

Slow-Drain Device Detection and SAN Congestion Prevention

Using Cisco MDS 9000 Series Multilayer Switches and Cisco Data Center Network Manager (DCNM)

Slow-drain device detection

Q **How do Cisco MDS 9000 Series Multilayer Switches detect SAN congestion and slow-drain devices, and at what granularity?**

A Cisco MDS 9000 Series Multilayer Switches detect SAN congestion and slow-drain devices using high-fidelity metrics that are available on Cisco® Fibre Channel port-ASICs. The following are a few important metrics:

- TxWait: The duration for which a switch port could not transmit a frame because of the lack of transmit buffer-to-buffer (B2B) credits. TxWait is collected every 2.5 microseconds, whereas automatic alerting is available at 1 second granularity.
- Slowport monitor: The duration for which a switch port could not transmit a frame due to the lack of transmit B2B credits. Slowport monitor is collected every 1 millisecond, whereas automatic alerting is available at 1 second granularity. Different from TxWait, slowport monitor alerts only on the continuous duration of transmit B2B credit unavailability.

- Transmit and receive B2B transitions to 0: The number of times when zero (0) B2B credits were available on a switch port.
- Transmit and receive link utilization: The percent utilization of a switch port. The raw metrics in bytes transmitted and received are collected in real time, whereas automatic alerting is available at a granularity of 10 seconds.
- Transmit and receive burst: The number of times switch port utilization was higher than the configured threshold. The raw metrics in bytes transmitted and received are collected in real time, whereas automatic alerting is available at a granularity of 10 seconds.

For more information, refer to the [congestion management chapter in the Cisco MDS 9000 Series Multilayer Switches Configuration Guide](#).

Q Can Cisco MDS 9000 switches send an automatic notification when congestion or a slow-drain device is detected?

A The Cisco MDS 9000 switches and Cisco Data Center Network Manager (DCNM) can send notifications using SNMP traps, syslog, email, and call-home capabilities. The Cisco Port-Monitor (PMON) feature provides a policy-based configuration to detect, notify, and take automatic port-guard actions to prevent congestion and slow drain.

For more information, refer to the [port-monitor section in the Cisco MDS 9000 Series Multilayer Switches Configuration Guide](#) and [port-monitor section in the Cisco MDS 9000 Series Multilayer Switches Configuration Guide](#).

Q How do I get historic trends and end-to-end correlation for slow-drain devices?

A Cisco Data Center Network Manager (DCNM) provides long-term trending and end-to-end correlation using the slow-drain analysis feature. I/O flow metrics, such as Exchange Completion Time (ECT), Data Access Latency (DAL), I/O sizes, IOPS, etc., are available using Cisco SAN Analytics technology.

For more information, refer to the [slow-drain analysis chapter in the Cisco DCNM SAN Management Configuration Guide](#) and [SAN Insights chapter in the Cisco DCNM SAN Management Configuration Guide](#).

Congestion prevention

Q How does the Cisco Dynamic Ingress Rate Limiting (DIRL) feature prevent SAN congestion and slow drain?

A In block-storage networks, the storage arrays don't send data by themselves, but the servers ask for data by initiating an I/O. Congestion in a SAN spreads because a server asks for more than it can ingest. The Cisco DIRL feature detects any symptoms of egress congestion on the switch ports. Then it limits ingress data to

prevent congestion in the egress direction. DIRL dynamically adapts the ingress traffic rate until the egress congestion goes away. By limiting ingress frames, DIRL also slows down the data-requesting frames (read I/O command) to the All-Flash Arrays (AFA). When the congested servers are allowed to ask for less, they will get less, and won't cause any congestion.

For more information, refer to the [congestion management chapter in the Cisco MDS 9000 Series Multilayer Switches configuration guides](#). [Congestion management chapter in the Cisco MDS 9000 Series Multilayer Switches Configuration Guide](#).

Q How does DIRL limit ingress traffic from the congested or slow-drain device? Does it drop the Fibre Channel frames?

A DIRL controls the flow of frames automatically by using the B2B credit pacing mechanism of the Cisco Fibre Channel ASICs. DIRL doesn't drop any frames.

Q What are the benefits of DIRL?

- A** The Cisco DIRL feature has the following benefits for the customer:
- End-device independent – Upgrading of end-devices is not needed to prevent spreading of congestion and slow drain. DIRL does not require the most current host bus adapter technology to mitigate congestion.
 - Adaptive – DIRL dynamically adjusts as per the traffic profile of the host.
 - Affordable – No additional license is required.
 - Easy adoption – DIRL is available on the Cisco MDS 9000 Series Multilayer Switches after a software-only upgrade.
 - No side effects – The rate-limiting is applied only to the congested host. Other non-congested hosts and storage ports are not affected.
 - Topology independent – DIRL works in edge-core, edge-core-edge, or collapsed core (single switch fabric) topologies.

Q Do Cisco MDS 9000 switches support Fibre Channel Notifications (FPIN) and signals?

A The Cisco MDS 9000 Series Multilayer Switches support Register Diagnostic Functions (RDF), Exchange Diagnostic Capabilities (EDC), Fabric Performance Impact Notifications (FPIN), and signals from Cisco MDS 9000 NX-OS Release 8.5(1) onward. After detecting the events, the Cisco MDS 9000 switches send notifications and signals to the end devices to inform them about congestion, peer congestion, link integrity, etc.

Q How is DIRL different from Fibre Channel Notifications (FPIN) and signals?

A The DIRL feature is completely integrated within the Cisco MDS 9000 Series Multilayer Switches and is available after a software-only upgrade. The switches detect congestion events and take automated preventive actions without any external dependency.

Fibre Channel notifications (FPINs) and signals have a dependency on the end devices. The Cisco MDS 9000 switches detect congestion events and then inform the end devices; taking an action in response is the responsibility of the end devices.

Q How is DIRL different from congestion isolation and de-isolation?

A The congestion isolation and de-isolation features work only in multi-switch fabrics with Inter-Switch Links (ISLs). Using congestion isolation, traffic that is going to a congested or slow-draining device is isolated in a different virtual link on the ISL.

DIRL is topology independent. It works in all types of environments - single-switch fabric, edge-core fabrics, or edge-core-edge fabrics. Using DIRL, traffic originating from a congested or slow-drain device is automatically paced to prevent the spreading of congestion.

Q Do DIRL, FPIN, congestion-isolation, and other congestion-prevention features require an additional license?

A Monitoring, alerting, and preventing congestion or slow drain (features such as DIRL, FPIN, congestion isolation, etc.) on Cisco MDS 9000 Series Multilayer Switches do not require any additional license.

Q How does a customer turn on DIRL?

A DIRL is enabled by leveraging the Port-Monitor (PMON) feature on Cisco MDS 9000 Series Multilayer Switches. DIRL can be a Cisco port-guard action for the threshold-based policies defined to detect transmit B2B credit unavailability (txwait), link overutilization (tx-datarate), traffic burst detection (tx-datarate-burst), etc.

For more information, refer to the [port-monitor section in the Cisco MDS 9000 Series Multilayer Switches Configuration Guide](#).