

HP StorageWorks

Using the QLogic HBA driver for
single-path or multipath failover mode
on Linux systems

application notes

Legal and notice information

© Copyright 1998, 2006 Hewlett-Packard Development Company, L.P.

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.

Red Hat® is a registered trademark of Red Hat, Inc.

Linux® is a registered trademark of Linus Torvalds.

About this document

This section describes the following topics:

- [Application notes information](#), page 3
- [Intended audience](#), page 3
- [Accessing future product updates](#), page 3
- [Other documentation](#), page 3

Application notes information

These application notes describe the following topics:

- [New features](#), page 3
- [Supported hardware and software](#), page 3
- [Installing the QLogic HBA driver](#), page 4
- [Removing the QLogic driver](#), page 5
- [Changing driver parameters](#), page 6
- [Understanding QLogic failover](#), page 7
- [Using QLogic failover with different storage system types](#), page 8
- [Using the QLogic HBA driver on a system with other multipathing products](#), page 8
- [Active/active EVA and autorestore](#), page 9
- [Troubleshooting HSV210 devices on Linux](#), page 10

Intended audience

This document is intended for customers who are using the QLogic host bus adapter (HBA) driver in failover mode on a system running Linux® operating software.

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.

Other documentation

Additional documentation, including white papers and best-practices documents, is available on the HP web site: <http://www.hp.com>.

New features

This release provides support for path preference.

Supported hardware and software

Table 1 lists supported hardware and software.

**NOTE:**

Depending on the storage array, existing Secure Path customers can transfer the path failover to the QLogic HBA driver. For more information, see the *HP StorageWorks LUN migration and persistence utilities application notes*.

Table 1 Supported hardware and software

Component	Description
HBAs	FCA2214DC — PN 321835-B21 (x86 only)
	A7538A
	A6826A (ia64 only)
	BL20P FC Mezzanine card — PN 300874-B21
	BL30P FC Balcony card — PN 354054-B21
	BL25P Mezzanine card — PN 381881-B21
	BL35P Balcony card
	FC1142SR PCI-X — PN AE311A
	FC1242SRDC PCI-X — PN AE312A
Storage systems	EVA3000/5000 EVA4000/6000/8000 MSA1000/1500 XP128/1024/12000
Linux 7.xx	Red Hat® Advanced Server 2.1, i86, ia64 SuSE Linux Enterprise Server 8 (Service Pack 3 or later) i386, ia64, x86_64 Red Hat® Linux 3.0, i386, ia64, x86_64
Linux 8.xx	SuSE Linux Enterprise Server 9, x86, x86_64, ia64 Red Hat Linux 4 x86, x86_64, ia64

Driver versions

If an EVA4000/6000/8000 storage system resides in a storage area network (SAN), the minimum driver version is 7.05.00a for 2.4 kernels and 8.00.02a for 2.6 kernels.

For an EVA3000/5000 active-active storage system, the minimum driver version is 7.07.03 for 2.4 kernels and 8.01.03 for 2.6 kernels.

Installing the QLogic HBA driver

You can use an installation script to install the QLogic HBA driver (see "[Using the scripted installation](#)" on page 5), or you can install the driver manually (see "[Using the manual installation](#)" on page 5).

Download the new driver from the HP storage web site:

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

Setting up driver parameters

Before you install the driver, set up the driver parameters using the `set_parm.2x` script for your kernel type:

- For 2.4 kernels, use `set_parm.24`.
- For 2.6 kernels, use `set_parm.26`.

See [Table 2](#) on page 7 for a list of parameters.

Using the scripted installation

NOTE:

You must install the GNU Compiler Collection (GCC) and the kernel source files to use the scripted installation.

The `INSTALL` script installs the following Red Hat Package Manager (RPM) files:

- QLA driver `hp_qla2x00src` RPM
- Fibre Channel utilities (`fibretils`) RPM for the `x86`, `x86_64`, or `ia64` architecture

You can use the `INSTALL` script for system configurations that meet any of the following conditions:

- No previous driver or platform kit was installed.
- The configuration was installed with an earlier version of `hp_qla2x00src` RPM.
- The configuration was installed with a platform kit that includes the `qla2x00` binary RPM.

If you use the `INSTALL` command without specifying options, it installs the driver in either failover mode, or non-failover mode, depending on the mode of the currently installed driver.

The following example shows the `INSTALL` command for single-path mode:

```
# ./INSTALL
```

To install the driver in multipath failover mode, use the `INSTALL` command with the `-f` option. For example:

```
# ./INSTALL -f
```

See [Removing the QLogic driver](#) to replace the drivers.

Using the manual installation

You can customize the installation by using the `RPM` command to select individual RPM packages:

- To install the driver package and the `fibretils` package, enter the following commands:

```
# rpm -ivh hp_qla2x00src-version-revision.noarch.rpm  
# rpm -ivh fibretils-version-revision.architecture.rpm
```
- To install the `hp_qla2x00src` RPM without running the driver build scripts, enter the following command:

```
# rpm -ivh -noscripts hp_qla2x00src-version-revision.noarch.rpm
```
- To manually upgrade the RPMs, enter the following commands:

```
# rpm -uvh hp_qla2x00src-version-revision.noarch.rpm  
# rpm -uvh fibretils-version-revision.architecture.rpm
```

Removing the QLogic HBA driver

You can use an installation script to remove the QLogic HBA driver, or you can remove components manually.

To remove the driver and Fibre Channel utilities, use the `INSTALL` script with the `-u` option:

```
# ./INSTALL -u
```

You can remove one, two, or all components manually by entering the appropriate commands:

- To remove the `fibreutils` package, enter the following command:
`rpm -e fibreutils`
- To remove the `hp_qla2x00` package, enter the following command:
`rpm -e hp_qla2x00`
- To remove all components, enter the following command:
`rpm -e hp_qla2x00src`

Changing driver parameters

HP sets the following parameter values in the nonvolatile RAM (NVRAM) of the QLogic adapters:

- Port down retry count (`qlport_down_retry`)
- Login retry count (`qlogin_retry_count`)
- Queue depth (`ql2xmaxqdepth`)
- Failover (`ql2xfailover`)
- Load balancing type (`ql2xlbType`)
- Exclude model (`ql2xexcludemodel`)
- Auto restore (`ql2xautorestore`)

Port down retry count

The port down retry count (`qlport_down_retry`) parameter sets the number of retries that the driver attempts before declaring the port down. Depending upon your configuration, ensure that the values are set as follows:

- Single Path: 64
- Secure Path: 1
- For QLogic failover mode, the value depends on the kernel:
 - For 7.x: 3
 - For 8.x: 30

Login retry count

The login retry count (`qlogin_retry_count`) parameter sets the number of logins the driver attempts before declaring the port down. Ensure that this parameter is set to one of the following values:

- For 7.x: 16
- For 8.x: 30

Queue depth

The queue depth parameter (`ql2xmaxqdepth`) sets the length of the command queue for each logical unit number (LUN). Valid values are from 1 to 256; however, setting the value too high can cause storage saturation problems. HP recommends that you set this parameter to 16.

Failover

The failover parameter (`ql2xfailover`) turns the QLogic driver failover on or off. A value of 0 disables failover; and a value of 1 enables failover.

Enabling failover automatically sets the following parameters:

- Load balancing type (`ql2x1bType`)—The driver attempts to distribute LUNs evenly across known active paths.
- Exclude model (`ql2xexcludemodel`)—The driver exports certain devices as non-failover devices. This parameter is used only if Secure Path is installed.
- Auto Restore (`ql2xautorestore`)—The QLogic failover driver attempts to failback active I/Os to the preferred path.

The `/etc/hp_qla2x00.conf` file contains the failover parameter values. HP recommends that you do not edit this file directly; the preferred method is to use the `set_parm` script.

Using the `set_parm` script

You can change the driver parameter values by using the `set_parm` script in either the `/opt/hp/src/hp_qla2x00src` directory or the `/opt/hp/hp_qla2x00` directory, depending which RPM you have installed. [Table 2](#) describes the driver parameter values.

You can also set the driver parameters manually using the `set_parm` command with the `manual` option.



NOTE:

If you change parameter values manually in the file, you must build a new `initrd` file by running the `edit_conf` and `make_initrd` scripts before the new values take effect.

Table 2 Driver parameters

Parameter	Switch	Description
Single Path	-s	Connect through one path to the storage without any redundancy.
Secure Path	-m	Use the Secure Path software for multipathing and failover.
QLogic Failover	-x	Use the multipathing and failover functionality built into the <code>qla2x00</code> RPM.

Understanding QLogic failover

The section describes the behavior of the QLogic failover driver.

Failover

The QLogic failover driver enables I/O to move from a failed path to a secondary path in a configuration with multiple I/O paths to supported HP StorageWorks arrays.

Failback

If the secondary path fails after a failover, the QLogic failover driver enables I/O to move to another path, which can be the original path if its functionality has been restored.



NOTE:

The secondary path must fail for failback to occur. I/O does not resume on the primary path automatically when it is restored unless the auto restore feature is enabled.

Preferred path

You can specify the path to a physical device that I/O uses. I/O uses the preferred path unless that path fails.

NOTE:

Preferred paths are supported on the XP arrays, EVA4000/6000/8000 and EVA3000/5000 AA. To set the preferred path, use the QLogic SANsurfer software.

Load balancing

Load balancing keeps all active paths equally busy by determining which path I/O takes. There are two types of load balancing:

- Static load balancing—You specify manually which I/O path to use.
- Dynamic load balancing—An algorithm determines which I/O path to use.

NOTE:

To set the dynamic load balancing I/O, use QLogic SANsurfer software.

Using QLogic failover with different storage system types

You must follow specific storage system procedures when using QLogic failover. The procedure that you use depends on your storage system type.

Using QLogic failover with EVA3000/5000 storage systems

If you are using the built-in failover functionality of the QLogic driver, set the failover preference to **Path A Failover** or **Path B Failover** in the Management Appliance host settings. Failure to do so can cause I/O errors when a failover occurs.

NOTE:

Set the host mode to **Sun Solaris**. Failure to do so can result in discovery errors.

Using QLogic failover with MSA1000/1500 storage systems

If you are using the built-in failover functionality of the QLogic driver, set the host mode to **Linux**. Failure to do so can cause I/O errors when a failover occurs.

HSG80 storage systems

QLogic failover is not supported on HSG80 storage systems.

Using the QLogic HBA driver on a system with other multipathing products

When you add new EVA4000/6000/8000 LUNs to a host that also presents LUNs to a legacy storage system managed by Secure Path, you must update the QLogic driver and utilities.

To update the QLogic driver and utilities that support EVA4000/6000/8000:

1. Enter the `/opt/hp/src/hp_qla2x00src/set_parm` command to set the driver parameters.
2. Select option 2, **Secure Path**. Enter **no**, when prompted to create a new `initrd` file.
3. Enter **1** for option 4, Failover.
4. Export models to 6. Enter **Yes**, when prompted to create a new `initrd` file.
5. Select option 5, **Quit**.
6. Enter the `/opt/hp/hp_fibreutils/pbl/pbl_inst.sh -i` command for the 2.4 kernels SCSI blacklist that does not include the HSV210 arrays.
This step is required for EVA XL LUN discovery to occur at system boot.
7. Reboot the system, and then verify that Secure Path is working.
8. Add your EVA4000/6000/8000 storage system to the SAN.

 **NOTE:**

To ensure that the LUNs are listed in sequential order, list your EVA4000/6000/8000 storage system before the lower switch port numbers of the legacy storage system managed by Secure Path.

9. Reboot the system, and then verify that the QLogic and Secure Path LUNs are present and accessible.

 **NOTE:**

Running `hp_rescan` and `probe-luns` script is not supported on systems with Secure Path. To discover newly added devices, reboot the system.

Active/active EVA and autorestore

Autorestore occurs when the I/O path is restored to its original, preferred path. Usually, if the preferred path is disrupted, the I/O moves to an alternate path. The I/O stays on this path until that alternate path is disrupted. With autorestore enabled, I/O returns to the original preferred path when that original path is restored to a normal working state.

With the QLogic failover driver, autorestore has to be enabled on each device (array) type.

Enabling autorestore for EVA

To enable autorestore for EVA active/active devices:

1. Open the `/etc/modprobe.conf` file for 2.6 kernel systems or `/etc/modules.conf` file for 2.4 kernel systems.
2. Locate the following lines that begins with:
 - For 2.6 kernel systems, the line begins with `options qla2xxx`
 - For 2.4 kernel systems, the line begins with `options qla2300`
3. Change the parameter `ql2xautorestore=0xXX` to `ql2xautorestore=0x20`.
The `XX` is the original value.
4. Save and close the file.
5. Run the `/opt/hp/src/hp_qla2x00src/make_initrd` script to implement the driver parameters changes and to persist the changes through reboots.
6. Reload the QLogic driver modules or reboot.

Autorestore and other devices

If you enabled autorestore for other arrays, you must add their values with the value for EVA active/active devices.

For example, if MSA active/active autorestore is enabled, the value for `ql2xautorestore` is 0x80. To enable autorestore for both MSA active/active and EVA active/active devices, add the two parameters together as follows:

`0x20 + 0x80 = 0xA0`

Therefore, the parameter value is now: `autorestore=0xa0`.

Troubleshooting HSV210 devices on Linux

In some of the earlier updates for Linux, the blacklist or HSV210 device list that is required by the SCSI midlayer of the Linux kernel does not exist. Blacklist or device list can prevent Linux from detecting listed devices during system boot. If you have trouble detecting these devices, run the `/opt/hp/hp_fibreutils/probe-luns` script during system boot.

- Enter the following commands to install the `probe-luns` startup code:

```
# cd /opt/hp/hp_fibreutils/pbl
# ./pbl_inst.sh -i
```
- Enter the following commands to remove the `probe-luns` startup code:

```
# cd /opt/hp/hp_fibreutils/pbl
# ./pbl_inst.sh -u
```

 **NOTE:**

If you are able to detect the devices, do not run `probe-luns` at system boot.
