AV Bridging and Ethernet AV™

Michael Johas Teener

mikejt@broadcom.com

Agenda

- IEEE 802.1 summary
- What is AV bridging?
  - and Ethernet AV?
- Why is it needed?
- Where will it be used?
- How does it work?

IEEE 802.1

- Responsible for “higher level” services within 802 networks
  - architecture
  - “interworking” - bridging between different 802 (and other) networks
  - link level security
  - management interface for layer 2 facilities
- Examples
  - 802.1D/Q bridging (usually called “switches”)
  - 802.1X, 802.1AE MAC security
  - 802.1ad Provider Bridges

New Work in 802.1

- Audio Video Bridging Task Group
  - provides the specifications that will allow time-synchronized low latency streaming services through 802 networks
  - http://www.ieee802.org/1/pages/avbridges.html
- Congestion Management Study Group
  - (Pat will be talking about this)

What is AV bridging?

- Precise timing in an 802 network
- E2E bandwidth and latency guarantees using QoS and Admission Control
- Do both with very low cost adder (approaching zero) to meet CE market requirements.
- Ethernet AV is the first wave
  … but not the last

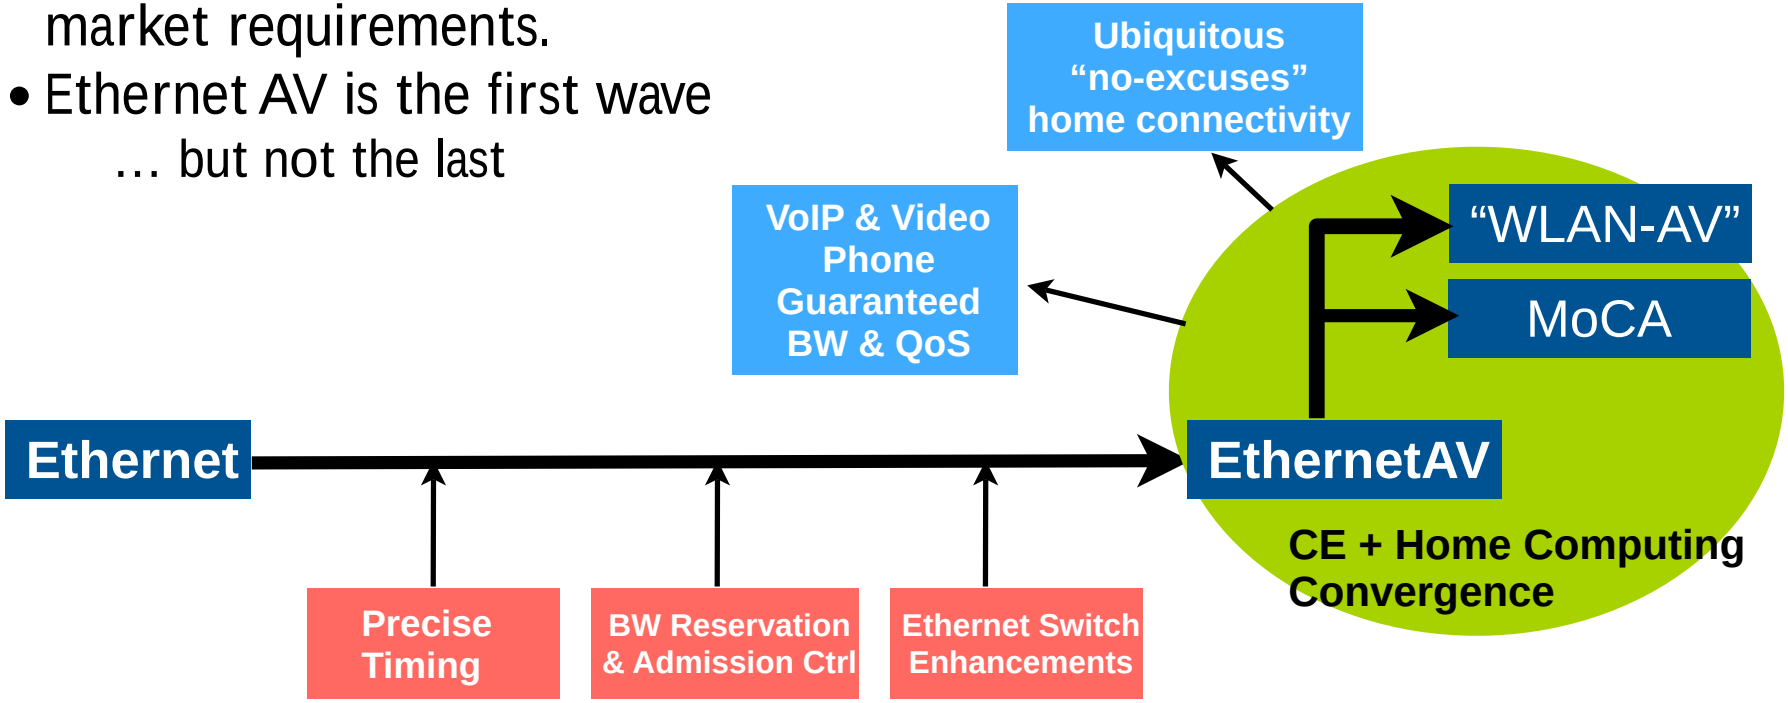

Performance goals

- Determined during the one-year life of the "Residential Ethernet Study Group"
  - http://www.ieee802.org/3/re_study/
- Timing synchronization to sub-microsecond
  - coordinating multiple audio devices (speakers and microphones) for proper phasing
- Clock jitter appropriate for applications
  - 100ns for "free", can be filtered down to 100ps
  - MTIE masks for most applications can be met
- Latency of no more than 2ms in a 7 hop Ethernet
  - viewed as worst case for most applications

Chaos in layer 2 steaming QoS

- "Isochronous" or "streaming" services required by some markets
  - Consumer electronics, professional A/V, telecom, and instrumentation in particular
- Always done in isolation with little thought of bridging to alternates
  - IEEE 1394 ≠ 802.11e ≠ HomePlug AV ≠ MoCA ≠ USB
  - good reasons (vastly different access protocols)
  - bad reasons (NIH)

Digital Home Media Distribution

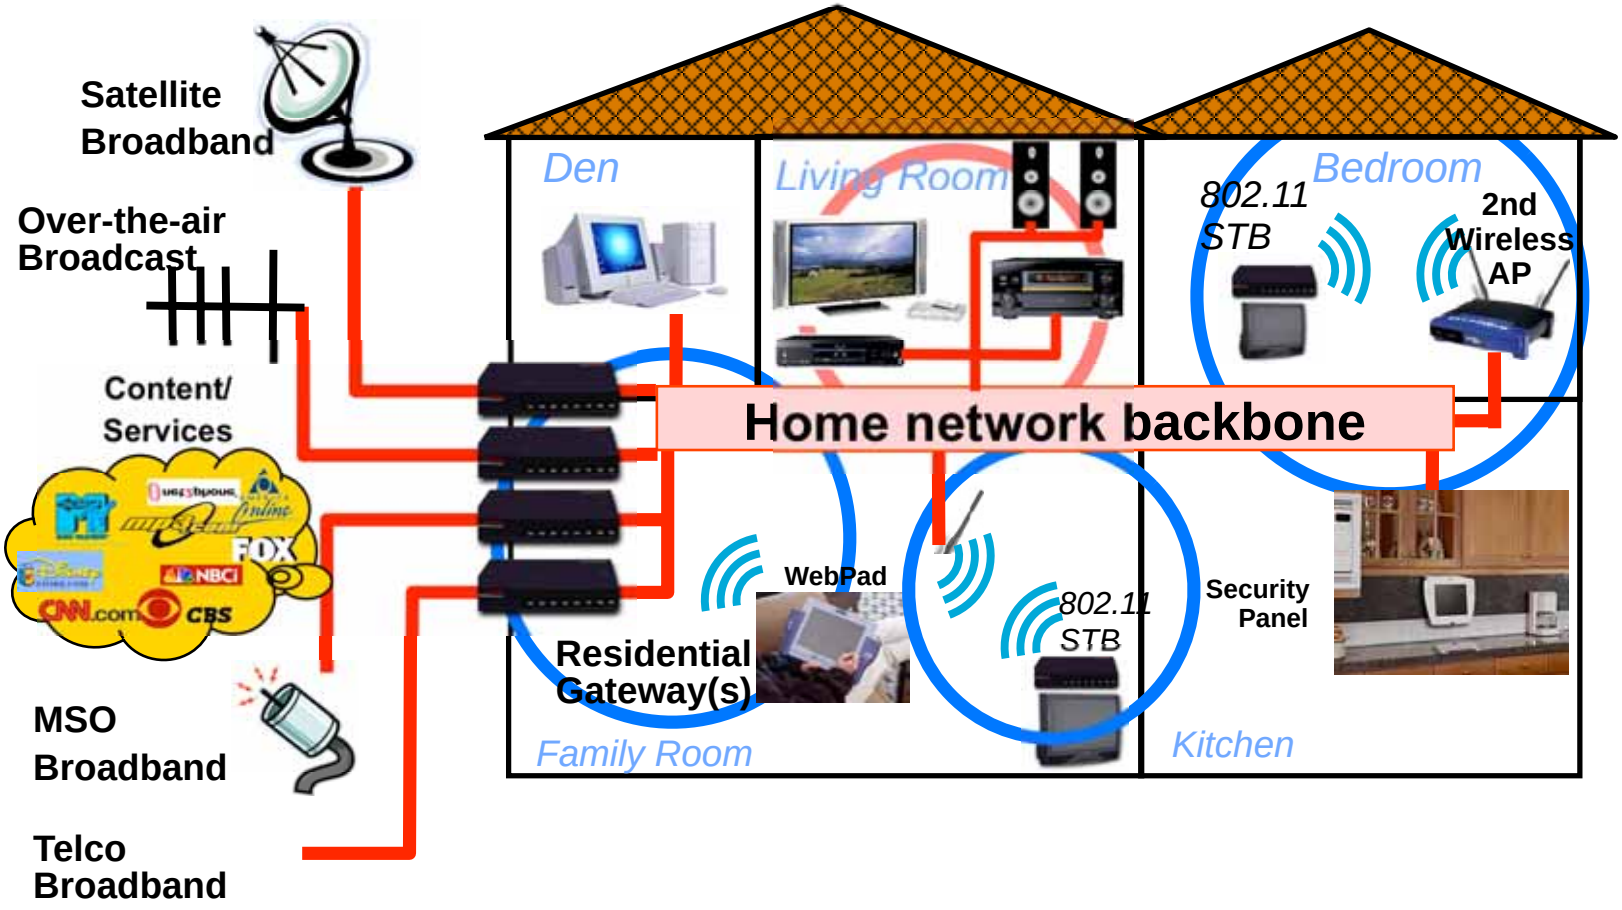

Unified Layer 2 QoS

- Enhance network bridging
  - Define common QoS services and mapping between different layer 2 technologies
    - E.g., 802.3 Ethernet, 802.11 WiFi, UWB, MoCA, etc
  - IEEE 802.1 is the common technology
    - Ethernet “switches” (IEEE 802.1D/Q bridges)
    - Basis of 802.11 A/P attachment to Ethernet
    - Basis of non-802 network bridging (e.g., FDDI, carrier nets)
- Common endpoint interface for QoS
  - “API” for QoS-related services for ALL layer 2 technologies
  - Toolkit for higher layers

The first step: Ethernet AV™

- Simple enhancement to IEEE 802.1 bridges to support streaming QoS
  - 2 ms guaranteed latency through 7 Ethernet bridges
  - Admission controls (reservations) for guaranteed bandwidth
  - Precise timing and synchronization services for timestamps and media coordination
    - < 1µs absolute synchronization between devices
    - jitter less than 100ns, filterable down to 100ps
- Change to IEEE 802.3 specs (but not implementations)
  - the actual time that a particular frame is sent or received
    - maybe use existing layer management services

Proposed architecture

- Changes to both IEEE 802.1Q and layer 2
  - 802.1Q - bridges/switches - most of work
  - 802.3 - Ethernet MAC/PHY - possible small change to MAC definition
  - 802.11n/HomeplugAV/MoCA - more work, but basic tools in place
    - Not discussed in this presentation
- Three basic additions to 802.1/802.3
  - Traffic shaping and prioritizing (p802.1Qav for 802.3 only),
  - Admission controls (p802.1Qat), and
  - Precise synchronization (p802.1AS)

Topology & connectivity

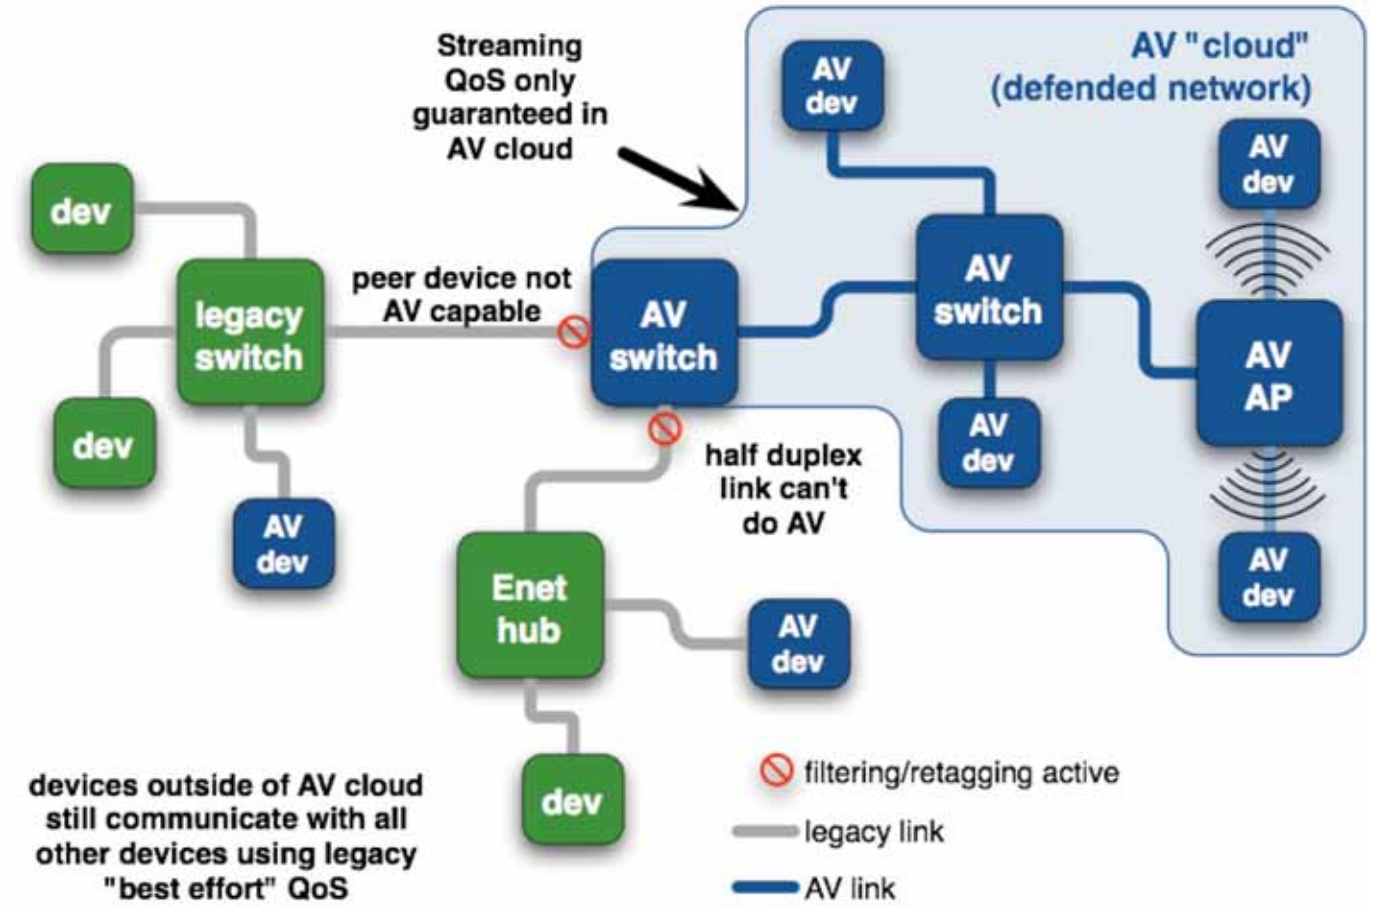

Changes to Ethernet Equipment

- MAC changes
  - Frame Timer – Accurately note time of RX/TX Ctrl Frame
- Queuing/DMA
  - Separate queues and DMA for class 4/5 frames to provide appropriate traffic shaping (scheduling)
- Admission Control (driver/bridge firmware)
  - Bandwidth allocation database associated with filtering database
  - Management using same methods (MRP) used for multicast addressing
- Real-time clock module
  - Master clock generator
  - Time Sync correction method

When?

- IEEE standardization process well under way
  - Originally an 802.3 study group, moved to 802.1 in November 2005 as "Audio/Video Bridging Task Group"
  - http://www.ieee802.org/1/pages/avbridges.html
  - Early drafts already available for 802.1AS and 802.1Qat
  - Expect technical closure in early 2007, WG ballots in mid-late 2007, IEEE standard in 2008
- First hardware/software soon after stablilization
  - Possibly a number of "pre standard" interations
- Will follow Ethernet-type product curve
  - 100M/1G/10G NIC/Switch all have markets for Ethernet AV

Interaction with IETF

- Unified services for L2 QoS and timing will be available
  - soon!
- IP-based services will likely be the most commonly used L3+ protocols on these new enhanced QoS L2
- 802.1 work is NOT finished
  - would like to make sure the new services have the right capabilities and interfaces

Specifics of interaction

(feel free to correct me, here)

- For transport:
  - RTP (precise timestamping might help using p802.1AS services, guaranteed BW/latency provided by p802.1Qav might also reduce the need for RTCP messages?)
- For QoS signaling:
  - The 802.1AS team is largely looking at the interactions with RSVP (as well as the non-IETF UPnP QoS protocols)
  - We just became aware of NSIS this week … perhaps this is the primary interaction point?
- For time synchronization:
  - Is NTP the right choice?

Thank you

Technical Details

Summary of protocols

Establishing the AV cloud

- IEEE Std 802.1AB defines "LLDP": Logical Link Discovery Protocol
  - Allows link peers to determine each other's characteristics
- Will be enhanced with P802.1AS service that gives a relatively precise round trip delay to a peer
  - Allows link peers to discover if any unmanaged bridges or other buffering devices are present on link

Traffic Shaping and Priorities
(p802.1Qav - rev to 802.1Q)

- Endpoints of Ethernet AV network must "shape traffic"
  - Schedule transmissions of streaming data to prevent bunching, which causes overloading of network resources (mainly switch buffers)
  - Shaping by limiting transmission to "x bytes in cycle n" where the cycle length is 125 µs or 1ms depending on traffic class
  - Traffic shaping in bridges will provide scalability
    - without it, all bridges need larger buffers
- Mapping between traffic class and priorities

Traffic Class?

- 802.1p introduced 8 different traffic classes
  - Highest (6 & 7) reserved for network management
    - low utilization, for emergencies
  - Next two for streaming (4 & 5)
  - Lowest four for “best effort”
- AV bridging:
  - Class 5 is for lowest latency streaming
    - Roughly 250 usec per bridge hop: interactive audio/video
  - Class 4 is for moderate latency streaming
    - Perhaps 1ms per bridge hop: voice over IP, movies

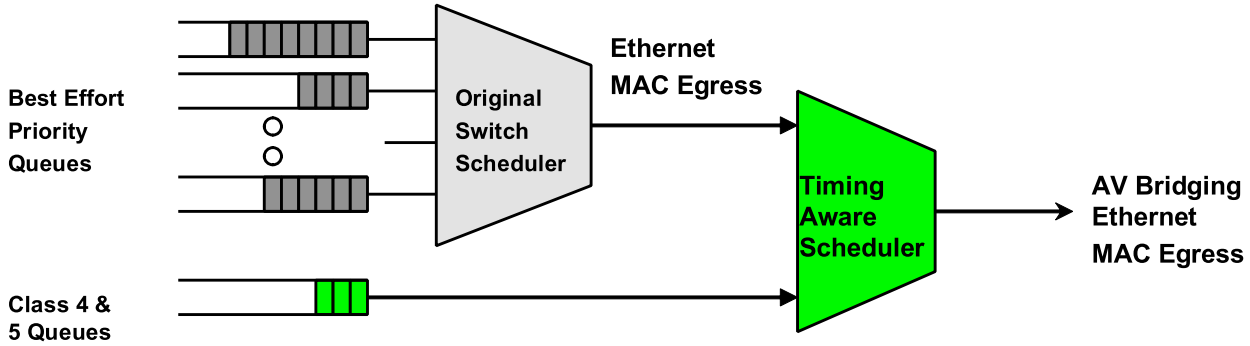

Admission controls

(p802.1Qat - added to 802.1Q)

- Streaming priority mechanism can reliably deliver data with a deterministic low latency and low jitter
  - but only if the network resources (bandwidth, in particular) are available along the entire path from the talker to the listener(s).
- For AV streams it is the listener's responsibility to guarantee the path is available and to reserve the resources.
- Done via a new 802.1ak "Multiple Registration Protocol" application: SRP ("Stream Registration Protocol")
  - Registers streams as multicast address/bandwidth/traffic class "tuples"
  - Perhaps other information useful for stream management such as path availability and a tag for higher layer mapping

Admission Control (I) (registration)

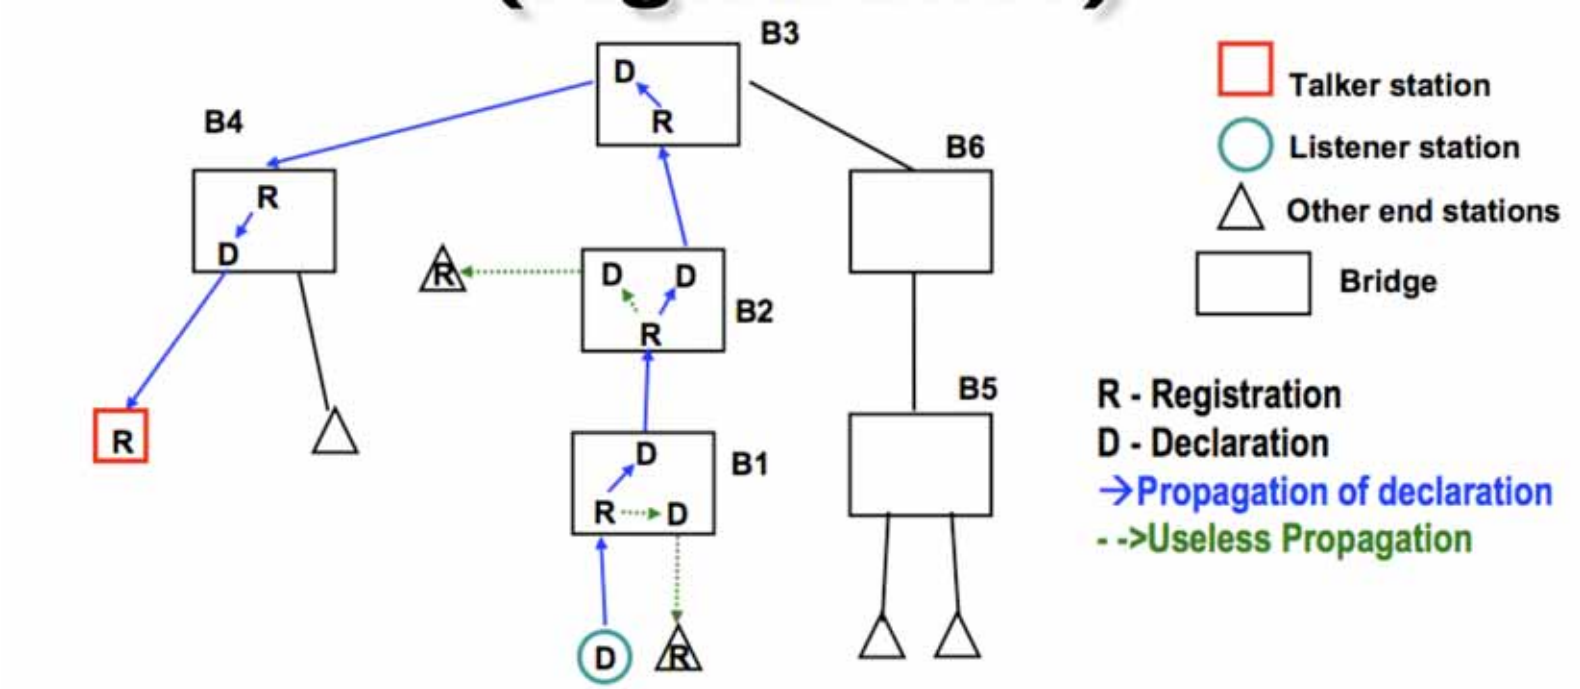

- With SRP registration, the talker and intermediate bridges know where are potential listeners and how to get to them
- Assume in the above figure, B3/B4 have learned the talker's address, and B1/B2 haven't, then:
  - MSRP floods the registration if the talker's address is not in the bridge FDB (eg. B1, B2)
  - MSRP relays the registration through specific outbound port if the talker's address is known by the bridge FDB (eg. B3, B4)

Admission Control (2)
(successful reservation)

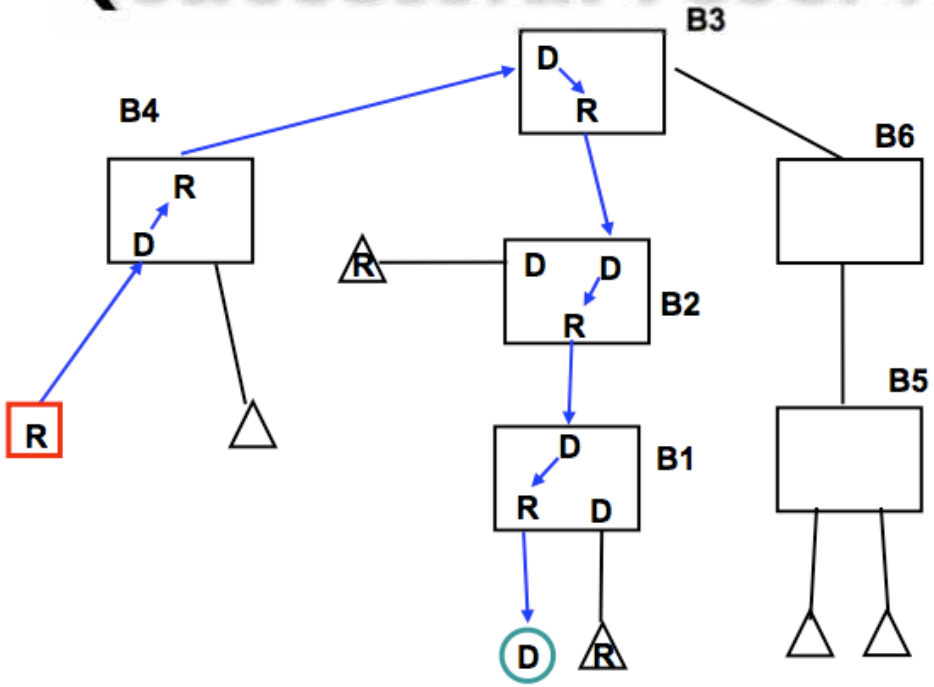

→Propagation of RESV

- RESV signaling triggers admission control operations in intermediate bridges. It also locks resources and updates isochronous filtering database if the admission control is successful.
- In this example, admission control is successful along the wholepath. RESV signaling servers as the end-to-end explicit ACK signaling to listener.

Admission Control (3)
(failed reservation)

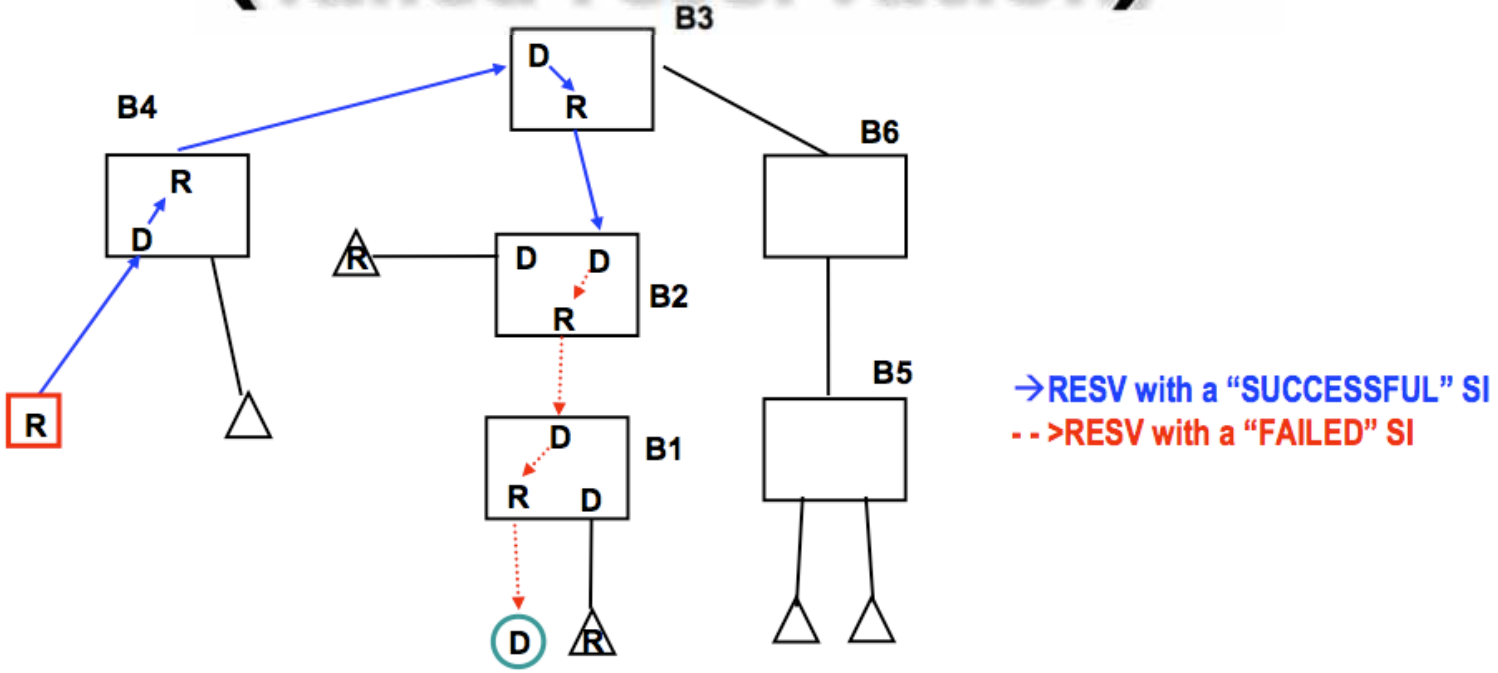

- In this example, admission control is failed at B2. The SI (Status Indication) bit of RESV signaling will be set to FAILED.
- The RESV is still forwarded to the listener. However, downstreambridges (i.e., B1, B2) will not lock resources for the RESV signaling whose SI is set to FAILED.
- Listener is noted of the failure since RESV with FAILED SI serves as an end-to-end explicit NACK

Precise synchronization (p802.1AS)

- AV devices will periodically exchange timing information
  - both devices synchronize their time-of-day clock very precisely
  - the delay time between devices is very precisely known
- This precise synchronization has two purposes:
  - to enable streaming traffic shaping and
  - provide a common time base for sampling data streams at a source device and presenting those streams at the destination device with the same relative timing
- Very similar to IEEE std 1588-2004, but much simpler
  - will be the “native IEEE 802 layer 2 profile” of new IEEE 1588v2

AV Grand Master clock

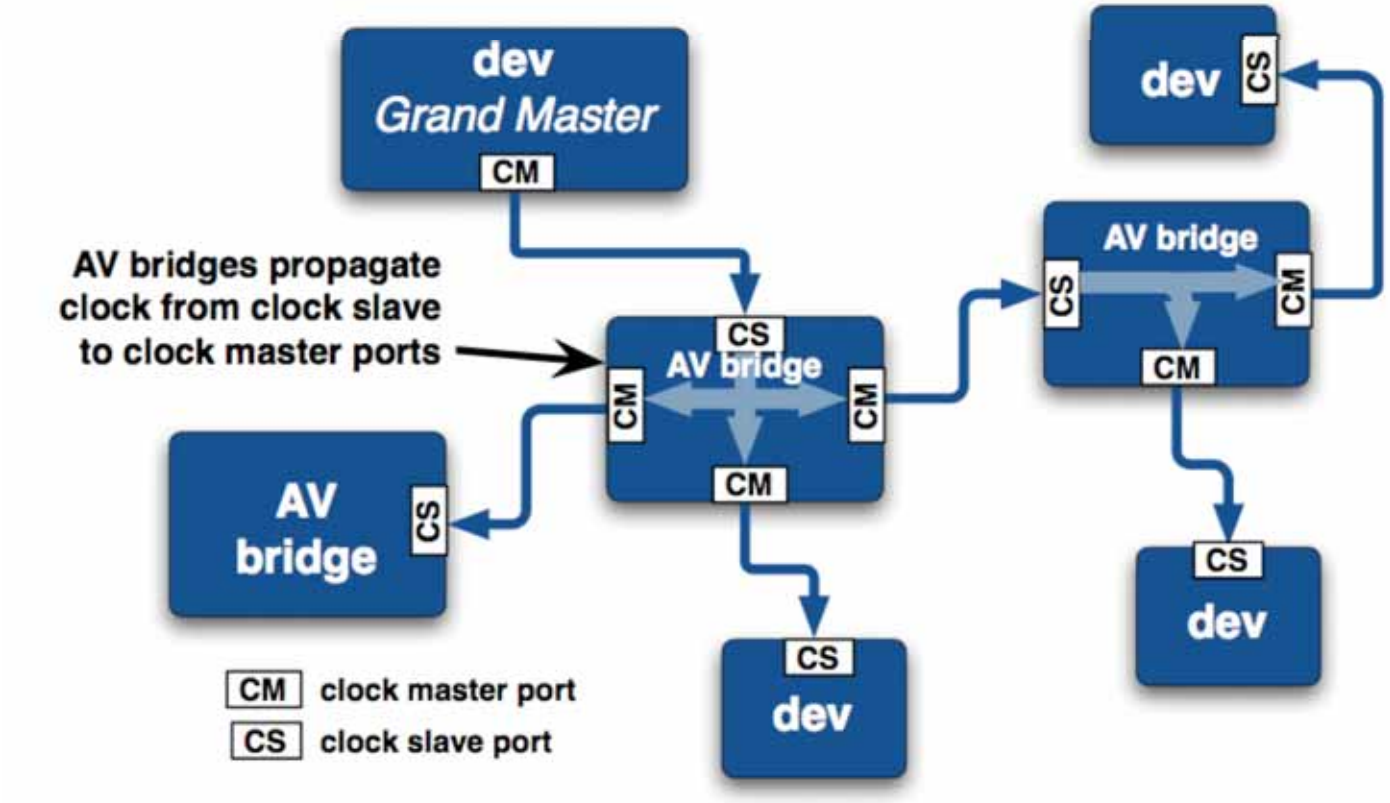

- There is a single device within an Ethernet AV cloud that provides a master timing signal.
  - All other devices synchronize their clocks with this master.

Master clock selection

- Selection of the master is largely arbitrary (all AV devices will be master-capable), but can be overridden if the network is used in an environment that already has a “house clock”.
  - Professional A/V studios
  - Homes with provider time-synchronization service
  - Carrier networks