Quality of Service

Quality of Service

Cisco UCS provides the following methods to implement quality of service:

- System classes that specify the global configuration for certain types of traffic across the entire system
- QoS policies that assign system classes for individual vNICs
- Flow control policies that determine how uplink Ethernet ports handle pause frames

Global QoS changes made to the QoS system class may result in brief data-plane interruptions for all traffic. Some examples of such changes are:

- Changing the MTU size for an enabled class
- Changing packet drop for an enabled class
- Changing the CoS value for an enabled class

Guidelines and Limitations for Quality of Service on Cisco UCS 6300 Series Fabric Interconnect

- Cisco UCS 6300 Series Fabric Interconnect uses a shared buffer for all system classes.
- Multicast optimization is not supported.
- When you change the QoS parameters for any class causes traffic disruption to all classes. The following table lists the changes in the QoS system class and the conditions that trigger a system reboot.

QoS System class status	Condition	FI Reboot Status
Enabled	Change between drop and no drop	Yes

QoS System class status	Condition	FI Reboot Status
No-drop	Change between enable and disable	Yes
Enable and no-drop	Change in MTU size	Yes

• The subordinate FI reboots first as a result of the change in the QoS system class , followed by the primary FI.

> ✎
>
> **Note**   Cisco UCS Manager displays a prompt that the fabric interconnect will reboot when the system policy is changed.

• **show queuing interface** command is not supported.

Guidelines and Limitations for Quality of Service on Cisco UCS Mini

• Cisco UCS Mini uses a shared buffer for all system classes.

• The bronze class shares the buffer with SPAN. We recommend using either SPAN or the bronze class.

• Multicast optimization is not supported.

• Changing the QoS parameters for any class causes traffic disruption to all classes.

• When mixing Ethernet and FC or FCoE traffic, the bandwidth distribution is not equal.

• Multiple streams of traffic from the same class may not be distributed equally.

• Use the same CoS values for all no-drop policies to avoid any FC or FCoE performance issues.

• Only the platinum and gold classes support no-drop policies.

• **show queuing interface** command is not supported.

Configuring System Classes

System Classes

Cisco UCS uses Data Center Ethernet (DCE) to handle all traffic inside a Cisco UCS domain. This industry standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. Two virtual lanes are reserved for internal system and management traffic. You can configure quality of service (QoS) for the other six virtual lanes. System classes determine how the DCE bandwidth in these six virtual lanes is allocated across the entire Cisco UCS domain.

Each system class reserves a specific segment of the bandwidth for a specific type of traffic, which provides a level of traffic management, even in an oversubscribed system. For example, you can configure the **Fibre Channel Priority** system class to determine the percentage of DCE bandwidth allocated to FCoE traffic.

The following table describes the system classes that you can configure.

Table 1: System Classes

System Class	Description
Platinum Gold Silver Bronze	A configurable set of system classes that you can include in the QoS policy for a service profile. Each system class manages one lane of traffic. All properties of these system classes are available for you to assign custom settings and policies. For Cisco UCS Mini, packet drop can only be disabled on the platinum and gold classes. Only one platinum and one gold class can be configured as a no drop class at a time.
Best Effort	A system class that sets the quality of service for the lane reserved for basic Ethernet traffic. Some properties of this system class are preset and cannot be modified. For example, this class has a drop policy that allows it to drop data packets if required. You cannot disable this system class.
Fibre Channel	A system class that sets the quality of service for the lane reserved for Fibre Channel over Ethernet traffic. Some properties of this system class are preset and cannot be modified. For example, this class has a no-drop policy that ensures it never drops data packets. You cannot disable this system class. **Note** FCoE traffic has a reserved QoS system class that should not be used by any other type of traffic. If any other type of traffic has a CoS value that is used by FCoE, the value is remarked to 0.

Configuring a System Class

The type of adapter in a server might limit the maximum MTU supported. For example, network MTU above the maximums might cause the packet to be dropped for the following adapters:

- The Cisco UCS M71KR CNA adapter, which supports a maximum MTU of 9216.

- The Cisco UCS 82598KR-CI adapter, which supports a maximum MTU of 14000.

SUMMARY STEPS

1. UCS-A# **scope eth-server**
2. UCS-A /eth-server # **scope qos**
3. UCS-A /eth-server/qos # **scope eth-classified** {**bronze** | **gold** | **platinum** | **silver**}
4. UCS-A /eth-server/qos/eth-classified # **enable**
5. UCS-A /eth-server/qos/eth-classified # **set cos** *cos-value*
6. UCS-A /eth-server/qos/eth-classified # **set drop** {**drop** | **no-drop**}
7. UCS-A /eth-server/qos/eth-classified # **set mtu** {*mtu-value* | **fc** | **normal**}
8. UCS-A /eth-server/qos/eth-classified # **set multicast-optimize** {**no** | **yes**}
9. UCS-A /eth-server/qos/eth-classified # **set weight** {*weight-value* | **best-effort** | **none**}

10. UCS-A /eth-server/qos/eth-classified # **commit-buffer**

DETAILED STEPS

	Command or Action	Purpose
Step 1	UCS-A# **scope eth-server**	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # **scope qos**	Enters Ethernet server QoS mode.
Step 3	UCS-A /eth-server/qos # **scope eth-classified** {**bronze** | **gold** | **platinum** | **silver**}	Enters Ethernet server QoS Ethernet classified mode for the specified system class.
Step 4	UCS-A /eth-server/qos/eth-classified # **enable**	Enables the specified system class.
Step 5	UCS-A /eth-server/qos/eth-classified # **set cos** *cos-value*	Specifies the class of service for the specified system class. Valid class of service values are 0 to 6.
		Important Use the same CoS values on UCS and N5K for all the no-drop policies. To insure that end-to-end PFC works correctly, have the same QoS policy configured on all intermediate switches.
		Note When the CoS value is set to 0 in any QoS class, this causes the adapter to use the same queue for best effort and the QoS class. When traffic congestion occurs, best effort and the QoS class will share the bandwidth equally instead of using the weight configured in the QoS class.
Step 6	UCS-A /eth-server/qos/eth-classified # **set drop** {**drop** | **no-drop**}	Specifies whether the channel can drop packets or not.
		Note Changes saved to the drop displays the following warning message: Warning: The operation will cause momentary disruption to traffic forwarding.
Step 7	UCS-A /eth-server/qos/eth-classified # **set mtu** {*mtu-value* | **fc** | **normal**}	The maximum transmission unit, or packet size to be used. The maximum value for MTU is 9216.
		Note If the vNIC has an associated QoS policy, the MTU specified here must be equal to or less than the MTU specified in the associated QoS system class. If this MTU value exceeds the MTU value in the QoS system class, packets might get dropped during data transmission.
		Changes saved to the MTU displays the following warning message: Warning: The operation will cause momentary disruption to traffic forwarding.
Step 8	UCS-A /eth-server/qos/eth-classified # **set multicast-optimize** {**no** | **yes**}	Specifies whether the class is optimized to for sending multicast packets.

	Command or Action	Purpose
Step 9	UCS-A /eth-server/qos/eth-classified # **set weight** {*weight-value*	**best-effort**	**none**}	Specifies the relative weight for the specified system class. Valid weight values are 0 to 10.
Step 10	UCS-A /eth-server/qos/eth-classified # **commit-buffer**	Commits the transaction to the system configuration.

Example

The following example shows how to enable the platinum system class, allow the channel to drop packets, set the class of service to 6, set the MTU to normal, set the relative weight to 5, and commit the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope qos
UCS-A /eth-server/qos # scope eth-classified platinum
UCS-A /eth-server/qos/eth-classified # enable
UCS-A /eth-server/qos/eth-classified* # set drop drop
Warning: The operation will cause momentary disruption to traffic forwarding
UCS-A /eth-server/qos/eth-classified* # set cos 6
UCS-A /eth-server/qos/eth-classified* # set mtu normal
Warning: The operation will cause momentary disruption to traffic forwarding
UCS-A /eth-server/qos/eth-classified* # set weight 5
UCS-A /eth-server/qos/eth-classified* # commit-buffer
UCS-A /eth-server/qos/eth-classified #
```

Disabling a System Class

If you disable a system class that is used in a QoS policy, Cisco UCS Manager uses the system class configured with CoS 0 for traffic on servers that are configured with the QoS policy. If no system class is configured as CoS 0, the Best Effort system class is used. You cannot disable the Best Effort or Fibre Channel system classes.

SUMMARY STEPS

1. UCS-A# **scope eth-server**
2. UCS-A /eth-server # **scope qos**
3. UCS-A /eth-server/qos # **scope eth-classified** {**bronze** | **gold** | **platinum** | **silver**}
4. UCS-A /eth-server/qos/eth-classified # **disable**
5. UCS-A /eth-server/qos/eth-classified # **commit-buffer**

DETAILED STEPS

	Command or Action	Purpose
Step 1	UCS-A# **scope eth-server**	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # **scope qos**	Enters Ethernet server QoS mode.
Step 3	UCS-A /eth-server/qos # **scope eth-classified** {**bronze**	**gold**	**platinum**	**silver**}	Enters Ethernet server QoS Ethernet classified mode for the specified system class.
Step 4	UCS-A /eth-server/qos/eth-classified # **disable**	Disables the specified system class.

	Command or Action	Purpose
Step 5	UCS-A /eth-server/qos/eth-classified # **commit-buffer**	Commits the transaction to the system configuration.

Example

The following example disables the platinum system class and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope qos
UCS-A /eth-server/qos # scope eth-classified platinum
UCS-A /eth-server/qos/eth-classified # disable
UCS-A /eth-server/qos/eth-classified* # commit-buffer
UCS-A /eth-server/qos/eth-classified #
```

Configuring Quality of Service Policies

Quality of Service Policy

A quality of service (QoS) policy assigns a system class to the outgoing traffic for a vNIC or vHBA. This system class determines the quality of service for that traffic. For certain adapters, you can also specify additional controls on the outgoing traffic, such as burst and rate.

You must include a QoS policy in a vNIC policy or vHBA policy and then include that policy in a service profile to configure the vNIC or vHBA.

Configuring a QoS Policy

SUMMARY STEPS

1. Switch-A# **scope org** *org-name*
2. Switch-A /org # **create qos-policy** *policy-name*
3. Switch-A /org/qos-policy # **create egress-policy**
4. Switch-A /org/qos-policy/egress-policy # **set host-cos-control** {**full** | **none**}
5. Switch-A /org/qos-policy/egress-policy # **set prio** *sys-class-name*
6. Switch-A /org/qos-policy/egress-policy # **set rate** {**line-rate** | *kbps*} **burst** *bytes*
7. Switch-A /org/qos-policy/egress-policy # **commit-buffer**

DETAILED STEPS

	Command or Action	Purpose
Step 1	Switch-A# **scope org** *org-name*	Enters org mode for the specified organization. To enter the default org mode, type **/** as the *org-name* .
Step 2	Switch-A /org # **create qos-policy** *policy-name*	Creates the specified QoS policy, and enters org QoS policy mode.

	Command or Action	Purpose
Step 3	Switch-A /org/qos-policy # **create egress-policy**	Creates the egress policy (for both vNICs and vHBAs) to be used by the QoS policy, and enters org QoS policy egress policy mode.
Step 4	Switch-A /org/qos-policy/egress-policy # **set host-cos-control** {**full** | **none**}	(Optional) Specifies whether the host or Cisco UCS Manager controls the class of service (CoS) for a vNIC. This setting has no effect on a vHBA. Use the **full** keyword to have the host control the CoS. If the packet has a valid CoS value, the host uses that value. Otherwise, it uses the CoS value associated with the specified class priority. Use the **none** keyword to have Cisco UCS Manager use the CoS value associated with the specified priority.
Step 5	Switch-A /org/qos-policy/egress-policy # **set prio** *sys-class-name*	Specifies the system class to be used for the egress policy. The *sys-class-name* argument can be one of the following class keywords: • **C**—Use this priority for QoS policies that control vHBA traffic only. • **Platinum**—Use this priority for QoS policies that control vNIC traffic only. • **Gold**—Use this priority for QoS policies that control vNIC traffic only. • **Silver**—Use this priority for QoS policies that control vNIC traffic only. • **Bronze**—Use this priority for QoS policies that control vNIC traffic only. • **Best Effort**—Do not use this priority. It is reserved for the Basic Ethernet traffic lane. If you assign this priority to a QoS policy and configure another system class as CoS 0, Cisco UCS Manager does not default to this system class. It defaults to the priority with CoS 0 for that traffic.
Step 6	Switch-A /org/qos-policy/egress-policy # **set rate** {**line-rate** | *kbps*} **burst** *bytes*	Specifies the expected average rate of traffic. Traffic that falls under this rate will always conform. The default is **line-rate**, which equals a value of 10,000,000. The minimum value is 8, and the maximum value is 40,000,000. Rate limiting is supported only on vNICs on the Cisco UCS VIC-1240 Virtual Interface Card and Cisco UCS VIC-1280 Virtual Interface Card. The Cisco UCS M81KR Virtual Interface Card supports rate limiting on both vNICs and vHBAs.
Step 7	Switch-A /org/qos-policy/egress-policy # **commit-buffer**	Commits the transaction to the system configuration.

Example

The following example creates a QoS policy for vNIC traffic, assigns the platinum system class and sets the rate limit (traffic rate and burst size) for the egress policy, and commits the transaction:

```
Switch-A# scope org /
Switch-A /org # create qos-policy VnicPolicy34
Switch-A /org/qos-policy* # create egress-policy
Switch-A /org/qos-policy/egress-policy* # set prio platinum
Switch-A /org/qos-policy/egress-policy* # set rate 5000000 burst 65000
Switch-A /org/qos-policy/egress-policy* # commit-buffer
Switch-A /org/qos-policy/egress-policy #
```

The following example creates a QoS policy for vHBA traffic, assigns the fc (Fibre Channel) system class and sets the rate limit (traffic rate and burst size) for the egress policy, and commits the transaction:

```
Switch-A# scope org /
Switch-A /org # create qos-policy VhbaPolicy12
Switch-A /org/qos-policy* # create egress-policy
Switch-A /org/qos-policy/egress-policy* # set prio fc
Switch-A /org/qos-policy/egress-policy* # set rate 5000000 burst 65000
Switch-A /org/qos-policy/egress-policy* # commit-buffer
Switch-A /org/qos-policy/egress-policy #
```

What to do next

Include the QoS policy in a vNIC or vHBA template.

Deleting a QoS Policy

If you delete a QoS policy that is in use or you disable a system class that is used in a QoS policy, any vNIC or vHBA that uses that QoS policy is assigned to the Best Effort system class or to the system class with a CoS of 0. In a system that implements multitenancy, Cisco UCS Manager first attempts to find a matching QoS policy in the organization hierarchy.

SUMMARY STEPS

1. UCS-A# **scope org**   *org-name*
2. UCS-A /org # **delete qos-policy**   *policy-name*
3. UCS-A /org # **commit-buffer**

DETAILED STEPS

	Command or Action	Purpose
Step 1	UCS-A# **scope org** *org-name*	Enters organization mode for the specified organization. To enter the root organization mode, type **/** as the *org-name*.
Step 2	UCS-A /org # **delete qos-policy** *policy-name*	Deletes the specified QoS policy.
Step 3	UCS-A /org # **commit-buffer**	Commits the transaction to the system configuration.

Example

The following deletes the QoS policy named QosPolicy34 and commits the transaction:

```
UCS-A# scope org /
UCS-A /org # delete qos-policy QosPolicy34
UCS-A /org* # commit-buffer
UCS-A /org #
```

Configuring Flow Control Policies

Flow Control Policy

Flow control policies determine whether the uplink Ethernet ports in a Cisco UCS domain send and receive IEEE 802.3x pause frames when the receive buffer for a port fills. These pause frames request that the transmitting port stop sending data for a few milliseconds until the buffer clears.

For flow control to work between a LAN port and an uplink Ethernet port, you must enable the corresponding receive and send flow control parameters for both ports. For Cisco UCS, the flow control policies configure these parameters.

When you enable the send function, the uplink Ethernet port sends a pause request to the network port if the incoming packet rate becomes too high. The pause remains in effect for a few milliseconds before traffic is reset to normal levels. If you enable the receive function, the uplink Ethernet port honors all pause requests from the network port. All traffic is halted on that uplink port until the network port cancels the pause request.

Because you assign the flow control policy to the port, changes to the policy have an immediate effect on how the port reacts to a pause frame or a full receive buffer.

Configuring a Flow Control Policy

Before you begin

Configure the network port with the corresponding setting for the flow control that you need. For example, if you enable the send setting for flow-control pause frames in the policy, ensure that the receive parameter in the network port is set to on or to desired. If you want the Cisco UCS port to receive flow-control frames, ensure that the send parameter is set to on or to desire on the network port. If you do not want to use flow control, you can set the send and receive parameters on the network port to off.

SUMMARY STEPS

1. UCS-A# **scope eth-uplink**
2. UCS-A /eth-uplink # **scope flow-control**
3. UCS-A /eth-uplink/flow-control # **create policy** *policy-name*
4. UCS-A /eth-uplink/flow-control/policy # **set prio** *prio-option*
5. UCS-A /eth-uplink/flow-control/policy # **set receive** *receive-option*
6. UCS-A /eth-uplink/flow-control/policy # **set send** *send-option*
7. UCS-A /eth-uplink/flow-control/policy # **commit-buffer**

DETAILED STEPS

	Command or Action	Purpose
Step 1	UCS-A# **scope eth-uplink**	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # **scope flow-control**	Enters Ethernet uplink flow control mode.
Step 3	UCS-A /eth-uplink/flow-control # **create policy** *policy-name*	Creates the specified flow control policy.
Step 4	UCS-A /eth-uplink/flow-control/policy # **set prio** *prio-option*	Specifies one of the following flow control priority options: • **auto** —The Cisco UCS system and the network negotiate whether PPP will be used on this fabric interconnect. • **on** —PPP is enabled on this fabric interconnect.
Step 5	UCS-A /eth-uplink/flow-control/policy # **set receive** *receive-option*	Specifies one of the following flow control receive options: • **off** —Pause requests from the network are ignored and traffic flow continues as normal. • **on** —Pause requests are honored and all traffic is halted on that uplink port until the network cancels the pause request.
Step 6	UCS-A /eth-uplink/flow-control/policy # **set send** *send-option*	Specifies one of the following flow control send options: • **off** —Traffic on the port flows normally regardless of the packet load. • **on** —The Cisco UCS system sends a pause request to the network if the incoming packet rate becomes too high. The pause remains in effect for a few milliseconds before traffic is reset to normal levels.
Step 7	UCS-A /eth-uplink/flow-control/policy # **commit-buffer**	Commits the transaction to the system configuration.

Example

The following configures a flow control policy and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope flow-control
UCS-A /eth-uplink/flow-control # create policy FlowControlPolicy23
UCS-A /eth-uplink/flow-control/policy* # set prio auto
UCS-A /eth-uplink/flow-control/policy* # set receive on
UCS-A /eth-uplink/flow-control/policy* # set send on
UCS-A /eth-uplink/flow-control/policy* # commit-buffer
UCS-A /eth-uplink/flow-control/policy #
```

What to do next

Associate the flow control policy with an uplink Ethernet port or port channel.

Deleting a Flow Control Policy

SUMMARY STEPS

1. UCS-A# **scope eth-uplink**
2. UCS-A /eth-uplink # **scope flow-control**
3. UCS-A /eth-uplink/flow-control # **delete policy** *policy-name*
4. UCS-A /eth-uplink/flow-control # **commit-buffer**

DETAILED STEPS

	Command or Action	**Purpose**
Step 1	UCS-A# **scope eth-uplink**	Enters Ethernet uplink mode.
Step 2	UCS-A /eth-uplink # **scope flow-control**	Enters Ethernet uplink flow control mode.
Step 3	UCS-A /eth-uplink/flow-control # **delete policy** *policy-name*	Deletes the specified flow control policy.
Step 4	UCS-A /eth-uplink/flow-control # **commit-buffer**	Commits the transaction to the system configuration.

Example

The following example deletes the flow control policy named FlowControlPolicy23 and commits the transaction:

```
UCS-A# scope eth-uplink
UCS-A /eth-uplink # scope flow-control
UCS-A /eth-uplink/flow-control # delete policy FlowControlPolicy23
UCS-A /eth-uplink/flow-control* # commit-buffer
UCS-A /eth-uplink/flow-control #
```

Configuring Slow Drain

QoS Slow Drain Device Detection and Mitigation

All data traffic between end devices in the fabric is carried by Fibre Channel services that use link-level, per-hop-based, and buffer-to-buffer flow control. These classes of service do not support end-to-end flow control. When slow devices are attached to the fabric, the end devices do not accept the frames at the configured or negotiated rate. The slow devices lead to an Inter-Switch Link (ISL) credit shortage in the traffic that is destined for these devices, and they congest the links. The credit shortage affects the unrelated flows in the fabric that use the same ISL link even though destination devices do not experience a slow drain.

Similarly, in End-Host Mode, if a server that is directly attached to the Fabric Interconnect receives traffic slowly, it may congest the uplink port shared by other servers. If a slow server is attached to a HIF port on FEX/IOM, it may congest the fabric port and/or uplink port.

Cisco UCS Manager Release 4.0(2) introduces the QoS Slow Drain Detection and Mitigation feature on Cisco UCS 6454 Fabric Interconnects. This feature provides various enhancements that enable you to detect slow drain devices that cause congestion in the network, and also mitigate it. The enhancements are mainly on the edge ports and core ports that connect to the slow drain devices. This is done to minimize the frames stuck condition in the edge and core ports due to slow drain devices that are causing an ISL blockage. To avoid or minimize the stuck condition, you can configure smaller frame timeout for the ports. A smaller frame timeout value helps to alleviate the slow drain condition that affects the fabric by dropping the packets on the edge ports sooner than the time they actually get timed out. This function frees the buffer space in ISL, which can be used by other unrelated flows that do not experience the slow drain condition.

In this release, slow drain detection and mitigation is supported on the following ports:

- FCoE
- Back-plane

Configuring Slow Drain Detection

SUMMARY STEPS

1. UCS-A# **scope eth-server**
2. UCS-A /eth-server # **scope qos**
3. UCS-A /eth-server/qos # **scope slow-drain**
4. UCS-A /eth-server/qos/slow-drain #**set fcoe-admin-state** {**disable** | **enable**}
5. UCS-A /eth-server/qos/slow-drain* # **commit-buffer**

DETAILED STEPS

	Command or Action	Purpose
Step 1	UCS-A# **scope eth-server**	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # **scope qos**	Enters Ethernet server QoS mode.
Step 3	UCS-A /eth-server/qos # **scope slow-drain**	Enters Ethernet server QoS slow drain mode.
Step 4	UCS-A /eth-server/qos/slow-drain #**set fcoe-admin-state** {**disable**	**enable**}	Sets the FCoE admin state to one of the following: • disable—Slow drain detection is disabled • enable—Slow drain detection is enabled
Step 5	UCS-A /eth-server/qos/slow-drain* # **commit-buffer**	Commits the transaction to the system configuration.

Example

The following example enables slow drain detection on FCoE ports and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope qos
UCS-A /eth-server/qos # scope slow-drain
UCS-A /eth-server/qos/slow-drain # set fcoe-admin-state enable
UCS-A /eth-server/qos/slow-drain* # commit-buffer
UCS-A /eth-server/qos/slow-drain #
```

Configuring Slow Drain Timers

While configuring slow drain timeout timers, you can select the timeout value from the list of allowed values. You cannot configure custom timeout values.

SUMMARY STEPS

1. UCS-A# **scope eth-server**
2. UCS-A /eth-server # **scope qos**
3. UCS-A /eth-server/qos # **scope slow-drain**
4. UCS-A /eth-server/qos/slow-drain #**set core-port-timer** {**100** | **200** | **300** | **400** | **500** | **600** | **700** | **800** | **900** | **1000**}
5. UCS-A /eth-server/qos/slow-drain* #**set edge-port-timer** {**100** | **200** | **300** | **400** | **500** | **600** | **700** | **800** | **900** | **1000**}
6. UCS-A /eth-server/qos/slow-drain* #**set backplane-port-timer** { **200** | **300** | **400** | **500** | **600** | **700** | **800** | **900** | **1000**}
7. UCS-A /eth-server/qos/slow-drain* # **commit-buffer**

DETAILED STEPS

	Command or Action	Purpose
Step 1	UCS-A# **scope eth-server**	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # **scope qos**	Enters Ethernet server QoS mode.
Step 3	UCS-A /eth-server/qos # **scope slow-drain**	Enters Ethernet server QoS slow drain mode.
Step 4	UCS-A /eth-server/qos/slow-drain #**set core-port-timer** {**100**	**200**	**300**	**400**	**500**	**600**	**700**	**800**	**900**	**1000**}	Sets the core FCoE port timeout to one of the listed values. The default timeout value is 500 ms.
Step 5	UCS-A /eth-server/qos/slow-drain* #**set edge-port-timer** {**100**	**200**	**300**	**400**	**500**	**600**	**700**	**800**	**900**	**1000**}	Sets the edge FCoE port timeout to one of the listed values. The default timeout value is 500 ms.
Step 6	UCS-A /eth-server/qos/slow-drain* #**set backplane-port-timer** { **200**	**300**	**400**	**500**	**600**	**700**	**800**	**900**	**1000**}	Sets the backplane port timeout to one of the listed values. The default timeout value is 1000 ms.
Step 7	UCS-A /eth-server/qos/slow-drain* # **commit-buffer**	Commits the transaction to the system configuration.

Example

The following example configures the slow drain timers and commits the transaction:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope qos
UCS-A /eth-server/qos # scope slow-drain
UCS-A /eth-server/qos/slow-drain # set core-port-timer 500
UCS-A /eth-server/qos/slow-drain* # set edge-port-timer 500
UCS-A /eth-server/qos/slow-drain* # set backplane-port-timer 1000
UCS-A /eth-server/qos/slow-drain* # commit-buffer
UCS-A /eth-server/qos/slow-drain #
```

Displaying Slow Drain Settings

SUMMARY STEPS

1. UCS-A# **scope eth-server**
2. UCS-A /eth-server # **scope qos**
3. UCS-A /eth-server/qos # **show slow-drain**

DETAILED STEPS

	Command or Action	Purpose
Step 1	UCS-A# **scope eth-server**	Enters Ethernet server mode.
Step 2	UCS-A /eth-server # **scope qos**	Enters Ethernet server QoS mode.
Step 3	UCS-A /eth-server/qos # **show slow-drain**	Displays QoS slow drain settings.

Example

The following example displays the slow drain settings:

```
UCS-A# scope eth-server
UCS-A /eth-server # scope qos
UCS-A /eth-server/qos # show slow-drain

QoS Slow Drain:
    Admin State for QoS Slow Drain for Physical FCoE Ports: Enabled
    QoS Slow Drain: Timer value for Core Physical FCoE Ports: 100
    QoS Slow Drain: Timer value for Edge Physical FCoE Ports: 100
    QoS Slow Drain: Timer value for Backplane Ports: 1000
UCS-A /eth-server/qos #
```

Priority Flow Control Watchdog Interval

A PFC storm may occur in the network from a malfunctioning NIC or switch, where the Priority Flow Control (PFC) frames are propagated to all senders causing a complete stall in traffic in the network. To mitigate a PFC storm, a PFC watchdog can be used. A PFC watchdog interval can be configured to detect whether packets in a no-drop queue are being drained within a specified time period. If packets are present in buffer longer than the configured time period and after the time period expires, all outgoing packets are dropped on the interfaces that match the PFC queue that is not being drained.

Note   For VIC 6332 Fabric Interconnects, Priority Flow Watchdog functionality does not operate on all 6332 Fabric Interconnect ports, due to ASIC limitations. These port limitations are as follows:

- For VIC 6332, it will not operate on Ports 1/28-32 (40G uplink-only ports).

- For VIC 6332-16UP it will not operate on the following ports: Ethernet1/1-16 (Combined Ethernet/FC ports) or 1/35-40 (40G uplink-only ports).

For VIC 6332 with Priority Flow Control Watchdog, use only supported ports as needed.

-

-

Configuring a Priority Flow Control Watchdog Interval

SUMMARY STEPS

1. UCS-A# **scope eth-server**
2. UCS-A /eth-server # **scope pfc**
3. UCS-A /eth-server/pfc # **set wd-admin-state {on | off}**
4. UCS-A /eth-server/pfc # **set wd-interval 500**
5. UCS-A /eth-server/pfc # **set wd-shutdown-multiplier 1**
6. UCS-A /eth-server/pfc* # commit-buffer

DETAILED STEPS

	Command or Action	Purpose
Step 1	UCS-A# **scope eth-server**	Enters the Ethernet server mode.
Step 2	UCS-A /eth-server # **scope pfc**	Enters the Ethernet server PFC mode.
Step 3	UCS-A /eth-server/pfc # **set wd-admin-state {on	off}**	Globally enables or disables the PFC watchdog interval for all interfaces. The default value is **on**.
Step 4	UCS-A /eth-server/pfc # **set wd-interval 500**	Specifies the watchdog interval value. The valid range is from 100 to 1000 milliseconds. The default value is 100.

	Command or Action	Purpose
Step 5	UCS-A /eth-server/pfc # **set wd-shutdown-multiplier 1**	Specifies when to declare the PFC queue as struck. The valid range is from 1 to 10. The default value is 1.
Step 6	UCS-A /eth-server/pfc* # commit-buffer	Commits the transaction to the system configuration.

The watchdog interval, polling interval, and shutdown multiplier are configured.

Example

The following example shows how to configure the watchdog interval, polling interval, and shutdown multiplier, and then commit the transaction.

```
UCS-A# scope eth-server
UCS-A /eth-server # scope pfc
UCS-A /eth-server/pfc # set wd-admin-state on
 UCS-A /eth-server/pfc # set wd-interval 500
UCS-A /eth-server/pfc # set wd-shutdown-multiplier 1
UCS-A /eth-server/pfc* # commit-buffer
```

Viewing the Watchdog Settings

SUMMARY STEPS

1. UCS-A# **scope eth-server**
2. UCS-A /eth-server # **show pfc details**

DETAILED STEPS

	Command or Action	Purpose
Step 1	UCS-A# **scope eth-server**	Enters the Ethernet server mode.
Step 2	UCS-A /eth-server # **show pfc details**	Displays the PFC watchdog settings.

Example

The following example displays the watchdog settings:

```
UCS-A# scope eth-server
UCS-A /eth-server # show pfc details

    Global PFC watchdog configuration details:
    PFC watchdog interval: On
    PFC watchdog poll interval: 500
    PFC watchdog shutdown multiplier: 1
    Current Task:
```