Class of Service in 802.1

Norman Finn

Cisco Systems, Inc.

CoS Bridging Example

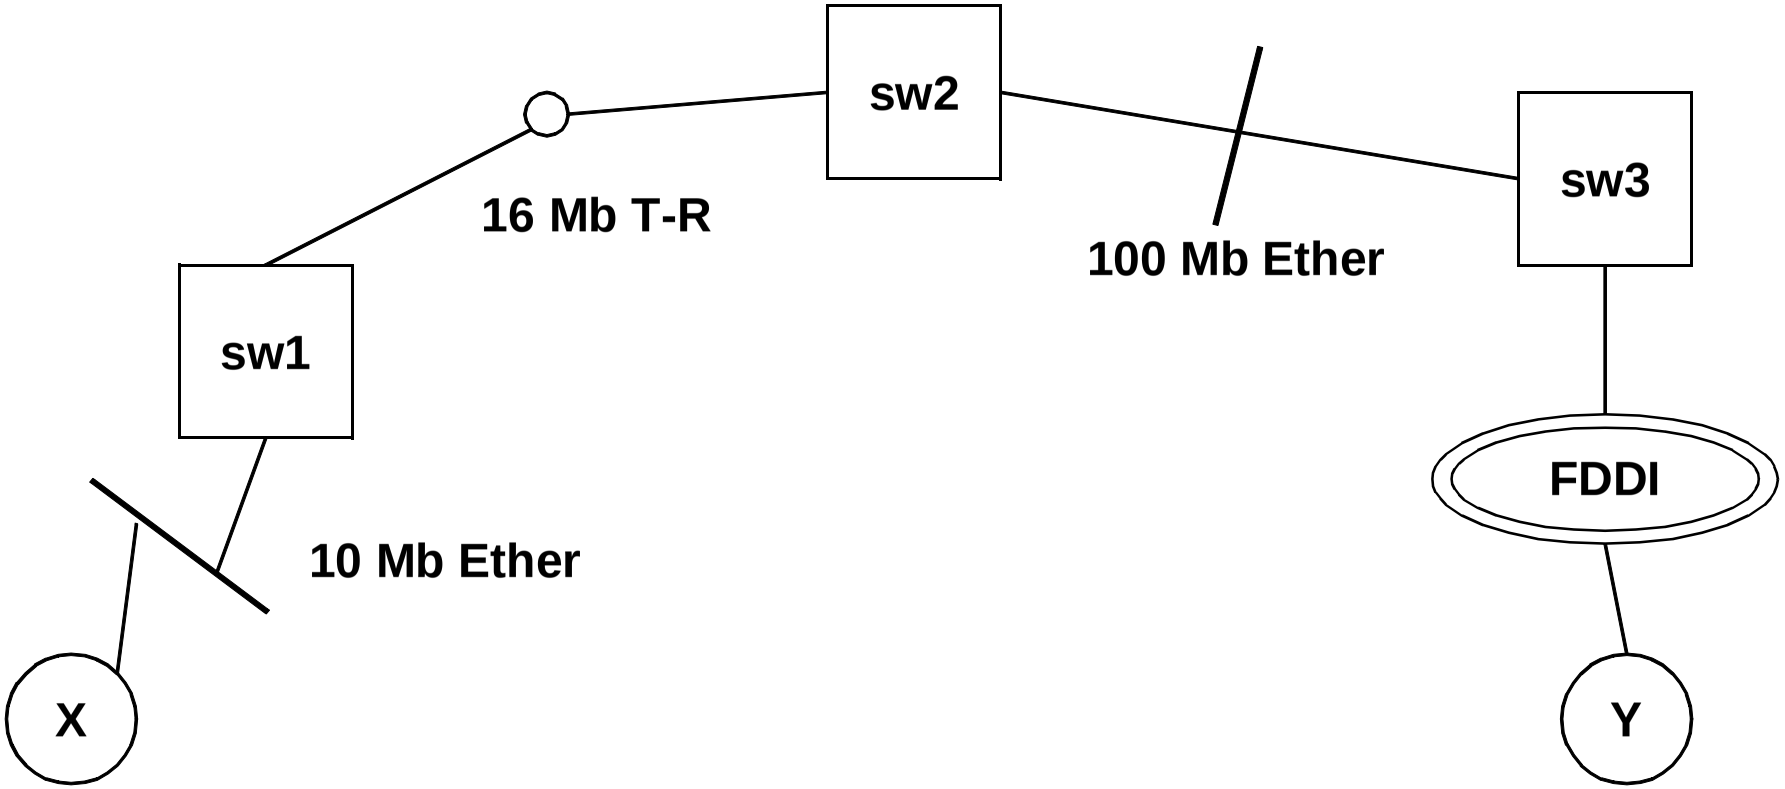

- **100 Mb and 10 Mb Ethernet prevent passing access priority, even if switches support it.**
- **An end-to-end CoS indication corrects this problem.**

Methods for End-to-end CoS Indication

At present, there are four methods on the table for an end-to-end CoS indication:

- **Use the high-order 3 bits of the VLAN-ID field.**
- **Using the GARP protocol, assign each multicast MAC address a 3-bit CoS value on a per-LAN segment basis.**
- **Redefine the Universal/Local bit in the destination MAC address as a 1-bit CoS indication.**
- **Create a protocol to assign blocks of locally-administered unicast addresses to end stations in such a manner that the CoS indication is in a fixed location in the MAC address.**

Summary of End-to-End CoS Solutions

	multi/ unicast	CoS levels	Frame length	bridge has CoS in:	CoS-unaware bridges	CoS-unaware end stations
high bits of VLAN-ID	multicast/ unicast	8	Adds 4 bytes	frame	transparent[a]	conflict avoided
assignment via GARP	multicast only	8	Same size	fwd table	transparent	conflict avoided
redefine U/L bit	multicast/ unicast[b]	2	Same size	frame or fwd table	transparent[c]	possible conflict[d]
blocks of local addrs	unicast only	many	Same size	frame or fwd table	transparent[c]	conflict[d]

a. If bridge does not parse above layer 2, and if frame length is no problem.

b. Subject to incompatibilities with IGMP-aware bridges.

c. Subject to issues with bridges learning all of a station's MAC addresses, and with "smart" but CoS-unaware bridges.

d. May conflict with existing protocols for allocating locally-administered MAC addresses.

Blocks Of Local Addresses

The “blocks of local addresses” should be ruled out:

- **No allocation algorithm exists**
- **Creating such an algorithm involves numerous problems:**
  - **—It would be difficult to add yet another protocol to a diskless end station’s (or bridge’s or router’s) boot PROMs.**
  - **—It would be difficult to change an end station’s MAC address if it used its L3 address with its own, universally allocated MAC address, before acquiring its locally allocated address.**
  - **—It is difficult to ensure that the same end station gets the same block of MAC addresses each time it boots, and difficult for some L3 protocols (IP, for example) to adjust to changed MAC addresses.**

Assign Multicast CoS with GARP

- **Applies only to multicast addresses.**
- **Why use two methods (one for multicast, one for unicast) if one will work?**
- **Variability of CoS with LAN segment:**
  - **—Partially attainable by varying the interpretation of embedded CoS on a per output port basis.**
  - **—Is it really necessary in the bridged, as opposed to the routed, environment?**

Use High-Order Bits of VLAN-ID for CoS

- **No adverse interactions between CoS and addressing.**
- **8 levels of CoS supported.**
- **End station address recognition not affected.**
- **Bridges' forwarding table size not affected.**
- **Either adds 4 bytes to a packet, or reduces MTU by 4 bytes. (802.1q shares this problem.)**
- **"Smart" bridges may be unable to parse past tag. (802.1q shares this problem.)**

Use High-Order Bits of VLAN-ID for CoS

- **No PAR changes required:**
  - **—802.1p defines:**

Ethertype(s) (same as 802.1q)	CoS	Undefined
16	3	13

  - **—802.1q defines:**

Ethertype(s) (same as 802.1p)	Undefined	VLAN-ID (0 = none)
16	3	13

Redefine U/L Bit for Unicast

- **Only provides 2 levels of CoS.**
- **Changes current meaning of U/L bit in a way that may be incompatible with existing uses.**
- **Requires solution to allow old bridges to learn dual MAC addresses:**
  - **—Random selection of source U/L bit? *Unreliable.***
  - **—Select source U/L bit by destination U/L bit used by each correspondent? *To complex for end station.***
  - **—Send broadcasts with both settings of U/L bit in source? *Unnecessary traffic. Old bridges' forwarding tables must grow to hold unused addresses.***
  - **—An answer for the source MAC problem would be required.**
- **"Smart" bridges may have difficulty with dual MAC addresses, e.g. existing bridges that prune IP multicast trees.**
- **Does not add to size of packet.**

Underlying Mechanisms

	CoS in bridge forwarding table	CoS explicit in frame
high bits of VLAN-ID	No	Yes
assignment via GARP	Yes	No
redefine U/L bit	Yes	Yes
blocks of local addrs	Yes	Yes

- To support CoS assignment via GARP, the bridge must include the CoS in its forwarding table.
- To support CoS in the VLAN-ID, the bridge must extract the CoS from the packet, itself, as already done in 802.5 and FDDI bridges.
- It is difficult to support both underlying mechanisms; we must not require both mechanisms to achieve interoperability.

Pick Just One

- **A bridge or end station, using only the mechanisms in 802.1p, must be interoperable with other such bridges and end stations, and fully support higher layer CoS requirements.**
- **The applications that drive CoS will use standard end-to-end layer 3 protocols (RSVP) and L3-to-L3 protocols, (ATM signaling, ISSLL). Applications that need CoS will use these standard mechanisms.**
- **The primary purpose of 802.1p should be to support (and influence) these standards.**
- **Two methods for end-to-end CoS leads to bridging nightmares:**
  - **—Translation between tagging methods.**
  - **—Per-end-station knowledge of tagging method (as for IP ARP and 802.3/Ethernet encapsulation).**
  - **—More than twice as much work!**
- **“Implicit” bridges must interoperate with “explicit” bridges to deliver the capabilities of 802.1p and higher-layer CoS standards.**

Conclusion

- **No solution is perfect.**
- **There should be only one standard solution.**
- **Using the high bits if the VLAN-ID provides the safest, cleanest, and simplest solution to the problem of End-to-end CoS indication.**
- **We should not require the forwarding data base to contain CoS in order to achieve interoperability with “implicit” bridges.**