

Fabric OS[®] Software

Brocade[®] Fabric Notifications – Installation Procedure for Linux

Installation Guide

Copyright © 2022 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to www.broadcom.com. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Table of Contents

Chapter 1: Description	4
Chapter 2: Fibre Channel Requirements	5
2.1 SAN Hardware and Firmware	5
2.2 HBA Hardware and Firmware	5
2.3 Operating System Software	5
2.4 Multipath Solutions	6
Chapter 3: Installation	7
3.1 RHEL	7
Chapter 4: Supported Fabric Notifications	9
4.1 Link Integrity FPIN	9
4.1.1 FPIN-LI Actions	9
4.1.2 Marginal Path Recovery	9
4.2 Congestion FPINs	10
4.3 Delivery FPIN	10
Revision History	11
FOS-FNIPL-IG100; March 8, 2022	11

Chapter 1: Description

This installation guide describes the steps necessary to install the Brocade® Fibre Channel Transport functionality for Linux. This functionality enables the Linux server to participate in and react to Fabric Notifications generated by the attached Fibre Channel SAN.

The Brocade® Fabric OS® (FOS) version 9.0.0 introduces the Fabric Notifications functionality that provides the ability for end devices to receive and react to events in the network. Registered end devices receive these events through a Fibre Channel Extended Link Service (ELS) function known as a Fabric Performance Impact Notifications (FPIN). Participating devices are able to use these events to evaluate error conditions and apply reactive measures to mitigate the condition.

In this document, we describe the requirements and configuration steps necessary to install the Fibre Channel (FC) Transport daemon for Linux, which enables the Linux multipath I/O (MPIO) layer to react to the FPIN events. Specifically, the daemon improves the functioning of the Linux dm-multipath component by processing the Fabric Notification events (FPINs) forwarded from the Fibre Channel host bus adapter (HBA).

Ultimately, the addition of the FC Transport daemon functionality improves the Linux MPIO technology by enabling reliable and deterministic response to transient fabric events.

Chapter 2: Fibre Channel Requirements

This section describes the required components necessary to enable the Fabric Notifications functionality in the SAN hardware and firmware, in the HBA hardware, firmware and driver, and in the operating system software.

2.1 SAN Hardware and Firmware

The following table shows what SAN hardware and firmware are required to support the Fabric Notifications features.

Switch Generation	Firmware	References
Gen6	FOS 9.0.0a	www.broadcom.com/products/fibre-channel-networking
Gen7		

2.2 HBA Hardware and Firmware

The following table shows what HBA hardware and firmware are required to support the Fabric Notifications features.

HBA	Model	Firmware	Driver	References
Support for FPIN ELS				
ECD/Emulex®	LPe3100x	12.6.240.40	12.6.240.33-1	www.broadcom.com/products/storage/fibre-channel-host-bus-adapters
	LPe3200x			
Marvell/Qlogic	QLE269x	9.3.40	10.1.0.63.x	www.marvell.com/products/fibre-channel-adapters-and-controllers.html
	QLE274x			
Support for FPIN ELS and Congestion Signal				
ECD/Emulex	LPe3500x-M2	12.6.240.40	12.6.240.33-1	www.broadcom.com/products/storage/fibre-channel-host-bus-adapters
Marvell/Qlogic	QLE277x	9.4.0	10.1.0.63.x	www.marvell.com/products/fibre-channel-adapters-and-controllers.html
	QLE28xx			

2.3 Operating System Software

The following table shows what operating system software levels are required to support the Fabric Notifications features.

Distribution	Release	References
IBM	AIX 7.2TL5	<ul style="list-style-type: none"> ■ www.ibm.com/products/aix ■ www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/1/897/ENUS220-381/index.html&lang=en&request_locale=en
Redhat	RHEL 8.3 / EPEL 8	<ul style="list-style-type: none"> ■ access.redhat.com/downloads ■ docs.fedoraproject.org/en-US/epel/

2.4 Multipath Solutions

The following table shows what multipath software solutions support the Fabric Notifications features.

Vendor	Release	References
DellEMC	PowerPath 7.4	<ul style="list-style-type: none">■ www.dell.com/en-us/dt/storage/powerpath-intelligent-multipathing-software.htm■ dl.dell.com/content/manual51097957-dell-emc-powerpath-for-linux-7-4-and-minor-releases-installation-and-administration-guide.pdf?language=en-us

Chapter 3: Installation

The following sections describe the installation instructions for the various distributions of the Linux operating system.

3.1 RHEL

The current offering from Redhat is an *out-of-box* solution that requires the installation of the [Extra Packages for Enterprise Linux](#) (EPEL) component. The package containing the Fabric Notifications functionality is the `fctxpd` package (such as, `fctxpd-0.2-4.20210326gitc4dba7f.e18.x86_64`).

The following steps outline the process for adding the [EPEL](#) component to the base RHEL system and starting the FC Transport daemon:

1. Install the EPEL package:

```
dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
dnf update
dnf install fctxpd.x86_64
```

2. Edit the `multipath.conf` file:

a. Enable the attribute `marginal_pathgroups` in the `/etc/multipath.conf` file.

For example: `marginal_pathgroups 1`

NOTE: If the `marginal_pathgroups` attribute is disabled, the FC Transport daemon will take **no action** on a path associated with a Fabric Notification event.

3. Start the FC Transport daemon:

a. Run `systemctl status fctxpd`

The result of this step should display the following:

```
fctxpd.service - Fiber Channel Transport Daemon
Loaded: loaded (/usr/lib/systemd/system/fctxpd.service; disabled; vendor preset: disable)
Active: inactive (dead)
```

NOTE: After it is loaded, the FC Transport daemon does not start by default. The next steps are required to start the daemon.

b. Run `systemctl enable fctxpd`

This step insures the FC Transport daemon is automatically loaded when the server reboots.

c. Run `systemctl start fctxpd`

4. Verify the FC Transport daemon status:

a. Run `systemctl status fctxpd`

The result of this step should display the following:

```
systemctl status fctxpd
fctxpd.service - Fiber Channel Transport Daemon
Loaded: loaded (/usr/lib/systemd/system/fctxpd.service; disabled; vendor preset: disabled)
Active: active (running) since Mon 2020-07-27 17:43:11 IST; 2s ago
Main PID: 64691 (fctxpd)
Tasks: 2 (limit: 403798)
Memory: 688.0K
```

At this point, the installation procedure is complete and the system is ready to participate in the Fabric Notifications functionality.

Chapter 4: Supported Fabric Notifications

The FC Transport (FCT) daemon processes the FPIN ELS events sent by the Fibre Channel fabric and forwarded by the HBA.

4.1 Link Integrity FPIN

The Link Integrity notification FPIN ELS (FPIN-LI) is processed.

When the FC Transport daemon receives an FPIN ELS for a link integrity event (FPIN-LI), it processes the event for the SCSI block device associated with the path indicated in the FPIN ELS. The FC Transport daemon interprets the event and sets the SCSI block device path to the marginal path state. When this occurs, the Linux dm-multipath driver will not select the path for active IO until the link condition is resolved (such as, toggling the associated device port).

The FC Transport daemon manages marginal paths by moving the path to the marginal path group. Paths that are associated with the marginal path group are not used for I/O operations unless no usable paths are available (such as, there are no paths in the active path group). In this case, the kernel uses the paths in the marginal path group for I/O as a last resort. Paths are removed from the marginal path group with the link condition resolved (see the following [Marginal Path Recovery](#) section).

4.1.1 FPIN-LI Actions

When a link integrity event occurs, the fabric generates an FPIN-LI (such as, an FPIN ELS with the link Integrity descriptor) and distributes it to the registered devices. On a Linux server with the FC Transport daemon enabled, the HBA receives the event and passes the ELS frame to the FC Transport daemon (FCTXPTD). The event payload includes the port name (such as, Port World Wide Name) of the port detecting the event and the port name of the port attached to the detecting port.

When the FC Transport daemon receives the event from the HBA, it extracts the HBA port identifier and translates it to the port name. The port name is used to look up the path information associated with the event (such as, the port identifier, Port World Wide Name, symbolic link to the Target LUNs, and so on). The FC Transport daemon also extracts the list of affected port names from the FPIN-LI payload as well. This list determines which paths to designate as marginal with the dm-multipath component. A path designated as marginal is associated with a low priority group by the dm-multipath component and is used only when no other active paths are available.

4.1.2 Marginal Path Recovery

When the FC Transport daemon moves a path to the marginal path group, user intervention is required to recover the path using one of the following actions:

- Toggle the port – The user toggles the state of the affected port (such as, the port identified by the fabric as the source of and impacted by the link integrity event). Thus, any action that causes the port to go offline then online toggles the state of the port.
- Recover the path – The user executes the following DM commands to restore the path to the active state: `multipathd reinstate path <pathname>`

4.2 Congestion FPINs

The congestion notification and peer congestion notification FPIN ELSs (FPIN-CN and FPIN-PN) are not processed and do not affect the dm-multipath driver.

When The FC Transport daemon receives an FPIN ELS for a congestion notification, it logs the event and includes it in the Linux server syslog.

4.3 Delivery FPIN

The delivery notification FPIN ELS (FPIN-DN) is not processed.

When The FC Transport daemon receives an FPIN ELS for a delivery event, it discards the event.

Revision History

FOS-FNIPL-IG100; March 8, 2022

Initial release.

