



P802.1Qau issues

Issues to be settled for P802.1Qau

Version 2

Norman Finn

Cisco Systems

P802.1Qau issues

- Questions in black
 - Answers in Red
-
- This slide deck documents the decisions made with regard to the preparation of D1.0 of P802.1Qau.
 - The terminology of this presentation does not match the terminology to be incorporated into P802.1Qau/D1.0.

Document Issues

- (p. x) Exactly how do CM-aware Bridges influence the selection of the active topology (MSTP, RSTP, P802.1aq, etc.)?
 - Make capability match part of the priority vector?
 - CM capable or CM enabled?
 - Add cost to links with unlike capabilities?
 - (802.1Qat/Qav have the same problem.)
- IF we do it at all, BPDUs carry a capabilities bit vector (1 bit for CM enabled).
 - Increase the port cost if capabilities differ.
 - But, it's not clear that this needs to affect MSTP, because this is always a managed network, so not for D1.0.
- LLDP could be useful to compare configurations (wholly or via a checksum like MSTP) to put defenses between bridges with unlike configurations.

Document Issues

- How do the stations know what priority values are congestion managed, and for what applications' use each is intended?
 - By configuration?
 - Through LLDP?
- No answers, yet.

Document Issues

- (p. 24) Should BCNs generated by stations be routed through the stations' output priority queues?
- Yes.

Document Issues

- (p. 28) Format of BCN (1):

Destination address is from sampled frame, of course

Source address: Could be CP, could be any MAC address belonging to the transmitting bridge.

Version, Fb: Obviously required.

Is size of Fb fixed? 6 bits? 8 bits?

CPID? Port number? (This is *not* required for the protocol, but for possible debugging purposes. Trade this against the fact that this frame may be generated by hardware.)

- CPID + Source address == unique identifier of CP, which is a queue on a port.
- Size of Fb is not known, yet. Also send Qoffset & Qdelta in the PDU.

Document Issues

- (p. 28) Format of BCN (2):
 - Returned frame is not necessarily available to the Congestion Point in on-the-wire format because:
 - Destination address, source address, priority, CFI, VLAN identifier, are parameters, not embedded in a frame.
 - *We can* choose to put them down in 802.3 format.
 - Is source address of encapsulated frame necessary? (It's the same as BCN's destination.) .3 format or .5 format?
- **Use fixed fields for parameters. Perhaps remove source address.**

Document Issues

- (p. 28) Format of BCN (3):
 - VLAN ID can only be inserted as a local-format tag, even though station's link has no VLAN ID, or has a different format (SNAP vs. EtherType). Isolated VLAN ID field might be better.
 - At least 8 bits of encapsulated frame length are required, because not all media have a 64-byte minimum length.
 - How many additional octets of frame are returned?
- Use isolated VLAN field. No length.
- Data PDU must be delivered at known place in the BCN.
- Need appendix pointing out that a CM-enabled TPMR will return a nonsense encapsulated frame in BCN.

Document Issues

- Link aggregation between RP and network:
 - Link aggregation queues are a general problem: Are queues placed above link aggregation, below link aggregation, or in both places?
 - Resource reservation (.1Qat, .1Qav) has issues with matching the reservation and queue configuration to the link that the data will use.
 - Congestion management (.1Qau, .1Qaz) has issues with matching BCNs to the flow queues on physical links.
 - The Bridge model places queues much higher in the stack, above Link Aggregation. This solves the both problems in theory, if not in fact.
- In P802.1Qau, there is no problem with Link Aggregation. Add editor's notes.

Algorithm Issues

- Which is more up-to-date? Pseudocode (au-rong-qcn-serial-hai-pseudo-code rev2.0.pdf) or descriptive slides (au-pan-qcn-benchmark-sims-0108.pdf)?
- **Descriptive slides are more up-to-date.**

Algorithm Issues

- Are RPs going to send probes?
- Not in P802.1Qau/D1.0.

Algorithm Issues

- Does a station assign every frame with a congestion-managed priority value to a rate limiter (flow queue)? Or, does it construct rate limiters when needed? Can you tell, from outside the station?
- Assume that rate limiters are created (destroyed) as (not) needed.

Which parameters are manageable?

	Constant in network?	Manageable (default)	Computed from
w			
W_EQ			
Fb quant.			
bytes_per_BCN			
bytes_per_BCN_jitter			
Gd			
R_AI			
R_HAI			
BC_LIMIT			
TIMER_PERIOD			
FAST_RECOVERY_TH			
MIN_DEC_FACTOR			
MIN_RATE			