Guidelines and Limitations

This chapter includes the following sections:

- Verified Limits for FCoE, on page 1
- FCoE, on page 1

Verified Limits for FCoE

The configuration limits are documented in the *Cisco Nexus 7000 Series NX-OS Verified Scalability Guide*.

FCoE

- VDCs apply only to Cisco Nexus 7000 Series Switches.
- You cannot enable FCoE on default VLAN.
- The QoS policy must be the same on all Cisco FCoE switches in the network.
- Beginning with Cisco NX-OS Release 6.1, FCoE is supported on F2 and F2e Series modules. F3 Series modules are supported from Cisco NX-OS Release 6.2(6) onwards.
  - FCoE supports only F2e (SFP+) modules.
  - FCoE does not support F2e (Copper) modules.

FCoE VDC

FCoE in a dedicated storage VDC has the following guidelines:

- Enable the FCoE feature set in only one VDC.
- Create VLANs in the FCoE allocated VLAN range.
- Do not enable any other features other than storage-related features in the dedicated FCoE VDC.
- Allocate resources for the dedicated FCoE VDC from an F Series module, such as the 32-port 10-Gigabit Ethernet I/O module (PID N7K-F132XP-15) .
- Rollback is not supported in a storage VDC.

- For Cisco NX-OS Release 7.2(0)D1(1), ports from only 24 FEXes can be shared to storage VDC. System will not restrict the user to go beyond 24 but, more than 24 is not tested and not supported.
- FCoE on F2, F2e, and F3 Series modules is supported with the Supervisor 2 module (N7K-SUP2 for Cisco Nexus 7000 Series devices) and the Supervisor 2E module (N77-SUP2E for Cisco Nexus 7700 Series devices and N7K-SUP2E for Cisco Nexus 7000 Series devices).
- In order to enable FCoE over FEX on the storage VDC, you must execute the **allow feature-set FEX** command from the Admin or default VDC beforehand for storage VDC. FCoE over FEX is available from 7.2(0)D1(1) and onwards.
- IVR (Inter VSAN route) zone configuration is not supported for FCoE over FEX.
- F3 Fiber Channel over Ethernet (FCoE) feature licensing is supported from 7.2.0 release onwards. To downgrade to the older version of the image 6.2.x, first uninstall the F3 FCoE license and then proceed. For more information about licensing, refer Cisco NX-OS Licensing Guide.
- F2, F2e, and F3 Series modules can co-exist in the same VDC. This applies to both LAN and storage VDCs.
- F1 and F3 Series modules cannot co-exist in the same VDC. This applies to both LAN and storage VDCs.
- F1 and F2 series modules cannot exist in the same VDC. This applies to both LAN and storage VDCs.
- Use the **limit-resource module-type** command in the admin or default VDC to assign module resources such as F1, F2, F2e and F3 to a storage VDC. The supported line card modules are F1, F2, F2e and F3.
- When you configure a multi-hop FCoE, ensure that you use the same no-drop classes on both sides. Priority flow control does not work when you use different no-drop classes. Use the **show interface priority-flow-control** command to verify the priority flow control operation.

Shared Interfaces

Note

- Any change in protocol state that flaps the parent port of a shared interface because of any port feature also affects the FCoE traffic on the storage vdc.
- 1500 MTU do not carry FCoE traffic in all FCoE supported platforms.

The following interface config modes are not allowed while sharing an interface from Ethernet vdc to a storage vdc:

- SPAN destination
- Private VLAN mode
- Port-channel interfaces
- Access mode
- mac-packet-classify
- Interfaces that are part of a VLAN that has an associated QoS policy

Shared Ethernet interfaces must be in trunk mode and only shared with one other VDC.

Storage VDC

Configuring a VDC for the Out-Of-Band (OOB) management interface mgmt0 is accomplished with the **vrf context management** command. However, a storage VDC does not support VRF, so configuring mgmt0 requires a different approach.

The following table shows how to configure mgmt 0 for a VDC and for a storage VDC:

Configuring mgmt 0 for VDC	Configuring mgmt 0 for storage VDC

`vrf context management`
`ip route 0.0.0.0/0 default_gateway`
	`interface mgmt 0`
	`ip address mgmt0_ip_address mgmt0_subnet_mask`
	`no shut`
	`ip route 0.0.0.0/0 default_gateway`
	Note The **ip route** command specifies the default route that points to the default gateway.

where

- *mgmt0_ip_address* is the mgmt0 IPv4 address.
- *mgmt0_subnet_mask* is the mgmt0 IPv4 netmask.
- *default_gateway* is the IPv4 address of the default-gateway.

For more information about VDC, see the Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide.

Multi-Hop FCoE Distance Configuration

In Multi-Hop FCoE, when a device sends a pause, the interface that generates the pause frame must have an ingress queue with a buffer space large enough to buffer twice the link distance. This is because, when the pause is generated the wire might get congested. By the time the adjacent device receives or processes the generated pause frame, the wire might get congested again. Therefore, the device that generates the pause must have the ability to buffer twice the link distance.

As per calculations, there can be more than 100 packets traveling on the 10 kilometer link. Due to an ASIC limitation, the F1 series line card does not support lossless FCoE on a link greater than or equal to 10 kilometers.

For more information about Multi-Hop FCoE distance limitations, see http://www.cisco.com/c/en/us/support/docs/switches/nexus-7000-series-switches/117785-probsol-nexus7000-00.html

The F3 line cards support long haul lossless distance of up to 40 kilometers. In Cisco NX-OS Release 7.2(0) and later, you can change the ingress queuing buffer configuration.

Table 1: Buffer Tuning Table for FCoE Long Distance on F2, F2E, and F3 Line Cards

Distance	Line Card	SFP	Ingress Buffer Queue-Limit
< 5 km	F2/F2e	LR	60% no-drop and 40% drop queue
> 5 km - 10 km	F2/F2e	LR	70% no-drop and 30% drop queue

Distance	Line Card	SFP	Ingress Buffer Queue-Limit
> 10 km - 40 km	F2/F2e	ER	80% no-drop and 20% drop queue
< 10 km	F3	LR	90% no-drop and 10% drop queue
< 40 km	F3	ER	90% no-drop and 10% drop queue