



Hewlett Packard
Enterprise

HPE Plexxi HCN™
Release Notes
Switch and Control Software
Release 4.1.0

Part Number: P12880-001a
Published December 2018
Edition: 3



Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Export of the information contained in this publication may require authorization from the U.S. Department of Commerce.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgements

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java and Oracle are registered trademarks of Oracle and/or its affiliates. UNIX® is a registered trademark of The Open Group.

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Table of Contents

Notices	2
Acknowledgements.....	2
Introduction	4
Software Versions Included.....	4
Supported Hardware.....	4
New Features.....	4
Related Documentation.....	5
Contacting Plexxi Support	5
Fabric LAGs	6
Overview.....	6
Creating Fabric LAGs for All Fabric Links in a Fabric	6
Creating a Fabric LAG for a Specified Number of Fabric Ports	6
Verifying the Fabric LAG Method	6
10GbE Copper NICs	7
Plexxi px-setup Utility Flashes Switch LED.....	7
Plexxi px-setup Utility SNMP Support	7
Expanded Syslog Event Reporting.....	7
Upgrade Considerations for Fabric Management	7
Overview.....	7
Switching to the IS-IS Fabric Management Mechanism	8
Verifying the IS-IS Fabric Management Mechanism.....	8
Monitoring the IS-IS Fabric Management Mechanism.....	9
Plexxi CLI Command Syntax Changes	9
Changes to Plexxi CLI Commands for Data Center Interconnect.....	10
New Plexxi CLI Commands for Fabric Management.....	10
Deprecated Plexxi CLI Commands	11
Issues Fixed in this Release.....	12
Known Issues in this Release	13

Introduction

This document contains information about Plexxi Switch and Control Release 4.1.0, which delivers software updates for Plexxi Switch and Plexxi Control.

Software Versions Included

The following versions of software updates are included:

- Plexxi Switch 4.1.0
- Plexxi Control 4.1.0

Supported Hardware

Release 4.1.0 is supported on the following Plexxi switches:

- Switch 1x (PX-S1X)
- Switch 2 (PX-S2)
- Switch 2s (PX-S2S)
- Switch 2p (PX-S2P)
- Switch 2sp (PX-S2SP)
- Switch 2e (PX-S2E) also known as HPE FM 2072
- Switch 3eq (PX-S3EQ)

New Features

The following new features and software changes are included in this release:

- Fabric LAGs
- 10GbE Copper NICs
- Plexxi px-setup Utility Flashes Switch LED
- Plexxi px-setup Utility SNMP Support
- Expanded Syslog Event Reporting
- Upgrade Considerations for Fabric Management

Related Documentation

The following additional documentation supports this release:

- *Plexxi Compatibility Matrix, Version 12* or greater. The *Plexxi Compatibility Matrix* contains version-specific software and hardware support information as well as cable and transceiver support information.
- *Plexxi Switch Administrator Guide using Linux and Plexxi CLI, Release 4.1.0*
- *Plexxi Control Installation, Upgrade and Administration Guide, Release 4.1.0*
- *Plexxi Switch 3eq Hardware and Installation Guide*
- *Plexxi Getting Started Guide, Small and Medium Deployments, Releases: Switch-Control 3.2.2 – 4.1.0 and Connect 2.4.0 – 3.1.0*

Documentation (except Online Help) is available on the Support > Technical Documentation page of <http://www.plexxi.com>.

- *Plexxi Control Online Help* is available while logged into the Plexxi Control UI

Contacting Plexxi Support

Plexxi Technical Support services are available to answer your questions and to make sure that your software and hardware continue to operate properly.

You can contact Plexxi Support at:

support@plexxi.com

1.888.415.9809 (US/Canada toll-free)

+1 603-782-0702 (US/International)

<http://support.plexxi.com>

Fabric LAGs

Overview

You can create Link Aggregation Groups (LAGs) for the fabric links between Plexxi switches. This functionality will combine all parallel fabric links between two switches into a single, larger LAG. This feature enables you to provide higher bandwidth single paths for IS-IS and Fitting Engine calculated topologies and will also allow for a higher bandwidth topology for Broadcast, Unknown unicast, and Multicast (BUM) traffic. After creating fabric LAGs, you must perform a fit using the Plexxi Control UI.

You configure fabric LAGs by executing the `fabric lag` CLI command from the Privileged Exec mode.

IMPORTANT: It is assumed that all of the parallel fabric links between the two peer switches have the same speed. If they do not, creating a LAG for the fabric links between the peer switches will fail.

Creating Fabric LAGs for All Fabric Links in a Fabric

To create LAGs that group the fabric links between all Plexxi switches in a fabric, enter the following command on a single switch in the fabric:

```
fabric lag all
```

This command gets executed on all of the Plexxi switches in the fabric. Each switch in the fabric creates a LAG that groups the fabric links to each of its peers, creating one LAG per peer, even when there is a single fabric link between the switch and a peer. This command provides the highest available bandwidth topology for Broadcast, Unknown unicast, and Multicast (BUM) traffic.

IMPORTANT: Fabric LAGs reduce the overall available fabric diversity and paths in a Plexxi fabric. As a result, fewer options are available to create isolated topologies in the fabric. It is therefore recommended not to combine `fabric lag all` with Plexxi fabrics that require isolated topologies for storage or other workloads.

Creating a Fabric LAG for a Specified Number of Fabric Ports

To create a LAG for a specified number of the fabric links used to connect a switch to a peer, enter the following command on both of the two peer switches:

```
fabric lag links [ number | none ]
```

where *number* is the number of fabric links to be included in the LAG. The number of fabric links to be included in the LAG can be a subset of the total number of fabric links used to connect the two peer switches. The same number of fabric links must be specified on both of the two peer switches. The default value is *none*, where no fabric LAG is created.

Verifying the Fabric LAG Method

To show the configured fabric LAG method, in the Exec or Privileged Exec modes of the Plexxi CLI, enter the following command:

```
show fabric
```

The command output includes a line for the configured LAG method, which can be *none*, *all*, and *links number*.

10GbE Copper NICs

10GbE copper NICs are supported through the use of ProLabs 10GbE copper SFP+ transceivers (Plexxi product code: PX-CBL-PL-10GBT-SFP-30M).

Plexxi px-setup Utility Flashes Switch LED

When configuring a switch using the Plexxi `px-setup` utility, an LED now flashes on the switch being configured.

Plexxi px-setup Utility SNMP Support

The `px-setup` utility now allows the creation of SNMP v3 user accounts for switches.

Expanded Syslog Event Reporting

Logging of the following types of events is turned on to enable the events to be reported in syslog messages.

- BGP events
- PIM events
- Route events
- IP mroute/MFIB events
- IGMP groups/snooping events

Upgrade Considerations for Fabric Management

Overview

IMPORTANT: If you are upgrading to Plexxi Switch Release 4.1.0 from a release prior to Release 4.0.0, read this section and follow the instructions for switching to the new IS-IS fabric management mechanism.

As of Plexxi Switch Release 4.0.0 and later, the mechanism that Plexxi switches use to discover each other, create a default forwarding topology for unicast and multicast routing, and maintain fabric connectivity during switch or link failures has been significantly changed. Plexxi switches now use the standard Intermediate System to Intermediate System (IS-IS) protocol to discover connectivity between switches and maintain default unicast and multicast forwarding topologies. The forwarding topologies created by the IS-IS fabric management mechanism use shortest path rules for switch-to-switch control plane traffic and serve as the default forwarding topology in the absence of Plexxi Control calculated topologies. The IS-IS fabric management mechanism also uses shortest path rules to create a Broadcast, Unknown unicast and Multicast (BUM) topology that is used for all BUM traffic in a Plexxi fabric.

In cases of switch or link failures, the IS-IS fabric management mechanism will automatically converge onto new topologies.

The topologies created by the IS-IS fabric management mechanism do not replace the forwarding topologies calculated by Plexxi Control through fitting. Fitting-created topologies will always take precedence. Topologies calculated by the IS-IS fabric management mechanism should be seen as the default topology for switch-to-switch control traffic, and as a topology of last resort in the case where no valid fitted topologies are available.

Switching to the IS-IS Fabric Management Mechanism

The IS-IS fabric management mechanism is not compatible with the fabric management performed in releases prior to Release 4.0.0/4.1.0. This means that a Plexxi fabric needs to run entirely in the prior fabric management mechanism or run entirely in the new IS-IS fabric management mechanism.

Fabrics that are upgraded to Release 4.1.0 from a release prior to Release 4.0.0 will continue to use the old fabric management mechanism to ensure full operation during the upgrade. After the upgrade to Release 4.1.0 has been completed and the fabric is stable, these fabrics should be switched to the IS-IS fabric management mechanism as soon as possible. Switching to the IS-IS fabric management mechanism is done across the fabric with no reboot required, but some traffic might be affected for a short period of time.

For Plexxi Switch Release 4.1.0 upgrade instructions, including instructions for switching to the IS-IS fabric management mechanism, refer to Chapter 8, Upgrading the Plexxi Switch Software to Release 4.1.0, in the *Plexxi Switch Administrator Guide using Linux and Plexxi CLI, Release 4.1.0*.

IMPORTANT: After upgrading a Plexxi switch to Release 4.1.0 and switching to the new IS-IS fabric management mechanism, Plexxi recommends that you use the Plexxi Control UI to configure and perform a new fit. Refer to the Plexxi Control Online Help for instructions.

New switches that have not been upgraded from a release prior to Release 4.0.0 will automatically run the IS-IS fabric management mechanism. This means that new Plexxi switches running Release 4.0.0 or later can only be added to a fabric after switching to the IS-IS fabric management mechanism.

Verifying the IS-IS Fabric Management Mechanism

After performing the upgrade to Release 4.1.0 and switching to the IS-IS fabric management mechanism, verify that the Plexxi switches in the fabric are using the IS-IS protocol for fabric management by issuing the following `show` commands:

```
show fabric
show topography
```

Output from these commands will display L2 IS-IS as the protocol used by the fabric. For example:

```
Plexxil> show fabric
Ring redirect setting: west-east
Confluent LAG link speeds (West-East): 10G-10G
Control channel protocol in use: L2-ISIS
Each link in confluent fabric has a speed of 10 Gbps
```

```
Plexxil> show topography
Topography Info
-----
Protocol In Use: L2-ISIS
Local station MAC address: xx:xx:xx:xx:xx:xx
Local alias: s2329
Operational stage: S2i

Fabric is stable.
Fabric UUID:
Controller: 0.0.0.0
Active List:
  xx:xx:xx:xx:xx:xx s2329 (S2i)
  R xx:xx:xx:xx:xx:xx s2332 (S2i)
BUM Ports: xp74
```

Monitoring the IS-IS Fabric Management Mechanism

In the Plexxi CLI, the status of the IS-IS fabric management mechanism can be examined using the following `show` commands:

Plexxi Switch CLI Command	Definition
<code>show l2-isis bum-graph</code>	Displays the L2 IS-IS BUM graph.
<code>show l2-isis counter</code>	Displays L2 IS-IS interface counters.
<code>show l2-isis database</code>	Displays the L2 IS-IS link state database.
<code>show l2-isis interface</code>	Displays L2 IS-IS interface information.
<code>show l2-isis neighbors</code>	Displays L2 IS-IS neighbor adjacencies.
<code>show l2-isis protocol</code>	Displays L2 IS-IS routing protocol process information.
<code>show l2-isis route</code>	Displays the L2 IS-IS routing table.
<code>show l2-isis isis topology</code>	Displays the L2 paths to the Intermediate Systems.
<code>show l2-isis summary</code>	Displays the L2 summary of port and packet counts.

Plexxi CLI Command Syntax Changes

After upgrading to Plexxi Switch Release 4.1.0 from a release prior to Release 4.0.0, the syntax of the following CLI commands has changed:

CLI Command Syntax before Release 4.0.0	CLI Command Syntax in Release 4.0.0/4.1.0
<code>show ring</code>	<code>show fabric</code>
<code>ring east-egress-rate (Switch 2e)</code>	<code>fabric east-egress-rate (Switch 2e)</code>
<code>ring east-speed (Switch 2e)</code>	<code>fabric east-speed (Switch 2e)</code>
<code>ring egress-rate (Switch 3eq)</code>	<code>fabric egress-rate (Switch 3eq)</code>
<code>ring lightrails (Switch 3eq)</code>	<code>fabric lightrails (Switch 3eq)</code>
<code>ring redirect (Switch 2e)</code>	<code>fabric redirect (Switch 2e)</code>
<code>ring redirect (Switch 2, 2s, 2p, 2sp)</code>	<code>fabric redirect (Switch 2, 2s, 2p, 2sp)</code>
<code>ring speed (Switch 3eq)</code>	<code>fabric speed (Switch 3eq)</code>
<code>ring west-egress-rate (Switch 2e)</code>	<code>fabric west-egress-rate (Switch 2e)</code>
<code>ring west-speed (Switch 2e)</code>	<code>fabric west-speed (Switch 2e)</code>

Changes to Plexxi CLI Commands for Data Center Interconnect

After upgrading to Plexxi Switch Release 4.1.0, on switches previously configured for Data Center Interconnect (DCI) mode using the `fabric redirect` command, the definition of “west” and “east” affected by the CLI command has changed and might decrease the number of “west” and “east” ports on a switch, as follows:

- On Plexxi Switches 2, 2s, 2p, and 2sp, the west ports are now mapped to the single SFP+ port xp65, and the east ports are now mapped to the single SFP+ port xp66.
- On Plexxi Switch 2e, the west ports are now mapped to the single SFP+ port xp25, and the east ports are now mapped to the single SFP+ port xp26.
- On Plexxi Switch 3eq, the west ports are now mapped to a single channel on the QSFP port xp105, and the east ports are now mapped to a single channel on the QSFP port xp106 (applies only to the `fabric egress-rate` and `fabric west-egress-rate` commands).

New Plexxi CLI Commands for Fabric Management

In addition to the syntax changes for the `ring` commands listed above, the following new CLI commands have been added in Release 4.0.0 and later for fabric management:

- `fabric protocol-change l2isis`—For fabrics that have been upgraded to Plexxi Switch Release 4.1.0 from a release prior to Release 4.0.0, this command changes the control protocol for the fabric to the new IS-IS protocol. Fabrics that are upgraded to Release 4.1.0 from a release prior to Release 4.0.0 will continue to use the old fabric management mechanism to ensure full operation during an upgrade. After the upgrade to Release 4.1.0 has been completed and the fabric is stable, these fabrics need to be configured to use the IS-IS based fabric management mechanism.

Issue this command from the `px-shell` on any Plexxi switch in the fabric and the command will change all switches in the fabric to the new IS-IS based fabric management mechanism.

IMPORTANT: This change can potentially create a short interruption in seconds in traffic forwarding.

- `fabric clear-fabric-id`—When connecting Plexxi switches that were previously in a fabric to a new fabric, these switches do not automatically join the new fabric until you issue this command. This behavior is to ensure that switches and fabrics are not connected to each other by mistake, potentially causing connectivity disruptions.

To join the new fabric, the switches must have no fabric identifier, just as a new switch has no fabric identifier. This command clears the fabric identifier from the switches that were in a previous fabric and are about to be connected to the new fabric. After the fabric identifier has been cleared, the switches that were already in the fabric will assign their shared fabric identifier to the newly-connected switches, and the newly-connected switches join the new fabric.

- `fabric learn-fabric-id`—When you break up a fabric into multiple smaller fabrics, the Plexxi switches in the smaller fabrics retain the fabric identifier assigned by the original fabric. Because Plexxi Control uses the fabric identifier to identify the switches in a fabric, it will view the switches now part of the smaller fabrics as part of a single segmented fabric.

This command creates and assigns a new fabric identifier to all the switches in a fabric. Issue this command on any switch in each of the smaller fabrics to enable Plexxi Control to recognize the smaller fabric as an entirely new and separate fabric.

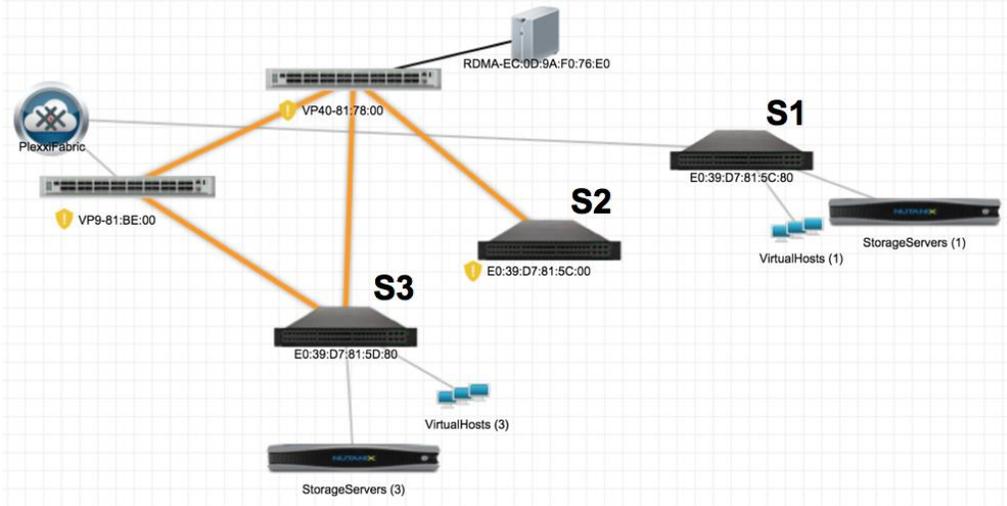
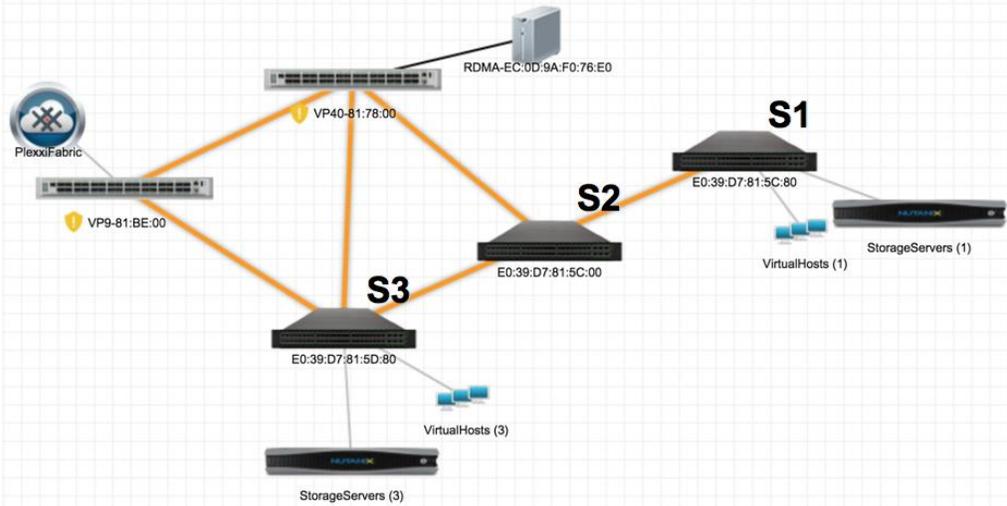
Note: If Plexxi Control has never detected and controlled these switches before, issuing this command is the same as issuing the `fabric clear-fabric-id` command.

Deprecated Plexxi CLI Commands

The following CLI commands are deprecated in Plexxi Switch Release 4.0.0 and later:

- `dci-behavior`—This command enables or disables east and/or west DCI behavior. When running the IS-IS fabric management mechanism, no port ever gets blocked if there is a temporary condition that prevents communication without taking down a link, so this command is no longer needed.
- `fabric confluent-fabric-links` (LightRail1 and LightRail2)—This command configures the number of links (1-12) to be used in the confluent fabric LAG. The IS-IS fabric management mechanism uses bypass links for Broadcast, Unknown unicast and Multicast (BUM) traffic rather than the confluent fabric links, so this command is no longer needed.
- `fabric control`—This command, which was available for Switches 2, 2s, 2p, and 2sp, configures the control plane traffic channel (in-band or out-of-band). This command is no longer needed when running the IS-IS fabric management mechanism because the control plane traffic is always in-band.

Issues Fixed in this Release

Issue ID	Description
PXCTL-3922	<p>In a leaf-spine configuration, the Plexxi Control UI's Networks view did not show connections between all of the peer switches.</p> <p>For example, the connections between peer switches S1, S2, and S3 below would not appear:</p>  <p>This issue has been fixed in this release, as show in the example below:</p>  <p>For both leaf-spine and ring configurations, Networks view in this release shows peer connections between all Plexxi switches rather than physical cable connections.</p>

Known Issues in this Release

Issue ID	Description
	<p>Issue when upgrading switch software on any Plexxi switch to Switch 4.1.0:</p> <p>Before upgrading a switch from a pre-3.3.0 release, if you have Link Aggregation Groups with member ports on different switches (MLAGs), and you are running STP/Loopguard in your environment connected to Plexxi, please contact Plexxi Support for a pre-upgrade script.</p>
	<p>Multicast traffic over an MLAG is not supported.</p>
PXCTL-3448	<p>L2VPN configuration tasks should not be executed at a rate higher than 1 every 30 seconds.</p>
PXSW-2501	<p>The following issue exists:</p> <p>When a switch that is on the In-Band Management LAG is rebooted, the LACP configuration is not restored on the In-Band Management LAG.</p>
PXSW-2443	<p>The following MC traffic issue exists:</p> <p>When an MC receiver host moves, MC traffic still gets delivered to the port it was attached to previously until the IGMP JOIN on that port times out.</p>
PXSW-2440	<p>The following MC traffic issue exists when increasing the confluent ring size:</p> <p>An access port that was part of an MC group prior to increasing the confluent ring size is removed from the group after the confluent ring size is increased.</p>
PXSW-2165	<p>An issue exists where:</p> <p>If the following User Defined Path conditions exist:</p> <ul style="list-style-type: none"> • the path uses confluent ring ports • the last switch in the path (egress switch) is the Ubervisor switch <p>And if a Flow Entry has the following characteristics:</p> <ul style="list-style-type: none"> • Path action is selected • The user defined path stated above is selected • The egress ports on the Ubervisor egress switch (last switch in path) are not selected <p>Broadcast, unknown unicast, and unknown multicast traffic will not be forwarded. Known unicast and known multicast traffic is forwarded as expected.</p> <p>To solve this issue, if the Ubervisor switch is the egress switch for a User Defined Path that uses the confluent ring, in Plexxi Control, when configuring the Flow Entry, explicitly select the egress ports on the egress (Ubervisor) switch.</p> <p>If the last switch in the path is not an Ubervisor switch, traffic flows as expected.</p>

Issue ID	Description
PXSW-2162	In a Flow Entry created in Plexxi Control that uses the Path action, the Flow Entry might match a flow that does not originate on the source switch defined in the user-defined path.
PXSW-1216	Before increasing the size of the confluent ring using the Plexxi Switch CLI, you must remove any configuration from the ports to be added to the confluent ring LAG. This configuration includes L1 Replication, Fabric Link Encapsulation and User Defined Paths.