



Data Center Quality of Service

Data Center Bridging (DCB) Overview

General QoS Tips

- + Classify as close to the source as possible
 - + Ensures that flows land in the proper queues
- + Police unwanted flows as close to the source
 - + Drop on ingress is more efficient than drop on egress
- + Enable QoS End-to-End
 - + QoS is only as good as the weakest link
 - + E.g. forgot to mark traffic from UCS vNIC
- + Visualize the traffic flow; which nodes requires QoS?
 - + Follow the packet hop by hop
- + Recommended Resource
 - + [End-to-End QoS Network Design: Quality of Service for Rich-Media & Cloud Networks, Second Edition](#)

What is DCB?

- + Data Center Bridging (DCB) is the standard for vendor inter-op of advertising QoS attributes from the end host to its attached access switch
- + For more information...
 - + [Data Center Bridging \(DCB\) Task Group \(TG\)](#)
 - + DCB is under IEEE 802.1

What features does DCB include?

- + Priority Flow Control (PFC)
 - + Standard per IEEE 802.1Qbb
- + Enhanced Transmission Selection (ETS)
 - + Standard per IEEE 802.1Qaz
- + Congestion Notification
 - + Standard per IEEE 802.1Qau
- + DCB eXchange (DCBX)
 - + IEEE 802.1Qaz combined with LLDP (802.1AB)
 - + LLDP is used to discover neighbors

Understanding Priority Flow Control

- + Priority Flow Control (PFC)
 - + Similar to Ethernet Flow Control
 - + https://en.wikipedia.org/wiki/Ethernet_flow_control
 - + Goal of PFC is to provide lossless fabric for the storage protocols
 - + Let LAN traffic forward as normal and see input/output drops
 - + SAN traffic reacts to network conditions to be lossless
 - + PFC sends pause frames on a per CoS (802.1p) basis
 - + For example CoS 3 for FCoE
 - + Result is SAN traffic pauses, LAN traffic continues

Understanding Enhanced Transmission Selection

- + Similar to priority queueing and bandwidth reservations
- + ETS provides virtual interface queues
 - + Strict priority (LLQ)
 - + Guaranteed Bandwidth
 - + Best Effort
- + ETS options are controlled by the access switch and dynamically advertised to the downstream end-host

Understanding Congestion Notification

- + IEEE 802.1Qau is an Ethernet based congestion notification
 - + Congestion notification normally happens at upper layers like TCP
- + Once congestion is detected, rate limiters are applied to try to reduce the connection, causing “back pressure” on the senders

Understanding DCBX

- + Data Center Bridging Exchange (DCBX) is the method of how negotiation occurs for...
 - + Priority Flow Control (PFC)
 - + Enhanced Transmission Selection (ETS)
 - + Congestion Notification

Deploying DCBX

- + Access Switch enables feature FCoE
 - + Ideally QoS profile is auto-enabled
- + Switch advertises QoS settings through DCBX to host
 - + E.g. what is the LAN/SAN bandwidth split
- + Host uses this QoS config for outbound traffic
 - + Upstream switch can now control both inbound and outbound policy



<https://t.me/learningnets>



Data Center Quality of Service

QoS on Nexus 5000 & UCS

QoS on Nexus 5000

- + Nexus runs best effort queueing by default
 - + E.g. all traffic is CoS 0
- + For Storage (FC/FCoE) Nexus must provide lossless fabric
 - + Switches use new Data Center Bridging (DCB) protocols to negotiate lossless fabric
 - + FCoE policy should be automatically applied when “feature fcoe” is enabled
- + Nexus 5k uses Virtual Output Queues to avoid congestion

How do VOQs work?

- + Ingress packets are ingress queued in the VOQ
- + Egress port signals ingress port that congestion has occurred
- + Traffic is queued at all ingress ports sending to the congested port
- + Traffic to uncongested ports travels freely
- + VOQ is handled in hardware by the “central arbiter” (an ASIC)
- + VOQs prevent “Head of Line Blocking” (HoL)



<https://t.me/learningnets>



Data Center Quality of Service

Control Plane Policing (CoPP)

What is CoPP?

- + Control Plane Policing
- + Protects the control plane and separates it from the data plane
- + Used to prevent DoS attacks against the switch itself
- + Supervisor separates traffic into three categories
 - + Control Plane
 - + Routing protocol operations, e.g. OSPF
 - + Data Plane
 - + Regular data forwarding, e.g. web client to server
 - + Management Plane
 - + Managing the platform, e.g. CLI access

Default CoPP Settings

- + Nexus 7000 admin VDC runs strict CoPP Profile by default
 - + **copp profile strict**
- + Verification through **show policy-map interface control-plane**
- + Essentially automatically accounts for all possible control plane protocols
 - + E.g. OSPF, OTV, MPLS, FabricPath, etc.
- + Can be modified but not recommended
 - + Either you forgot to account for something, or policed something too strictly
 - + Can result in dropped legitimate control plane protocols



<https://t.me/learningnets>