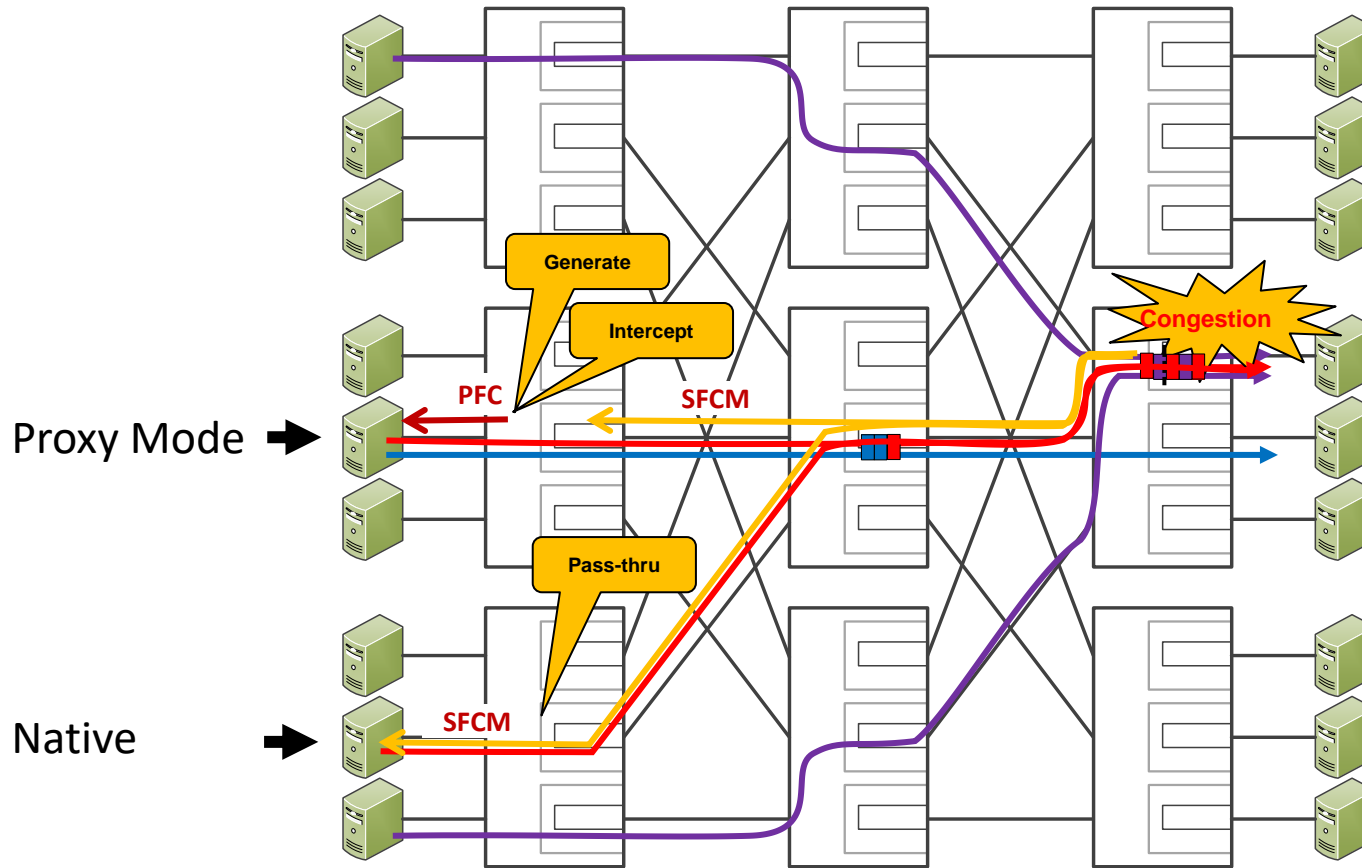


DCBX for SFC

Paul Congdon

Two operating modes for SFC

P802.1Qdw - Source Flow Control



Native

- End-station supports SFC
- SFCM passes through ToR to end-station

Proxy Mode

- End-station supports PFC, but not SFC
- SFCM intercepted by ToR and converted to PFC

How to automate the configuration of these two operating modes?

DCBX

- Data Center Bridging eXchange Protocol (DCBX)
- Operates over LLDP
- Goals
 - Discover capabilities
 - Detect miss-configuration
 - Automate configuration
- Current dcbxSet of TLVs
 - NOTE: to support a ‘set’ of organizationally specific TLVs you support the entire ‘set’

Table D-1—IEEE 802.1 Organizationally Specific TLVs

IEEE 802.1 subtype	TLV name	TLV set name	TLV reference	Feature clause reference
01	Port VLAN ID	basicSet	D.2.1	6.9
02	Port And Protocol VLAN ID	basicSet	D.2.2	6.12
03	VLAN Name	basicSet	D.2.3	12.10.2.1.3
04	Protocol Identity	basicSet	D.2.4	D.2.4
05	VID Usage Digest	basicSet	D.2.5	D.2.5
06	Management VID	basicSet	D.2.6	D.2.6
07	Link Aggregation TLV	basicSet	IEEE Std 802.1AX	IEEE Std 802.1AX
08	Congestion Notification	cnSet	D.2.7	Clause 33
09	ETS Configuration TLV	dcbxSet	D.2.8	Clause 38
0A	ETS Recommendation TLV	dcbxSet	D.2.9	Clause 38
0B	Priority-based Flow Control Configuration TLV	dcbxSet	D.2.10	Clause 38
0C	Application Priority TLV	dcbxSet	D.2.11	Clause 38
0D	EVB TLV	evbSet	D.2.12	D.2.12
0E	CDCP TLV	evbSet	D.2.13	D.2.13
10	Application VLAN TLV	dcbxSet	D.2.14	Clause 38

Proposed use of DCBX for SFC

- Advertise the capability of supporting SFC
- Advertise which TCs are enabled for SFC
- If end-station does NOT support SFC, but supports PFC – enable Proxy
- If end-station supports SFC – enable Native
- Consider allowing enabled TCs to be configured by peer

How to incorporate into DCBX?

1. Modify an existing TLV in the dcbxSet (e.g. PFC configuration)
2. Define a new TLV for the dcbxSet

Option 1: Update PFC TLV




Figure D-10—Priority-based Flow Control Configuration TLV format

- Willing – indicates if sender is willing to accept configuration from remote
- PFC cap – number of TCs that can simultaneously support PFC
- PFC Enable – bit mask of TCs that have PFC enabled
- **NOTE:** P802.1Qdt plans to add a fixed amount to this TLV – this fixed amount would need to be present and have the ability to be specified as N/A if Qdt is not implemented or not supported.
- SFC cap – number of TCs that can simultaneously support SFC
- SFC Enable – bit mask of TCs that have SFC enabled
- SFC Willing?
- **NOTE:** Backward compatibility rules for TLVs are specified in 802.1AB Clause 6.6.1

Option 2: New SFC TLV

TLV type = 127	TLV information string length = 6	802.1 OUI 00-80-C2	802.1 subtype = ??	Willing	Re- served	SFC cap	SFC Enable
7 bits	9 bits	3 octets	1 octet	1 bit	3 bits	4 bits	1 octet

- SFC cap – number of TCs that can simultaneously support SFC
- SFC Enable – bit mask of TCs that have SFC enabled
- SFC Willing?

Pros/Cons

- Option 1 – Update PFC TLV
 - Pro: Saves space – existing TLV, backward compatible rules support additions
 - Pro: required for DCBX and PFC support anyway – part of dcbxSet
 - Con: Qdt fields must be present and puts requirement on their definition
- Option 2 – Dedicated new SFC TLV
 - Pro: Could be independent of other DCBX TLVs (e.g. make it its own set)
 - Pro: No risk in backward compatibility
 - Con: Yet-another-TLV with more overhead and new data (6 bytes of overhead for 2 bytes of information)
 - Con: must correlate information from multiple TLVs (PFC TLV and SFC TLV)

Preference: Option 1