installation](#), page 14
 - [Installing Red Hat Enterprise Linux](#), page 14
 - [Installing SUSE Linux Enterprise Server](#), page 14
 - [Option 2: Use the automated Rapid Deployment Pack installation](#), page 14

Perform the following procedure for XP128/1024-based storage to boot from a SAN:

[Boot setup notes for XP](#), page 15

Step 1: Set up the storage arrays

Before you begin, make sure that you are starting with a clean, unpartitioned virtual disk.

To set up the storage:

1. Create a Virtual Disk (LUN) that is at least 9 Gb.

 **NOTE:** When deploying supported Linux distributions, the only unit that is bootable to a Linux host connection is LUN 1. For example, with HSG80-based storage, Unit D11 is visible and bootable to a host connection on Port 1 that has an offset of 10 on LUN 1. See the documentation for the HP StorageWorks HSG80 platform software for detailed information. In general, the LUN number equals the unit number minus the offset. Therefore, LUN 1=11-10.

2. Present the LUN to the Boot from SAN (BFS) host.

 **NOTE:** You must assign one bootable LUN to the BFS host.

3. Set and/or verify the controller SCSI version.

 **NOTE:** On HSG80-based storage, Linux is supported in SCSI-3 mode only. In this mode, LUN 0 is assigned to the Command Console. Therefore, the next available LUN, the one you assign as your boot disk, must be LUN 1. Changing offsets allows for multiple occurrences of LUN 1.

Step 2: Set up the host systems

Use the following steps to set up all host systems:

- [Step 2A: Configure the host boot order](#), page 10
- [Step 2B: Set up switch zoning](#), page 10
- [Step 2C: Create a LUN](#), page 11
- [Step 2D: Configure the HBAs](#), page 12

Step 2A: Configure the host boot order

1. Verify that you have the latest BIOS.
2. Set the boot order:
 - a. While the system is booting, press **F9** to start the BIOS Setup Utility.
 - b. Select **Boot Controller Order**.
 - c. Select the primary HBA (that is, the HBA dedicated to your SAN or presented to your boot LUN) and move it to **Controller Order 1**.
 - d. Disable the Smart Array Controller.
3. Press **F10** to save your configuration and exit the utility.

Step 2B: Set up switch zoning

Set up initial zoning

1. Create a zone that includes the boot HBA port and at least one active host port from the target storage.
2. Enable the zone.

 **NOTE:** See “[Troubleshooting](#)” on page 16 for information about active-standby controller types. HSG80-based storage and Enterprise Virtual Arrays use this architecture.

Finalize zoning

Configure your adapter as described in “[Step 2D: Configure the HBAs](#)” on page 12.

Step 2C: Create a LUN

Creating a LUN on Enterprise Modular Arrays

1. Using a terminal emulator or a terminal connected to the storage subsystem, enter the following CLI commands to configure a device with no access allowed:

```
HSG> initialize disk10000
HSG> add unit d1 disk10000 disable_access_path=all
HSG> show d1
```

2. Note the controller to which the LUN is recognized as online.
3. Verify that the HBA is online by entering the following command:
4. Optionally, rename the connection.
5. Allow the HBA to access the LUN by entering the following command:

```
HSG> set d1 enable_access_path=!newcon01
```

For more information about LUN Configuration, see the *HP StorageWorks HSG80 installation and configuration guide*.

Creating a LUN on Enterprise Virtual Arrays

1. Access Command View EVA or the HSV Element Manager through a web browser.
2. Create the LUN and assign a host. You must create a host if one does not exist.
3. Return to the server.

For detailed information about using the Enterprise Virtual Array Command View, see the *HP StorageWorks Command View EVA getting started guide*.

Creating a LUN on Modular Smart Arrays

1. Create a virtual disk (LUN) using the Array Configuration Utility (ACU) on a separate server.
2. Using the Selective Storage Presentation option, assign the LUN to the HBA connection.
3. Return to the server.

Creating a LUN on Blade servers attached to Modular Smart Arrays

If you are creating a boot LUN for a blade server, consider using the following procedure to create the LUN if there is no separate server for the ACU. Before you begin:

1. Obtain iLO with firmware 1.40 or later. The latest iLO is available on the server support page: www.hp.com/go/blades.
2. Obtain the latest MSA Support Software CD-ROM from the MSA web site: www.hp.com/go/msa1000.

 **NOTE:** The server you use to browse to the iLO web interface contains the CD-ROM drive for the virtual CD option.

3. Browse to the iLO interface.

4. Install the MSA Support CD-ROM in the drive.
5. Power cycle the Blade server. The server now boots to the MSA Support Software CD-ROM.
6. After creating the LUN, exit ACU and exit the MSA Support Software CD-ROM.
The server reboots.
7. Disconnect the Virtual CD from the Virtual Media page.

For detailed information about using ACU, see the *HP StorageWorks array configuration guide*.

Creating a LUN on XP

Using a web browser, start Command View XP and perform the following steps:

1. Select a port (such as CL1-A).
2. Set LUN security for the array port.
3. Create a host group with the proper host mode (for Linux, host mode= 0x00).
4. Assign the first boot HBA port WWN to the newly created host group.
5. Assign LUN 0 to that host group.

For detailed information about using Command View XP, see the *HP StorageWorks Command View XP for XP disk arrays user guide*.

Step 2D: Configure the HBAs

Qlogic HBAs

1. While the server is booting, press **Ctrl Q** to enter *Fast!UTIL*.
2. Perform the steps in the “[Step 2B: Set up switch zoning](#)” on page 10, and then go to [step 3d](#).
3. From the Select Host Adapter menu, select the HBA you want to boot, and then press **Enter**.
4. From the *Fast!UTIL* Options menu, select **Configuration Settings**, and then press **Enter**.
5. From the Configuration Settings menu, select **Host Adapter Settings**, and then press **Enter**.
6. From the Host Adapter Settings menu, change **Host Adapter BIOS** to **Enabled** by pressing **Enter**.
7. Press **ESC** to go back to the Configuration Settings menu.
8. Select **Selectable Boot Settings**, and then press **Enter**.
9. From the Selectable Boot Settings menu, enable the Selectable Boot option, move the cursor to **Primary Boot Port Name, Lun**, and then press **Enter**.
10. Create a LUN by following the appropriate procedure in “[Step 2C: Create a LUN](#)” on page 11.
11. Select **Primary Boot Port Name, Lun**, then press **Enter**.
12. From the Select Fibre Channel Device menu, select the device to boot from, and then press **Enter**. If more than one device has access, select the supported LUN (the LUN number you created in [step 10](#)).
13. Save the changes by pressing **ESC** twice.

Emulex HBAs

1. Record the IEEE number (WWN) for all HBAs.

2. Configure the HBA:
 - a. While the server is booting, press **Alt E** when the following message displays:
Press <Alt E> To Go To Emulex BIOS Utility.
 - b. Perform the steps in "[Step 2B: Set up switch zoning](#)" on page 10.
 - c. At the Enter a Selection prompt, select the first HBA.
 - d. Select **Option 2, Configure This Adapter's Parameters**.
 - e. Select **Option 1, Enable or Disable BIOS**. Press **1** to enable BIOS. (These steps are repeated for the other HBAs after the Linux operating system has been installed.)
 - f. Select **Option 4, Topology Selection, <Advanced Option>**.
 - g. Select **Point to Point**. Press **Esc**.
 - h. If you are using XP or EVA4000/6000/8000, omit this step. Otherwise, select **Option 8, Enable Start Unit Command, <Advanced Option>**. Press **1** to enable.
 - i. Press **X** to exit and reboot the server.
3. Configure the boot devices:
 - a. While the server is booting, press **Alt E** when the following message displays:
Press <Alt E> To Go To Emulex BIOS Utility.
 - b. At the Enter a Selection prompt, select the first HBA.
 - c. Select **Option 1, Configure Boot Devices**.
 - d. Create a LUN by following the appropriate procedure in "[Step 2C: Create a LUN](#)" on page 11.
 - e. In the list of Saved Boot Devices, select **1** as the primary boot path or **2** as the secondary boot path.
 - f. At the Select Two Digit Number of the Desired Boot Device prompt, select the WWPN for the desired boot port.

 **NOTE:** Each controller port has a WWPN. The WWPN must be identical for all ports except for the last number. Use the last number to determine the controller port from which you are booting.

- g. At the prompt, enter the two digits of starting LUN (hexadecimal). Use the LUN number you created in [step 3d](#).
- h. Select **1, Boot This Device via WWPN**.
- i. Reboot the server.

Step 3: Select the operating system installation method

Select one of the following methods for installing the Linux operating system on your BFS hosts:

- CD-ROM installation
- Automated installation with the Rapid Deployment Pack

Option 1: Use CD-ROM installation

 **NOTE:** Before you begin, make sure you download the latest HBA drivers from <http://h18006.www1.hp.com/storage/saninfrastructure.html>.

Installing Red Hat Enterprise Linux

1. Insert the Red Hat Enterprise Linux CD1 in the BFS host CD-ROM drive and reboot. The installation begins.
2. Follow the installation instructions.
3. Load the operating system to the Boot LUN on the SAN.
4. When the Installation Complete window appears, press **Enter** to complete the installation and reboot the system, completing the configuration.
5. Install the latest HBA drivers that you downloaded.

Installing SUSE Linux Enterprise Server

1. Insert SUSE Linux Enterprise Server CD1 in the BFS host CD-ROM drive and reboot. The installation procedure prompts you for additional CD-ROMs.
2. Follow the instructions.
3. Load the operating system to the Boot LUN on the SAN.
4. Install the latest HBA drivers.

Option 2: Use the automated Rapid Deployment Pack installation

To configure ProLiant Essentials Rapid Deployment Pack to provide automated server deployment and redeployment, follow the instructions in the Rapid Deployment Pack Knowledge Base Article 127 on the HP storage web site:

<http://h18013.www1.hp.com/products/servers/management/rdp/knowledgebase/00000127.html>.

Boot setup notes for XP

For XP, if one server is booting from a SAN, one host group is required for that HBA. If more than one server is booting from a SAN and accessing the same physical port, a unique host group is required for each HBA. Every host group must be assigned on a port with a unique LUN 0. If LUN security is turned off for this port (for example, port CL1-A), all the servers defined in the host groups are suddenly given access to all LUNs contained in the default host group.

For detailed information about Disk Array XP host groups, LUN security, and host modes, see the *HP StorageWorks Command View XP User Guide* available on the following web site:

http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?contentType=SupportManual&locale=en_US&docIndexId=179911&taskId=101&prodTypeId=12169&prodSeriesId=64820

Use the following boot setup procedure for XP128/1024:

1. Perform [Step 1: Set up the storage arrays](#), page 9.
2. Perform [Step 2: Set up the host systems](#), page 10.
3. Perform [Step 3: Select the operating system installation method](#), page 14.

Maintenance

This section includes HBA maintenance and troubleshooting information.

Server replacement in the event of a server failure

You must retain the HBAs to ensure that the replacement server hardware components are identical in every way to the component that you are replacing. Place each HBA into the PCI slots of the new server in the same order as they were placed in the old server.

If you are using an internal disk for your swap file system, you can move this disk to the replacement server. You can also use a new disk, but you must partition it and manually reset the swap to that disk.

If you are replacing a BL20p, you can move the mezzanine HBAs to a replacement server and the WWPNs remain the same. However, you also retain two MAC addresses when you move the mezzanine card. Therefore, you must remap all the LUNs to the onboard Fibre Channel adapters.

Troubleshooting

Install in single-path mode

HP recommends that you initially install the operating system to the SAN in single-path mode. To do so, map the root directory (/) and /boot to /dev/sda (or to the first Fibre Channel device).

After installing the operating system to the SAN, install the supported failover driver, and then configure multiple paths.

Active paths and boot LUNS

In boot from SAN configurations, the boot paths must be on the active path to the boot LUN. Because of a failover event, storage controller reboot, or some other event in the SAN, it is possible for the currently configured boot path to become the standby path to the LUN. In this case, attempting to boot using the configured boot path fails.

To correct this condition, either configure a new boot option using the second path to the LUN, or switch the standby path to the boot LUN back to the active path.

Target port order

When the primary boot path fails and the host is rebooted, the host may fail to successfully find the boot LUN through the alternate path if the order of target ports (as seen by the HBAs) is not consistent throughout different SANs.

You must arrange the target ports on the SAN so that all SANs have the same target port order.

 **NOTE:** When rearranging target ports on the SAN, you must ensure that the connections will not be moved at a later time. This approach limits the flexibility of where targets can be connected on the SAN.

Proper configuration ensures that if the primary path fails during a host reboot, proper failover occurs, and the secondary path functions as the boot path as expected.

Image file error message

The following message may occur during installation on the BFS host:

```
Loading.....  
Can't open image file for initrd  
C:\>
```

This usually occurs due to a bad or corrupted LUN on your storage array. To fix this problem, rebuild your LUN, and then proceed with the installation.

Configuring the boot loader during operating system installation

For Red Hat EL3 configurations, you may need to configure the boot loader during the operating system installation. The following example illustrates this procedure.

During the operating system installation, the installer detects all the attached devices.

1. If an error messages about `/dev/sda`, appears, ignore the error message, press **CANCEL**, and continue.
2. Follow the directions on the basic installation screens.
3. Select the Manual partition scheme to ensure that the partitioning is correct.
4. At the Boot Loader screen do the following:
 - a. Confirm that the operating system will be installed on the proper `/dev/sd` device of the boot LUN.
 - b. Check the Configure Advanced boot loader options box.
 - c. On the next screen, select **Change drive order** and move the LUN you are using to boot to the top.
 - d. Accept the changes to close the pop-up window.
5. Complete the operating system installation procedure.