

Catalyst 6500 Series Switches PFC, DFC, and CFC FAQ

Document ID: 107258

Questions

Introduction

What happens when you mix different versions of the PFC3x and DFC3x?

How can I determine the current PFC mode?

Can the Supervisors with different PFC versions form redundancy?

What are the benefits of a DFC?

For what is the CFC used?

What happens if line cards with different PFC modes are inserted in a Virtual Switching System (VSS) with PFC3C?

If the WS-X67xx line cards are required in Virtual Switching System (VSS), is the DFC3C or DFC3CXL also required, or could the default CFC be used?

Are the ingress and egress interface queues supported on the WS-X6704-10GE when used with a DFC3CXL in 12.2(33)SRB4? Also are these queues specific to the module/port, DFC type, or both?

Related Information

Introduction

This document addresses the frequently asked questions on the Policy Feature Card (PFC), Distributed Forwarding Card (DFC), and Centralized Forwarding Card (CFC) of the Cisco Catalyst 6500 Series switches.

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Q. What happens when you mix different versions of the PFC3x and DFC3x?

A. The PFC3 is the ASIC-based forwarding engine daughtercard for the Sup720; the DFC3 is the ASIC-based forwarding engine daughtercard for various fabric-enabled linecards (CEF256, CEF720). The PFC3/DFC3 generation is built upon a forwarding architecture known as EARL7. Within this generation, there are three different versions 'A', 'B', and 'BXL' that are all based on the same fundamental technologies but that each have incremental functionality. 'A' is the standard offering; 'B' is the intermediate option, and 'BXL' is the high-end option.

Since these versions are all within the same generation, it is possible that these versions can coexist within the same chassis with similar operation, but there are functional differences between them. A system with a mixture of forwarding engines only operates with the capabilities of the least-capable forwarding engine in the chassis. We cannot allow each forwarding engine to operate independently in its own mode. The BXL to run in BXL mode, the B to run in B mode, etc. within the same chassis is not allowed. This is because forwarding tables cannot be synchronized if we allow each PFC3/DFC3 to have different capabilities within the same system. For example, if the FIB table grows to 500K entries, the system is not able to provide consistent operation if it downloaded that 500K FIB table to the

PFC3BXL/DFC3BXL modules but not to the PFC3A/DFC3A and PFC3B/DFC3B modules. The same situation applies for an ACL configuration that requires more than 512 ACL labels. It is for these reasons that a mixture of PFC3/DFC3 versions must operate in a least–common–denominator mode so that tables can remain synchronized, and features can be applied consistently across interfaces.

PFC Card	DFC Card	Resultant Mode	Comments
PFC3A	DFC3A	PFC3A mode	No restrictions
	DFC3B		The PFC3A restricts DFC3B functionality
	DFC3BXL		The PFC3A restricts DFC3BXL functionality
	DFC3C		The PFC3A restricts DFC3C functionality
	DFC3CXL		The PFC3A restricts DFC3CXL functionality
PFC3B	DFC3A	PFC3A mode	PFC3B functionality is restricted by the DFC3A
	DFC3B	PFC3B mode	No restrictions
	DFC3BXL		The PFC3B restricts DFC3BXL functionality
	DFC3C		The PFC3B restricts DFC3C functionality
	DFC3CXL		The PFC3B restricts DFC3CXL functionality
	DFC3A	PFC3A mode	PFC3BXL functionality is restricted by the DFC3A
	DFC3B	PFC3B mode	PFC3BXL functionality is

			restricted by the DFC3B
	DFC3C	PFC3B mode	Each restricts the functionality of the other: the PFC3BXL functions as a PFC3B, and the DFC3C functions as a DFC3B (PFC3B mode).
	DFC3BXL		No restrictions
	DFC3CXL	PFC3BXL mode	The PFC3BXL restricts DFC3CXL functionality

Mixing DFCs and PFCs also impacts on-line insertion and removal (OIR). If a line card with DFC3B is inserted into a switch with SUP720-3BXL, the line card does not power up.

In order to use DFC3A-equipped switching modules with a PFC3BXL or PFC3B, the DFC3A-equipped switching modules must be installed at boot up. In order to use DFC3B-equipped switching modules with a PFC3BXL, the DFC3B-equipped switching modules must be installed at boot up.

The system also displays a similar error message:

```
OIR-6-DOWNGRADE_EARL: Module [dec] DFC installed is not identical
to system PFC and will perform at current system operating mode.
```

This is an example of the error message:

```
%OIR-SP-6-DOWNGRADE_EARL: Module 4 DFC installed is not identical
to system PFC and will perform at current system operating mode.
```

Q. How can I determine the current PFC mode?

A. With Release 12.2(17d) SXB and later releases, enter the **show platform hardware pfc mode** command to display the PFC3 mode.

With Release 12.2(17b) SXA and Release 12.2(17b)SXA2, enter the **show platform earl-mode** command to display the PFC3 mode.

Q. Can the Supervisors with different PFC versions form redundancy?

A. You cannot use one type of PFC3 (PFC3BXL, PFC3B, or PFC3A) on one supervisor engine and a different type on the other supervisor engine for redundancy. You must use identical policy feature cards for redundancy.

Q. What are the benefits of a DFC?

1. Performance is the biggest and most obvious reason to implement DFCs. You move from a 30 Mpps centralized forwarding system anywhere up to a 400 Mpps distributed forwarding system. This forwarding performance is for all L2 bridging, L3 routing, ACLs, QoS, and Netflow features, i.e., not just L3.
2. The performance benefit of a DFC is most applicable when you use the 67xx series modules. This is because these modules have enough ports and bandwidth to generate much more than the 30Mpps centralized forwarding engine has available. A 67xx-series module without a DFC is subject to the same centralized performance characteristics of all other centralized forwarding modules, i.e., 30 Mpps maximum for the whole system.

Consider a 6704 module; the 4x 10G ports have the ability to receive up to 60 Mpps of traffic (4x 14.88 Mpps, which is the linerate for 10GE at 64B frames). This is 2 times the amount of traffic that the centralized forwarding engine can handle, and this only account for one slot in the system. More modules in the system can also contribute to this oversubscription.

The addition of a DFC3 to the 6704 module increases the forwarding performance of that module to 48 Mpps; this is just for that slot, so the forwarding performance of the DFC3 is dedicated to the module on which it resides, i.e., not shares. The overall system over-subscription rate is greatly reduced when you add DFCs. The same principles applies to the 65xx modules, which also have optional DFC support, albeit at lower performance levels.

3. Minimize the impact that a classic module has in a system. Classic modules do affect the centralized forwarding performance of a system, limiting the maximum centralized forwarding rate to 15Mpps. Modules enabled with DFCs have their own forwarding engine and are not subject to this performance degradation. If a classic module used, the inclusion of a DFC mitigates any performance issues/concerns. Any non-DFC modules are still subject to the available 15 Mpps of forwarding available when a classic-module is present.
4. Increase the number of Netflow entries in the system. The system learns Netflow entries on a per DFC/PFC-basis; there is no synchronization between the Netflow tables. If we have 256K Netflow entries on a PFC3BXL/DFC3BXL, we can scale the system to 256K multiplied by the number of PFC3BXL/DFC3BXLs.
5. Increase the number of port-based QoS aggregate policers. A single PFC/DFC can support 1023 aggregate policers. We only download the policy for a port-based agg-policer to the PFC/DFC that manages the port in question. With 'x' number of PFC/DFCs, we can support 1023 times 'x' number of policers.
6. Increase the number of UBRL/Microflow policer entries. Since Microflow policer entries are stored in the Netflow table, the maximum capacity of these is directly related to the aggregate capacity of the Netflow table.
7. The addition of a DFC module effectively disconnects a module from the Data Bus. As such, a DFC-enabled module is not subject to the bus stall mechanism that occurs when a module is inserted or removed from the chassis. Throughout these Online Insertion and Removal (OIR) events, the Data Bus is temporarily paused for just enough time to ensure that the insertion/removal process does not cause any data corruption on the backplane. This protection mechanism causes a very brief amount of packet loss (sub-second, but dependent on the time it takes to fully insert a module). A module with a DFC onboard is not directly affected by this stall mechanism and does not have any packet loss on OIR.

Q. For what is the CFC used?

A. The WS-F6700-CFC is a daughtercard that provides centralized forwarding for the 67xx linecards. The CFC is the base requirement for 67xx linecard operation and is a zero cost option. The daughtercard houses two ASICs that function only as a bus interface. That is, these ASICs form truncated or compact headers, which are sent to the central PFC3* for forwarding decisions.

As the name implies, the CFC is only used for centralized forwarding. The centralized forwarding rate for the Catalyst 6500 is 30 Mpps, maximum. The CFC does not provide any local forwarding capabilities. This is different from any other type of module on the Catalyst 6500; typically, the centralized forward capability is inherent to the baseboard, and any daughtercards provide additional (non-standard) functionality.

When a DFC3** is added to a 67xx linecard, the CFC needs to be removed. The DFC3 daughtercard provides distributed forwarding (dCEF). The CFC and DFC3 each use the same linecard connector, so they are mutually exclusive on a particular module.

* PFC3 can be PFC3A, PFC3B, or PFC3BXL.

** DFC3 daughtercard options are WS-F6700-DFC3A, WS-F6700-DFC3B, WS-F6700-DFC3BXL, WS-F6700-DFC3C, or WS-F6700-DFC3CXL.

Q. What happens if line cards with different PFC modes are inserted in a Virtual Switching System (VSS) with PFC3C?

A. The VSS mode is currently supported in PFC3C or PFC3CXL only, and is not supported with any DFC3A, DFC3B, or DFC3BXL in the system. When you have a PFC3C/CXL and a line card with one of those DFCs inserted, the line card stays offline until the system is reloaded. This reload allows for system-wide PFC mode to be renegotiated, but it causes the system to run at the lower common denominator mode. All VSS functionality is disabled when the DFC-3B or lesser mode is enabled.

Q. If the WS-X67xx line cards are required in Virtual Switching System (VSS), is the DFC3C or DFC3CXL also required, or could the default CFC be used?

A. There is no restriction from the VSS to use DFC3C / DFC3CXL. Similar to normal switch requirements, either line cards with DFC3C/CXL can be used, or only the CFC can be used.

Note that the 8-port 10G card (WS-X6708-10G-3C/XL) must have a DFC installed for it to be operational.

Q. Are the ingress and egress interface queues supported on the WS-X6704-10GE when used with a DFC3CXL in 12.2(33)SRB4? Also are these queues specific to the module/port, DFC type, or both?

A. On WS-X6704-10GE in DFC3CXL, Transmit queues 1p7q8t and Receive queues 8q8t are supported as ingress and egress queues.

WS-X6704-10GE has the same queues as DFC-3CXL and DFC-3BXL. The difference is only between 6704/CFC and 6704/DFC. The **show interfaces interface interface-number**

capabilities *module number* can display this information.

Related Information

- **Cisco Catalyst 6500 Architecture White Paper**
 - **Cisco Catalyst 6500 Supervisor Engine 32 Architecture**
 - **LAN Switches Product Support**
 - **LAN Switching Technology Support**
 - **Technical Support & Documentation – Cisco Systems**
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jan 26, 2009

Document ID: 107258
