



Real-time networks and preemption

More to it than latency

Rev. 5

Norman Finn

nfinn@cisco.com

<http://www.ieee802.org/1/files/public/docs2012/new-avb-nfinn-real-time-networks-1111-v05.pdf>

What is a real-time network?

- In a real sense, all networks are “real-time” except for simulations of networks.
- Video or voice data is certainly a kind of “real-time”
- Priority, resource reservation, and other methods work for many networks that have tight latency and/or jitter requirements.
- In this slide deck, “real-time” means a guaranteed response time to any given input or combination of inputs. **No excuses, no exceptions.**
- Typical examples are automatic automobile braking systems and robot control.

The goal: Three networks in one

- Three levels of service: **Critical**, **Reserved**, and **Best-Effort**.
- **Critical** traffic uses preemption, time-gated queues, and maybe cut-through forwarding, so that other classes do not disturb it.
- **Critical** traffic uses also uses time synchronized transmissions to ensure that 1) critical flows do not interfere with each other, and 2) critical flows do not overly disrupt Reserved traffic.
- **Reserved** traffic uses bandwidth reservation and shaping to guarantee audio/video requirements.
- **Best-effort** traffic gets what's left.


What do real-time networks lack?

- Some excellent presentations have been made this year on requirements from users and designers of real-time automotive and industrial networks.
- There are common threads that we can address:
 - Topology**
 - Delivery**
 - Predictability**
- But, we cannot address them in isolation, either from each other, or from more general uses of Ethernet networks.

Topology

Topology

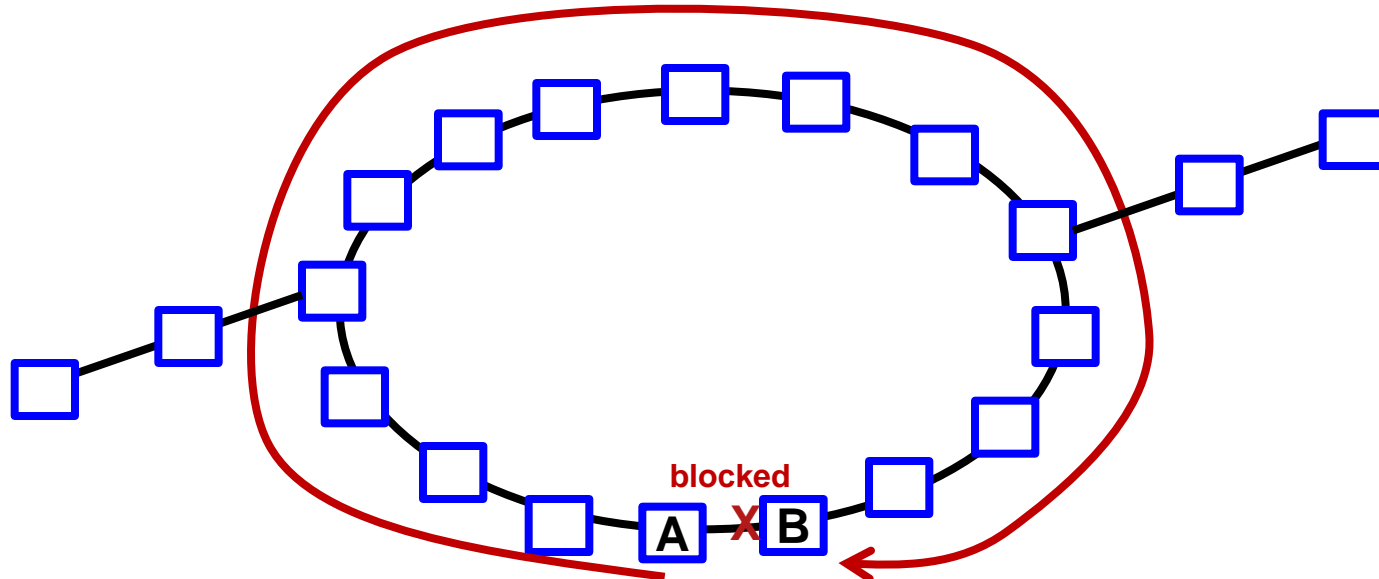
- As has been known for a long time, spanning tree has issues in simple networks with links of widely disparate data rates.



- This diagram illustrates the problem in the home.

Topology

- Similarly, large rings are the least-favored topology for spanning tree.
 - Rings (with tails) exhibit the worst case reconfiguration times.
 - Rings exhibit the worst case penalty for blocking a link.



Topology

- We could build on spanning tree. But ...
 - Bridges running MSTP lack a view of the whole network, and this may useful information to applications.
 - Using MSTP requires that MSRP or similar protocols must converge *after* MSTP converges, instead of simultaneously.
- For these reasons, and because the blocked-link problems in the previous slides are solved, this author believes that a **link-state protocol should be the basis for real-time networks.**

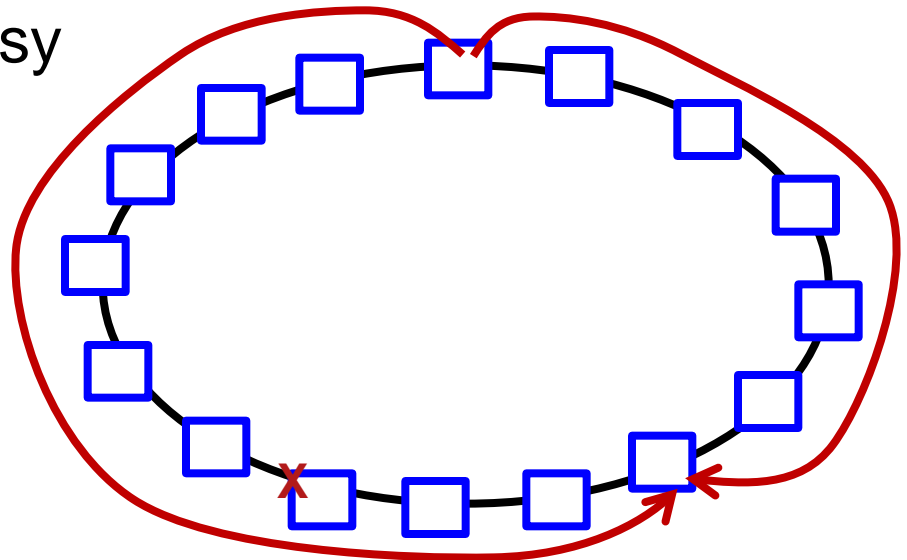
Shortest Path Bridging

- Coincidentally, SPBV (VLAN-mode Shortest Path Bridging) can be made plug-and-play for networks in the size range we're interested in.
- Some work would still be needed:
 - We must balance the number of VLANs against number of bridges ($[\text{number of bridges}] * [\text{number of VLANs}] < 4096$).
 - Learning MAC address can preclude the use of two paths between two stations.
- It is true that SPBV is more complex than alternatives that are based on a fixed topology. But, one must ask whether the topology is *really* fixed.

Delivery

Delivery

- For ultra-reliable communications between consenting stations, delivery of frames along two paths would be very helpful, and there are documented methods for it.
- This cannot be easily done by current bridging/routing protocols: paths are not equal cost, overriding the topology to slip past blocked links breaks address learning, and it is not easy to discovery maximally-disparate paths.
- But, if we can do it, the value will be significant!



Delivery

- It is worth pointing out that P802.1Qbf Segment Protection can route frames **outside** the spanning tree or SBP framework, including simultaneous delivery along multiple paths.
- We can observe that SPB could be used to distribute the paths from the ingress switch to other switches.
- SPB can also greatly increase the efficiency of the Stream Reservation Protocol.

Predictability

Time synchronization

- There is a long history of real-time networking, especially in the aerospace industry.
- In this world, “real time” does not mean interrupts and preemptive process scheduling. It does not mean “best effort delivery.”
- “Real-time” means **scheduling**: scheduling processes within a station, scheduling communications between stations, and coordinating the stations’ schedules.
- Scheduling guarantees that all processing and communications happen within the required time limits.
- Even network recovery is accounted in scheduling alternatives.

Predictability


- The real-time network scheduling model is:
communicate, **compute**, **communicate**, **compute**, ...



- Communication can be **concentrated** into a small window, in order to leave compute time unhindered by interruptions.
- This concentration:
 1. Is essential for the critical applications to work.
 2. Is essential to enable the bandwidth reserved applications.

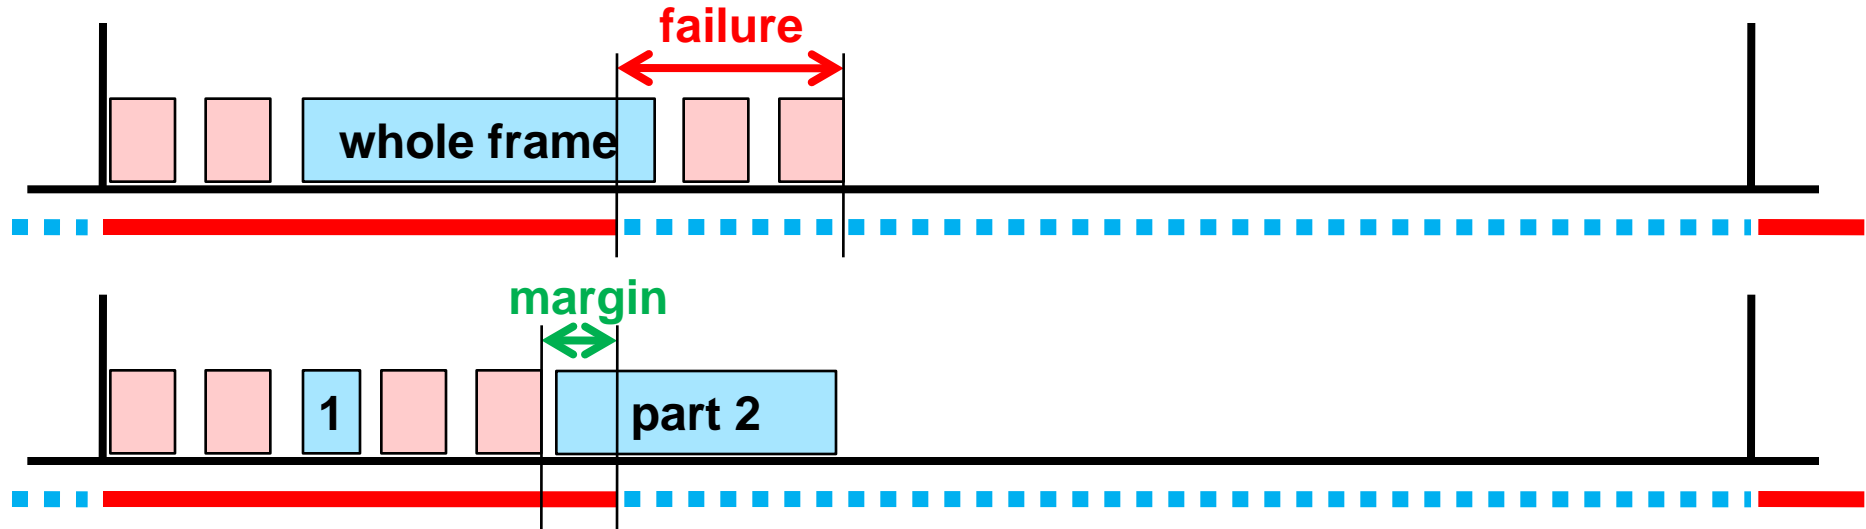
Although...

- Critical traffic must live with bandwidth reserved traffic, also.
- If scheduled critical traffic takes enough percentage of the bandwidth for a long enough time, it will starve the bandwidth reserved (audio or video) traffic.




- Critical/reserved requirements **could be incompatible**.
- **Applications developers understand this.**

Concentration solely via preemption



- **Small gaps** inevitably occur between critical frames because they take different paths through the network.
- Preemption prevents large non-critical frames from acting as a wedge to stretch out the critical communications period.
- Queuing delayed critical frames helps to prevent excessive wedge insertion.

Concentration via time-gated queues



- **Time-gated queues** in the bridges can be used to prohibit any but critical frames in the critical windows.
- But, the critical gate must be **extended ahead** of the transmission point **T** in order to prevent long frames from delaying the start of the critical data.
- **Preemption** eliminates the need for excessively-long pre-**T** extensions, which would disrupt reserved traffic.

Predictability

- It is true that preemption reduces queue size, and thus latency, by **only one frame**.
- But, that one frame makes a **big difference** when concentrating the critical traffic, leaving room for the both computing by critical applications, and bandwidth for reserved traffic.

Summary

Real-time networks: 3 networks in 1

- **Scheduling** of application transmissions is required, both to meet application requirements and to avoid disrupting bandwidth reserved traffic.
- **Preemption** and **time-gated queues** are required to prevent interference between critical traffic and non-critical traffic.
- **Cut-through forwarding** (of critical traffic only) **may** be needed to minimize latency.
- Existing **bandwidth reservation** and **shaping** are required to meet audio / video requirements.
- Existing **priorities** support best-effort service.
- **Shortest Path Bridging** offers a convenient means for implementing all of the above.