Ethernet AV™ summary

Michael Johas Teener
(mikejt@broadcom.com)
April, 2006

Agenda

- What is Ethernet AV?
- Why is it needed?
- Where will it be used?
- How does it work?
- How will Broadcom support it?
- Beyond?

What is 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
- Trade group to provide trademark "enforcement" of otherwise optional features
  - Require useful bridge performance, network management, PoE management, auto-configuration features

Why is it needed? (1)

- Common IT-oriented networks have inadequate QoS controls
    - All use 802.1 “priority” (actually, “traffic class”)
- Ethernet is the best
    - … but it’s easy for the customer to misconfigure or overload
    - … no guarantees
- Wireless has inadequate bandwidth and excessive delays for whole-home coverage
    - … 802.11n and UWB work for non-critical applications, or short range
    - … latencies through multiple A/Ps may be too much for interactive applications
    - … no guarantees
    - … and we still need a backbone for the wireless attachment points

Why is it needed? (2)

- Proposed CE-based networks need new media or are expensive
  - MoCA requires coax everywhere, and is not cheap, and does not carry power, and has modest performance

    … but it's part of the solution
  - Power line is not cheap, has modest performance, is susceptible to interference, and is blocked by protection circuits

    … but it's part of the solution
  - 1394b/c long distance has limited developer base & infrastructure, is not cheap

    … but even this is part of the solution

Digital Home Media Distribution

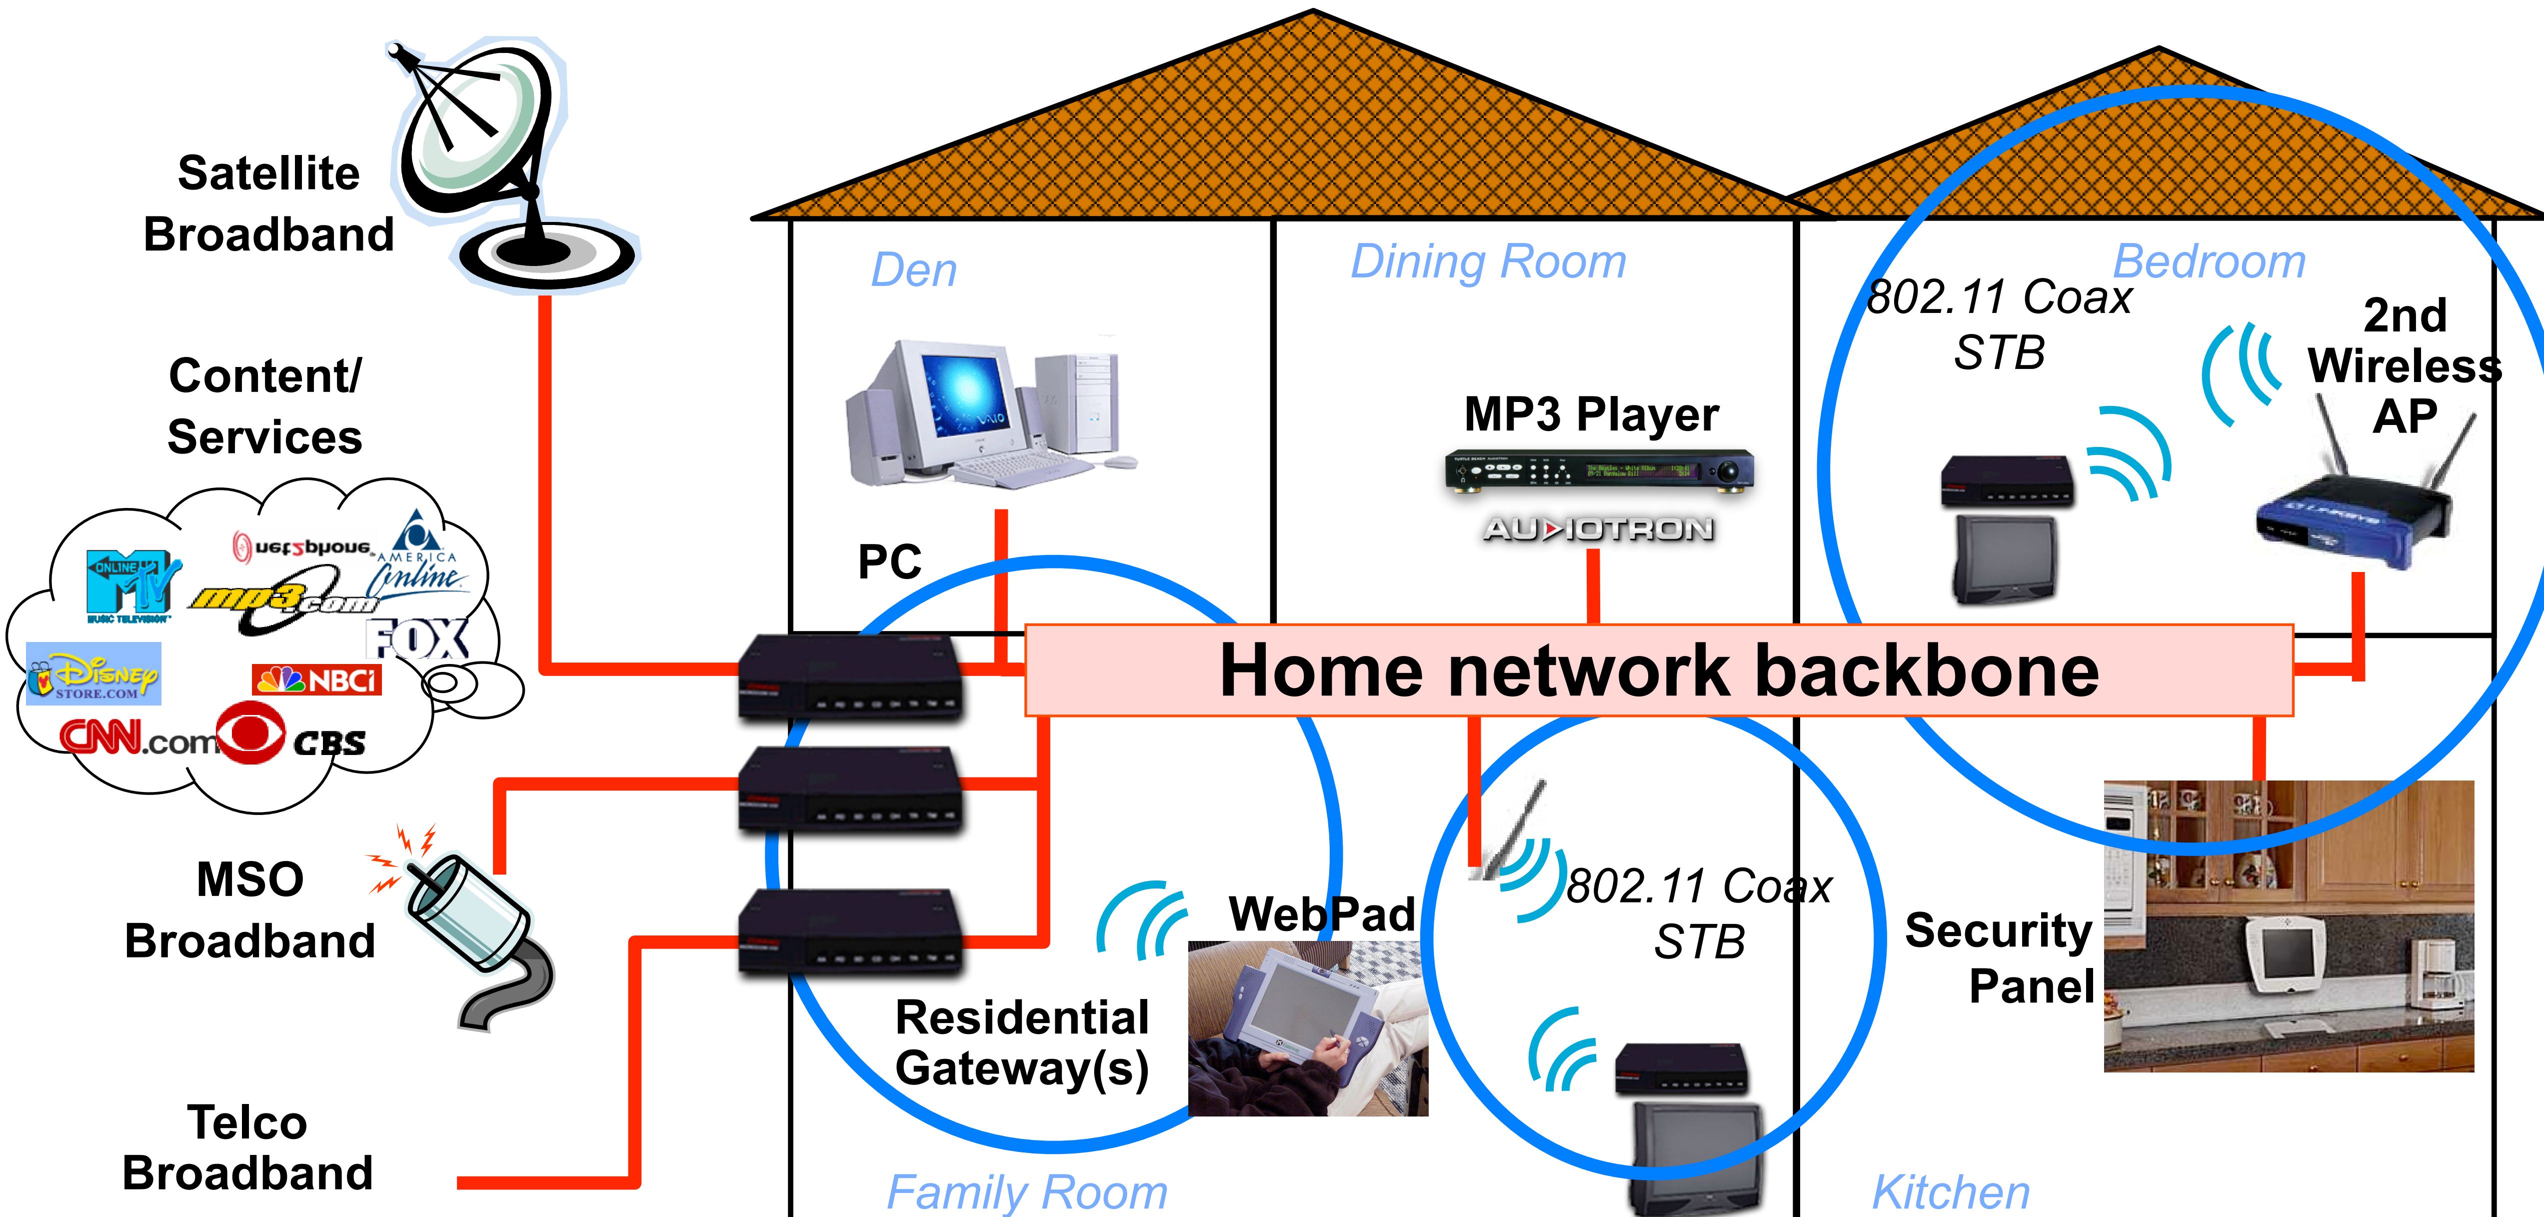

Where will Ethernet AV be used?

- Backbone for home
  - The "Gold Standard"
    - Highest quality/lowest cost way to interconnect wireless A/Ps
    - "Perfect" QoS, requires the least customer interaction
- Within the entertainment cluster
  - Trivial wiring, no configuration, guaranteed 100/1G/2.5G+ ***per device***, not just ***per room*** or ***per house***
  - PoE for speakers, extra storage (HD/optical), wireless A/Ps, other lower-power devices
  - Ideal long-term replacement for 1394
- Numerous non-"residential" applications
  - Professional audio/video studios, industrial automation, test and measurement

Proposed architecture

- Changes to both IEEE 802.1Q and IEEE 802.3
  - 802.1Q - bridges/switches - most of work
  - 802.3 - Ethernet MAC/PHY - small change to MAC
- Three basic additions to 802.3/802.1
  - Traffic shaping and prioritizing,
  - Admission controls, and
  - Precise synchronization

Topology & connectivity

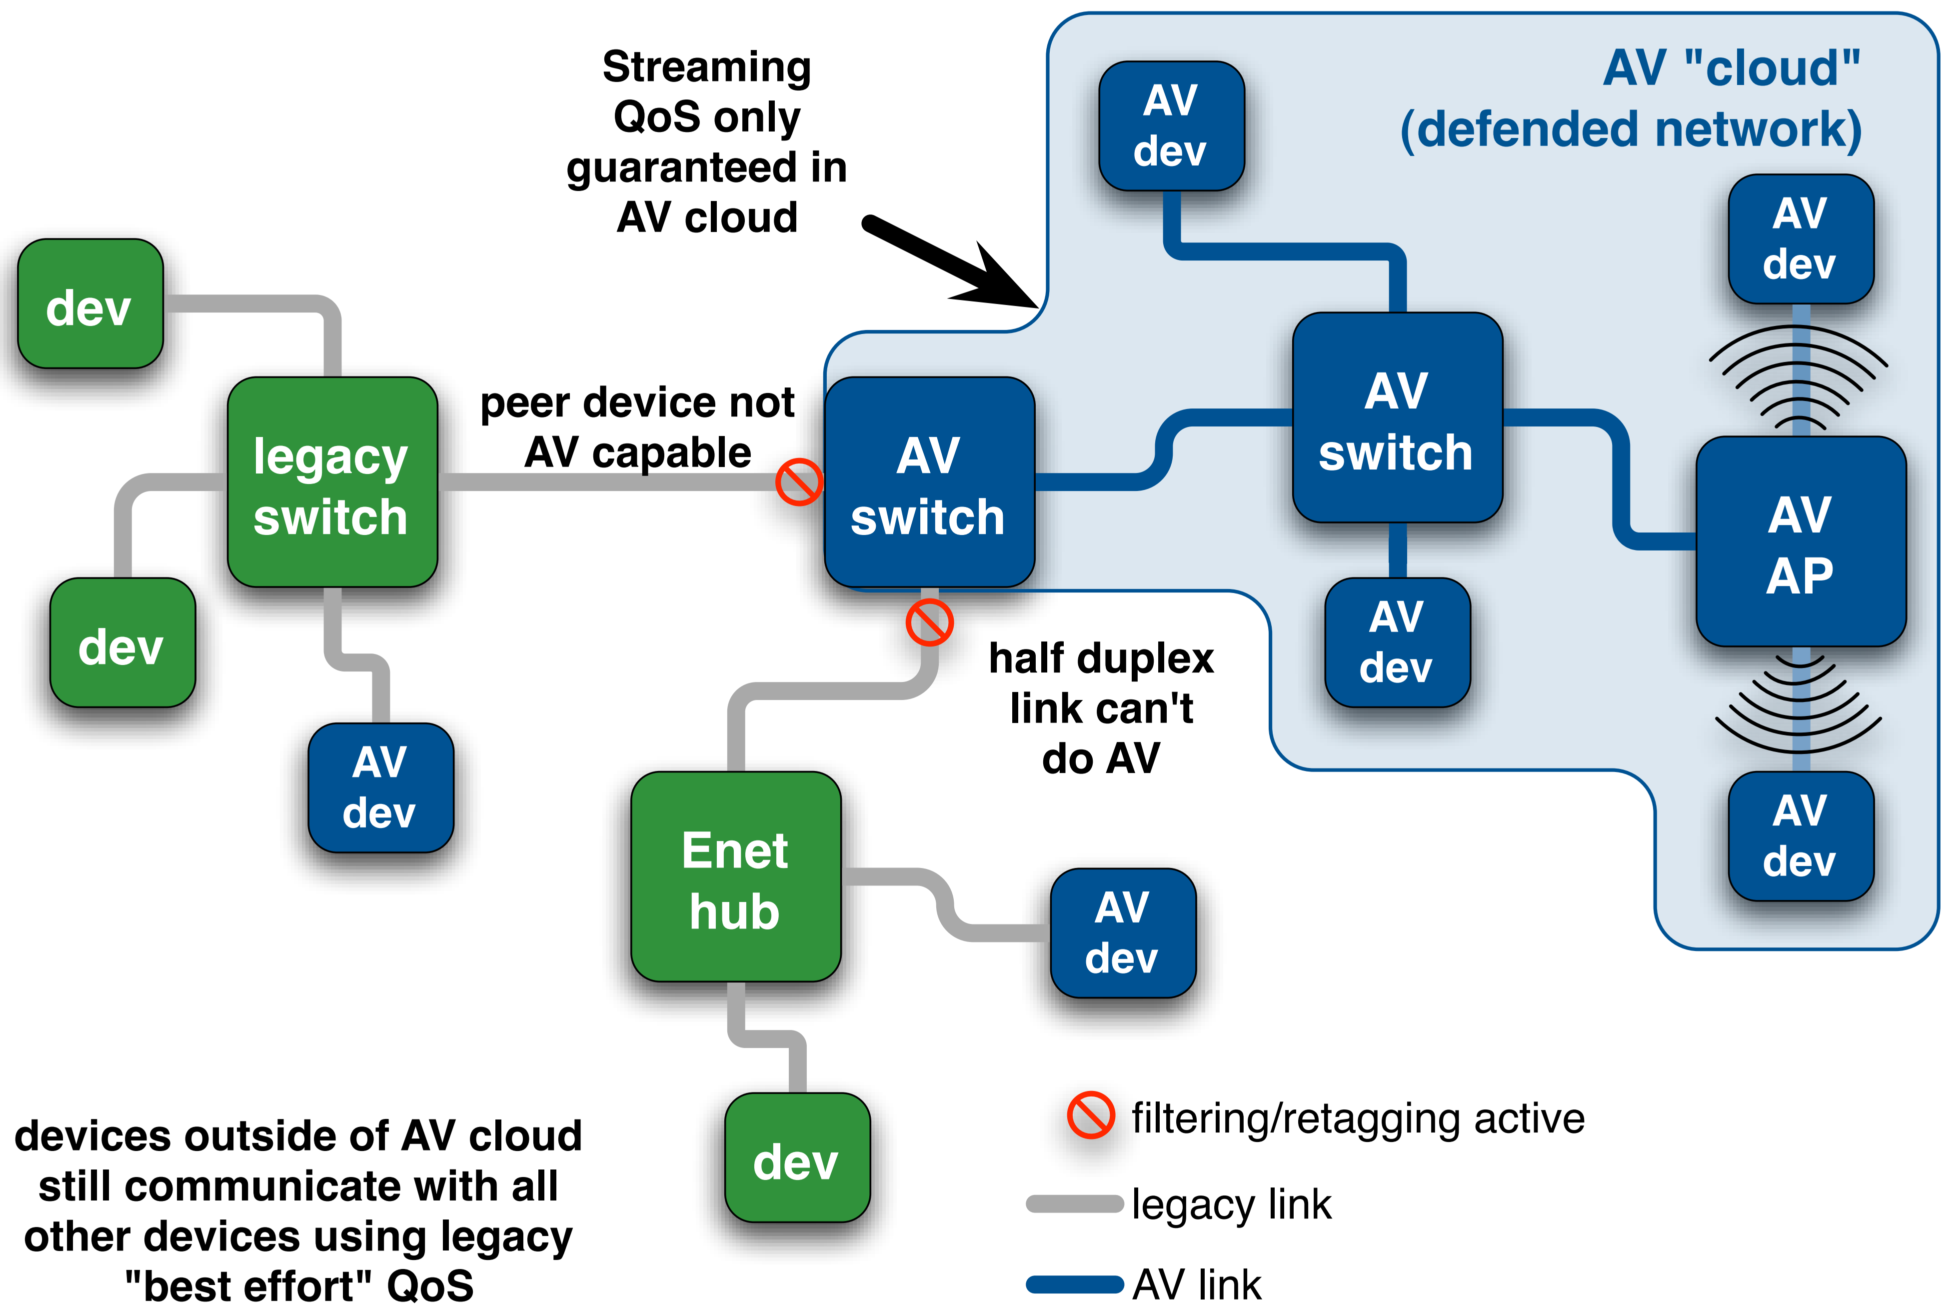

Establishing the AV cloud

- 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 Priorties

- 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 worst case buffers
- Mapping between traffic class and priorities

Traffic Class?

- 802.1p introduced 8 different traffic classes
  - Usually implemented as strict priorities
    - tagged frames, priority field as "PCP"
  - 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

- 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 needed tuples
  - Perhaps other information useful for stream management such as path availability

Admission Control (I)
(registration)

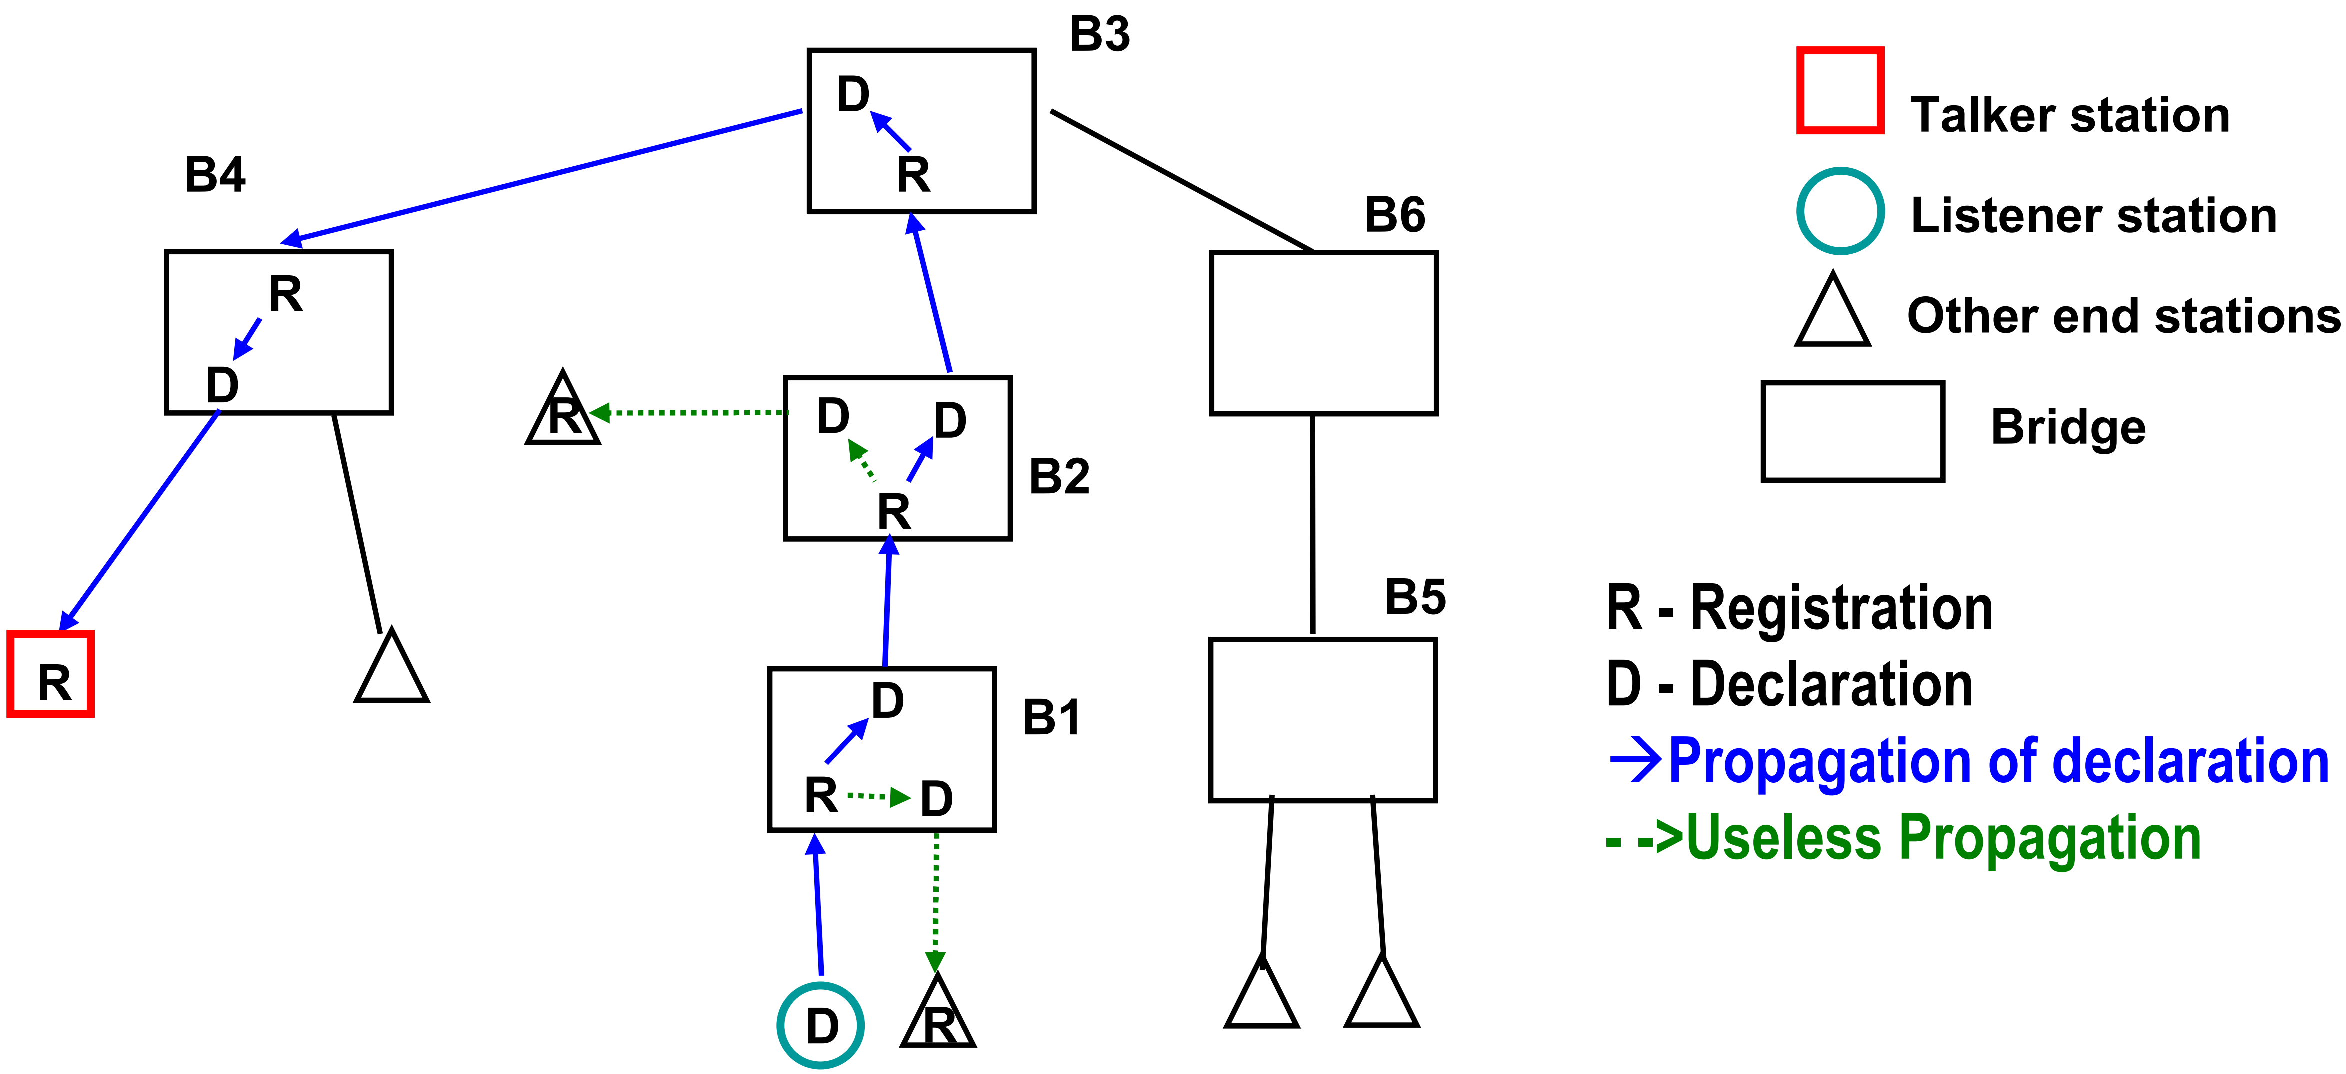

- With MSRP registration, the talker and intermediate bridges know where are potential listeners and how to get to them
- Assume in the above figure, B3/B4 have learnt the talker's address, and B1/B2 haven't, then:
  - MSRP floods the registration if the talker's address is not in the bridge FDB (eg.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

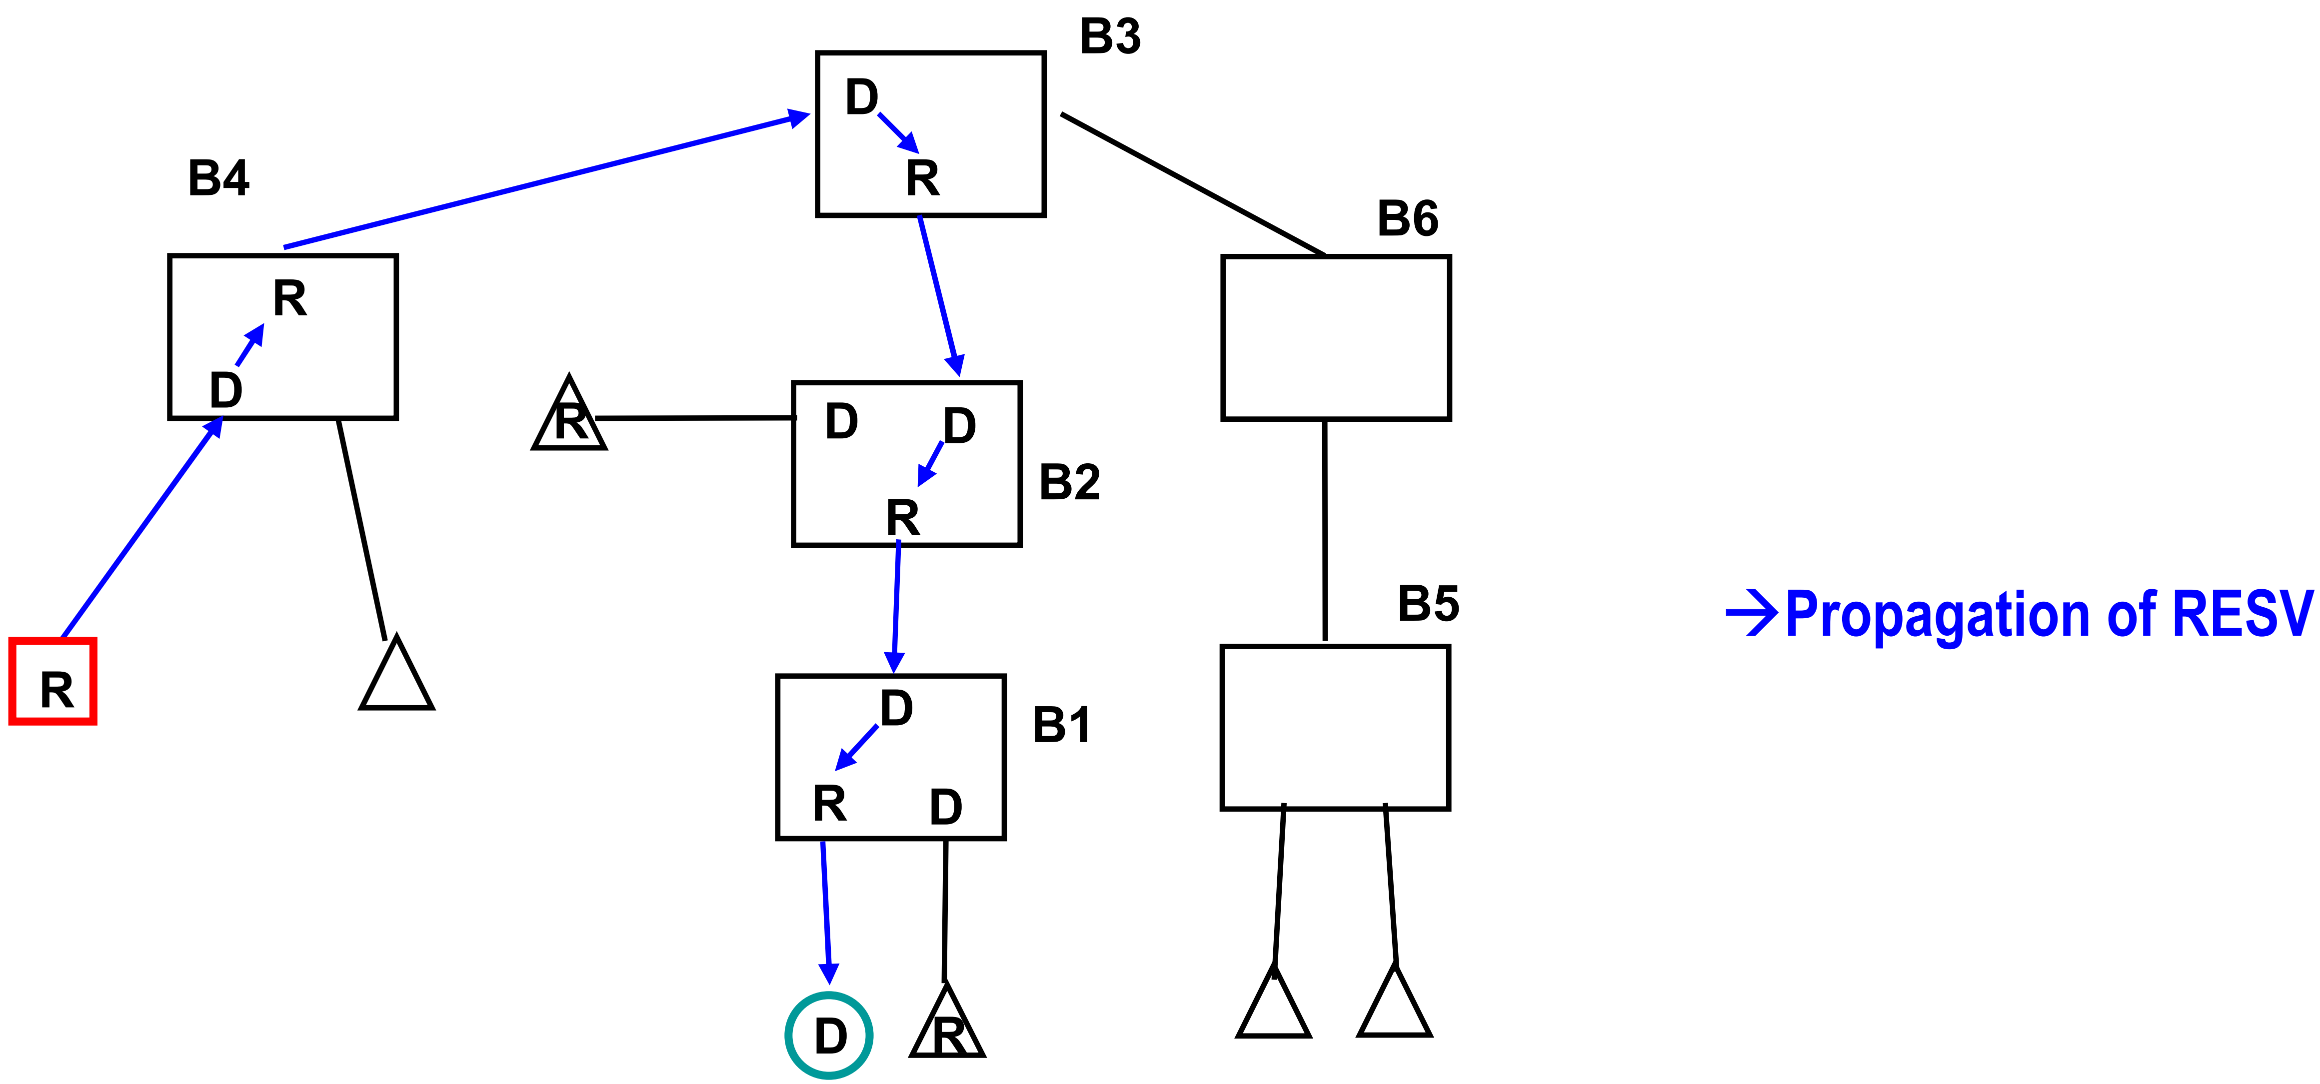

- RESV signaling triggers admission control operations in intermediate bridges. It also locks resources and updates isochronous filtering database if the admission control is successful.
- In this example, admission control is successful along the wholepath. RESV signaling servers as the end-to-end explicit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

Admission Control (3)
(failed reservation)

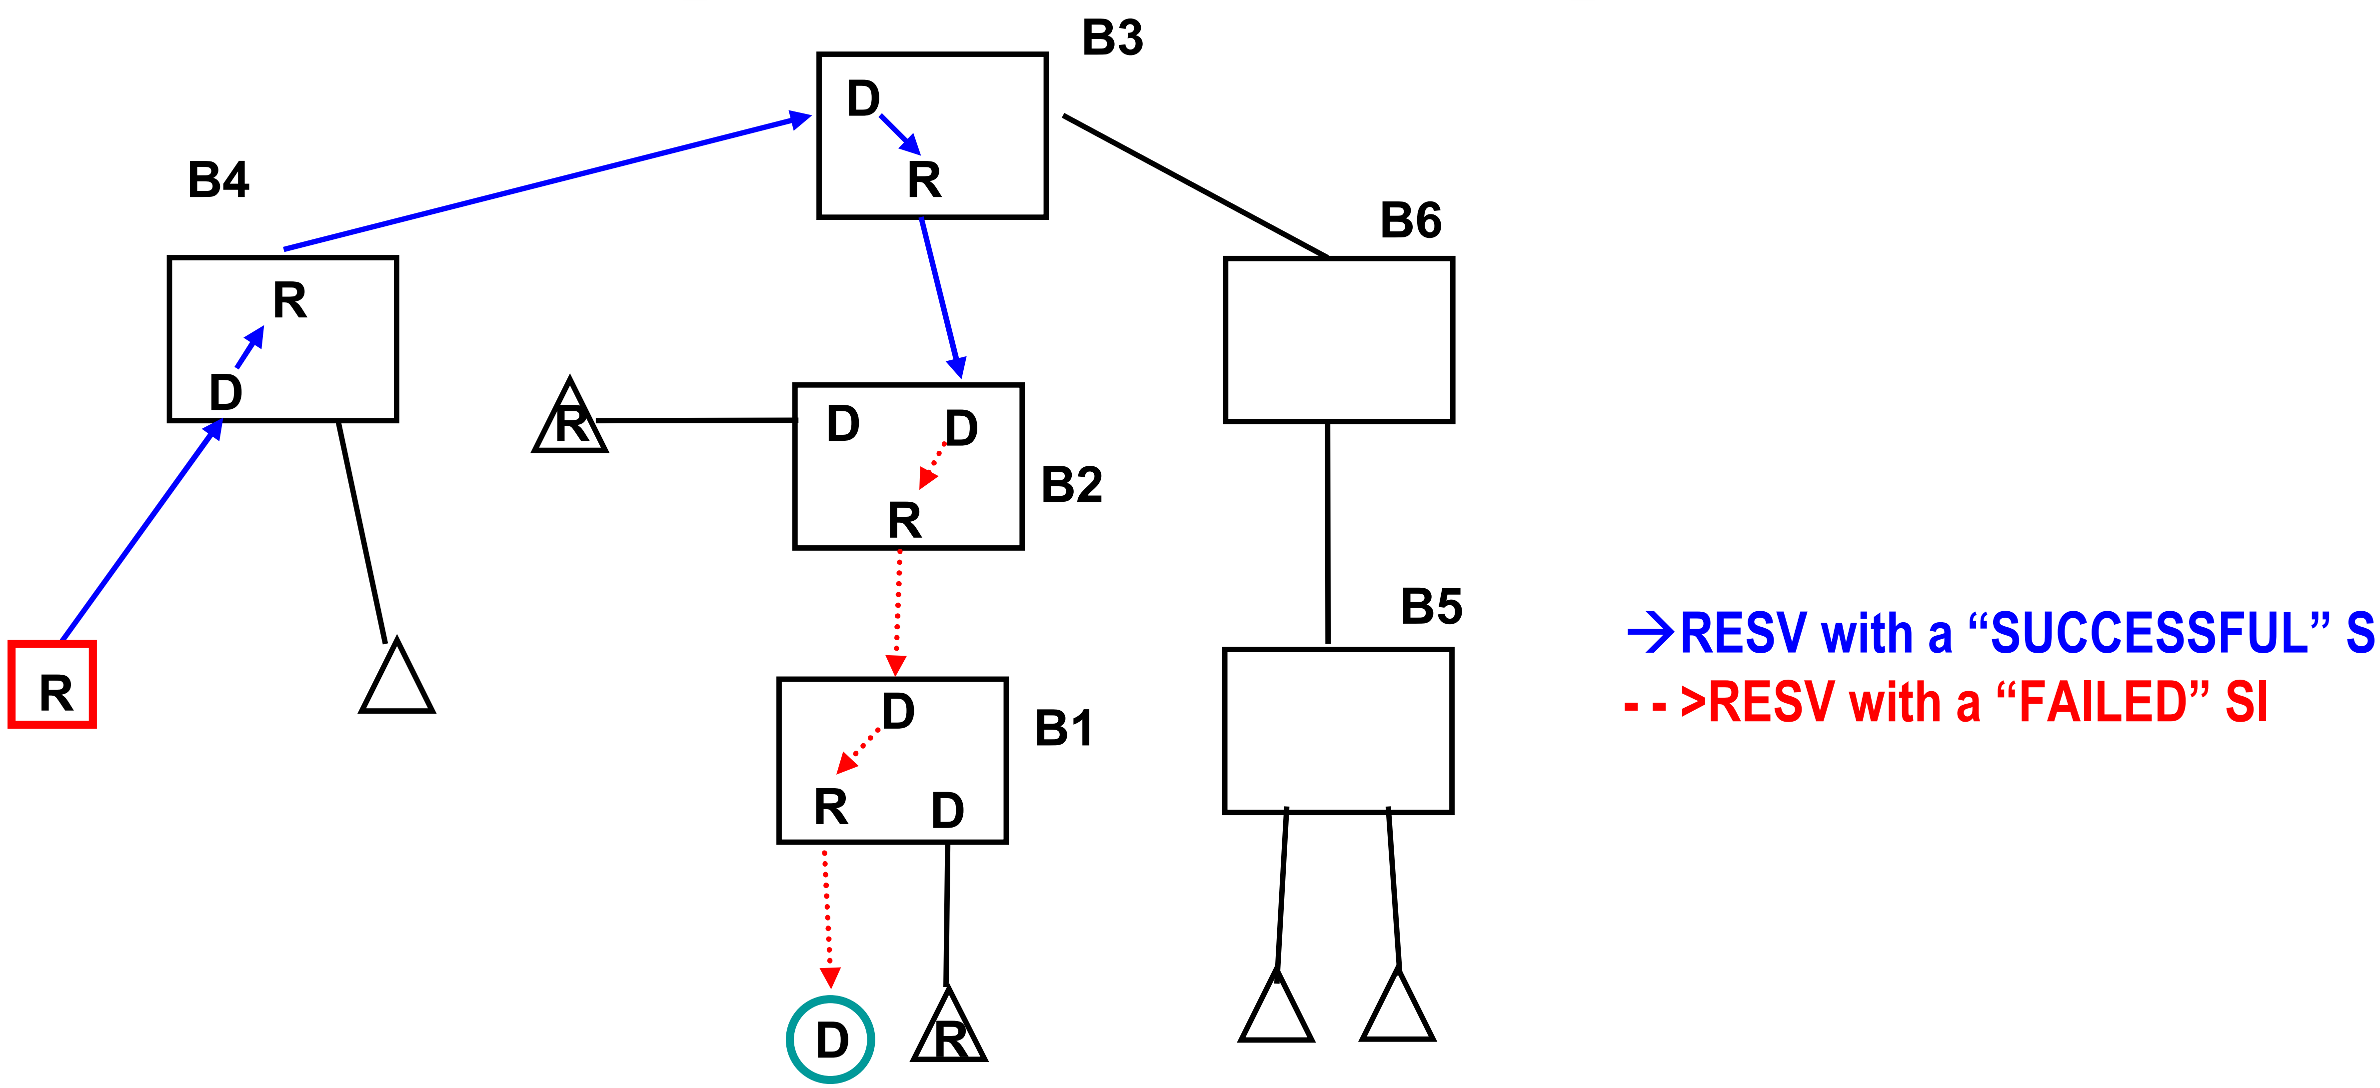

- In this example, admission control is failed at B2. The SI (Status Indication) bit of RESV signaling will be set to FAILED.
- The RESV is still forwarded to the listener. However, downstreambridges (i.e., B1, B2) will not lock resources for the RESV signaling whose SI is set to FAILED.
- Listener is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

Precise synchronization

- AV devices will periodically exchange timing information
  - both devices synchronize their time-of-day clock very precisely.
- 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 1588, but much simpler
  - likely to be a part of new IEEE 1588v2

17

AV Grand Master clock

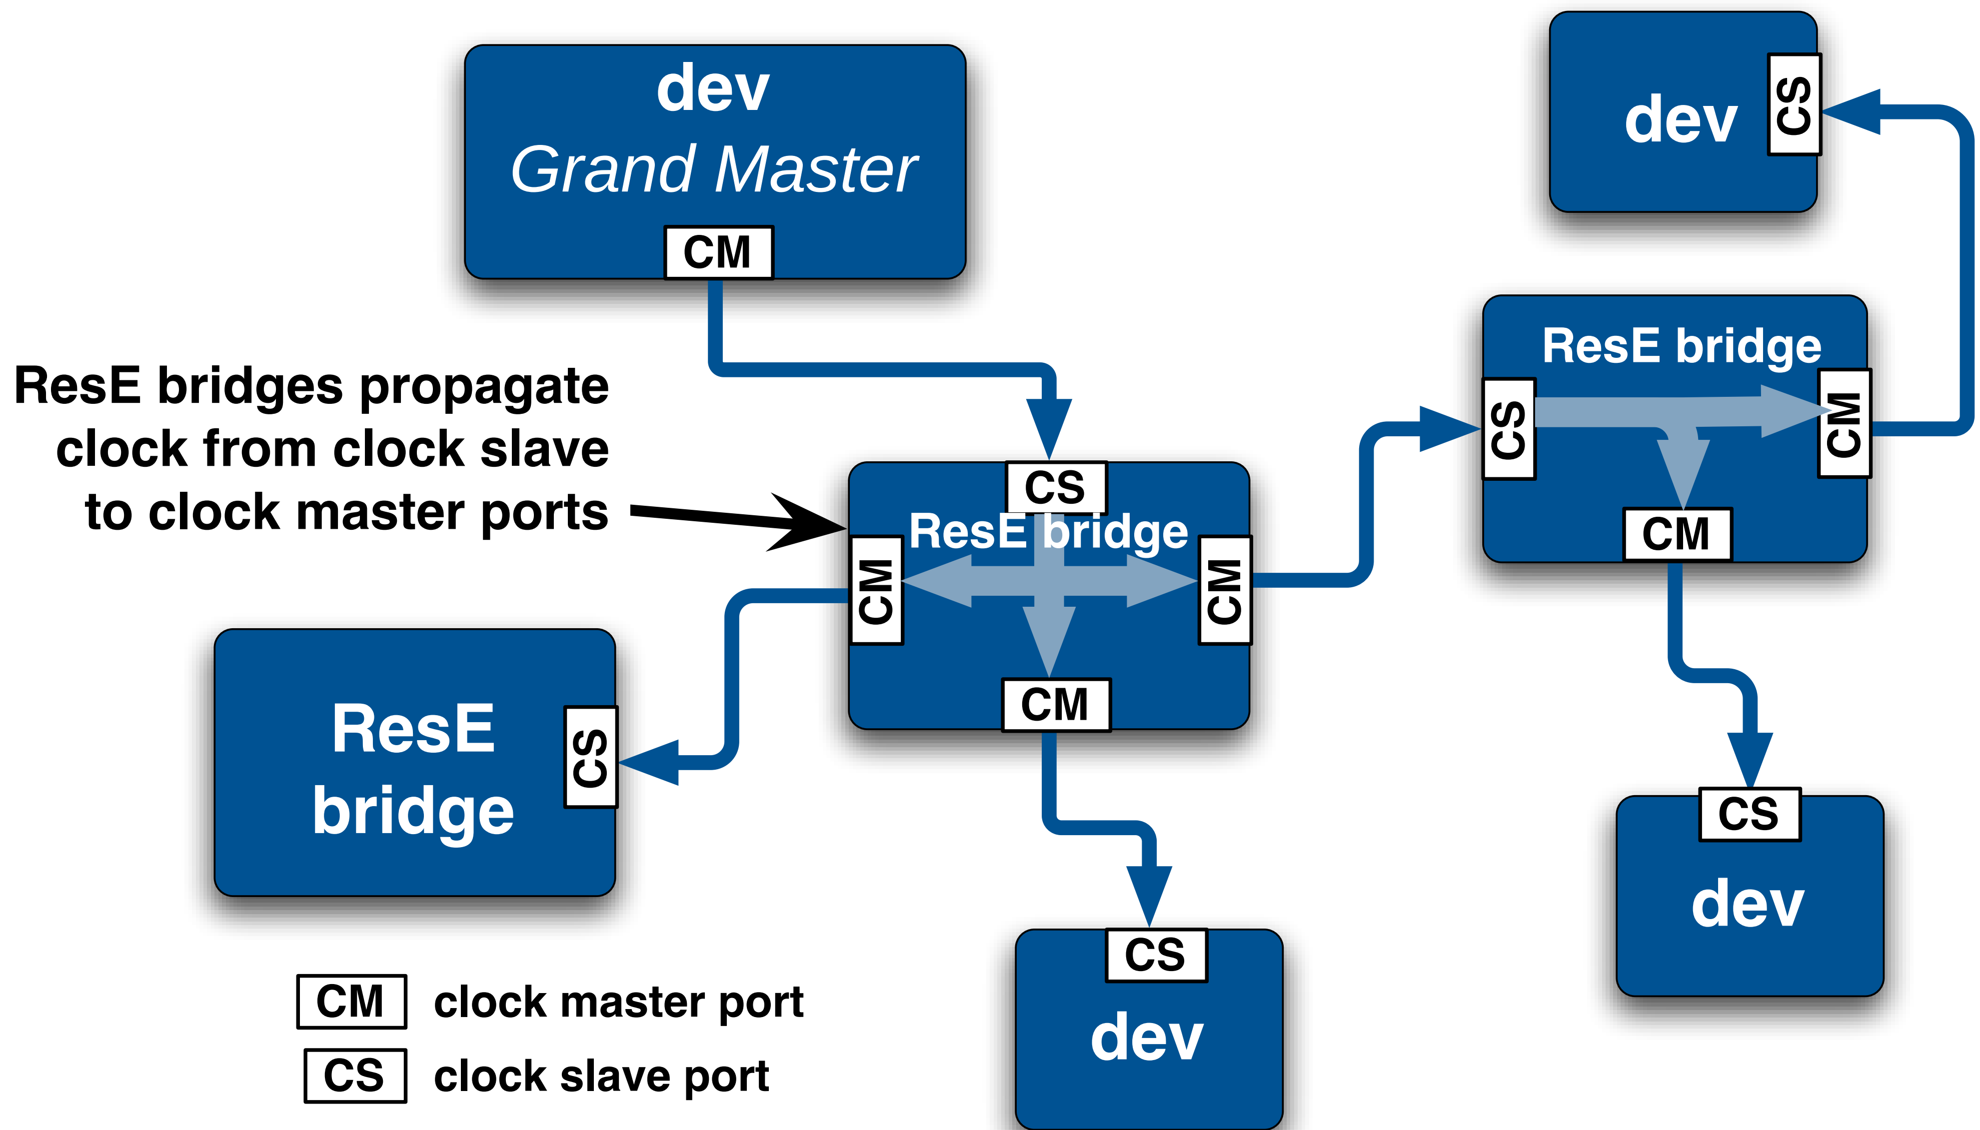

- There is a single device within an Ethernet AV "cloud" that provides a master timing signal.
  - All other devices synchronize their clocks with this master.

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

19

Changes to Ethernet NIC

- MAC changes
  - Frame Timer – Accurately note time of RX/TX Ctrl Frame
    - Not really a change to "MAC", but to buffers for the MAC
- Queuing/DMA
  - Separate queues and DMA for class 4/5 frames to provide appropriate traffic shaping (scheduling)
    - One extra queue/DMA channel possible
- Admission Control (driver 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

Changes to Ethernet Switch

- MAC changes
  - Frame Timer – Accurately note time of RX/TX Ctrl Frame
- Bridging
  - Ingress filtering/shaping at edge of network to ensure proper traffic shaping for class 4/5 (streaming) frames
  - Egress filtering to ensure that streaming CoS not over-utilized
- Admission Control
  - 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 engine per port - only if wanted to reduce switch CPU processing
- Reasonable Microprocessor Cycles
  - Scales with # of ports similarly.

When will Ethernet AV be available?

- IEEE standardization process started
  - Originally an 802.3 study group, moved to 802.1 in November 2005 as "Audio/Video Bridging Task Group"
  - Early drafts already available
  - Expect technical closure in 2006, final standard in 2007
- First hardware/software soon after stablilization
  - Possibly a number of "pre standard" interations
- Later editions support uncompressed HD video
  - "multiGigabit" NIC/Switch ("Ethernet HD")

Example EnetAV NIC

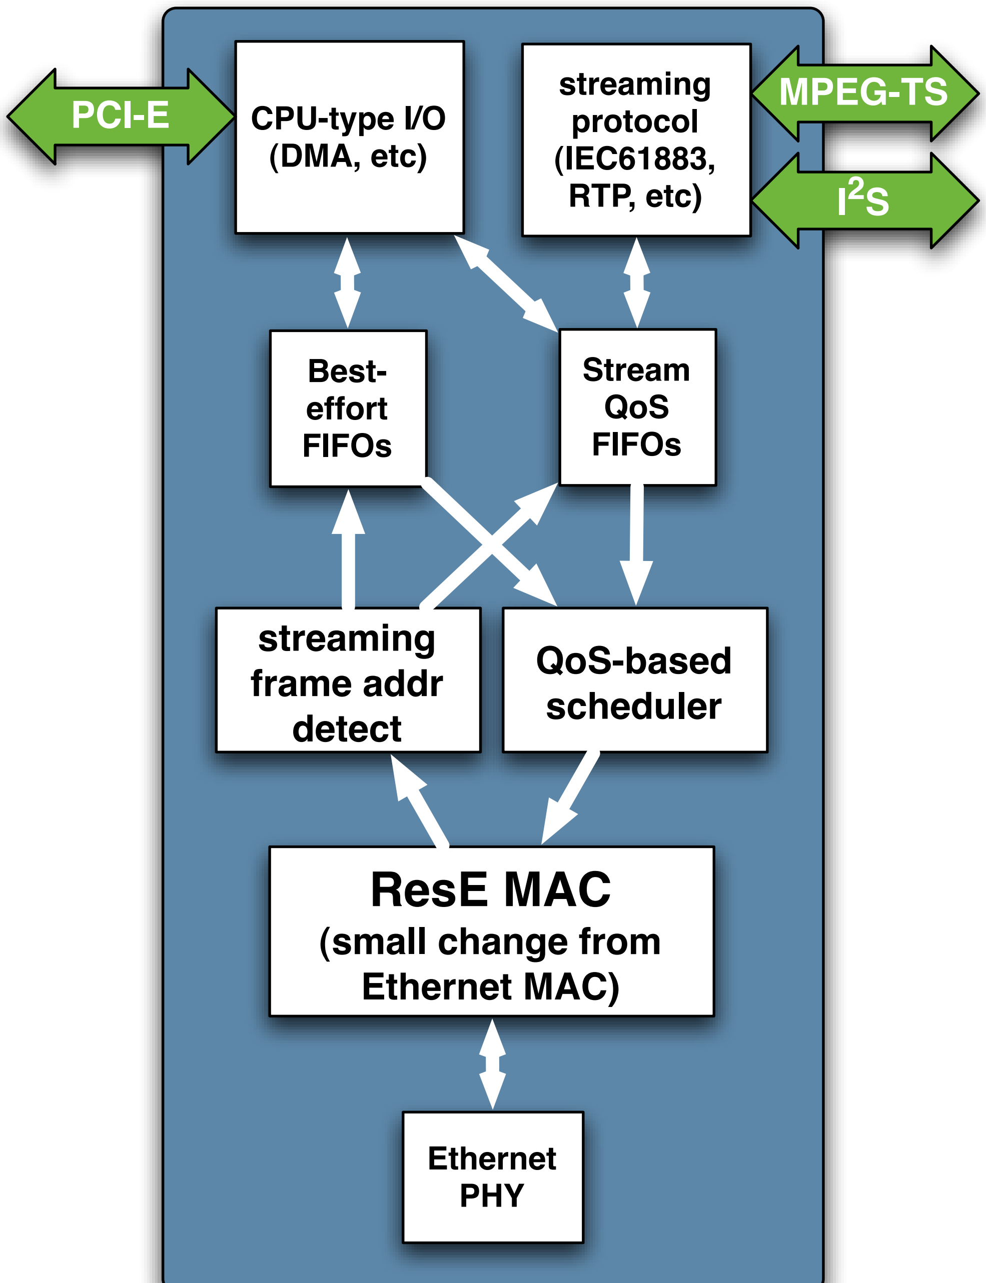

- PHY is Fast Ethernet or better
  - 1G for backbone, 2.5G/10G for uncompressed video
- CPU interface is PCI-E
  - Streaming frames on PCI-E use virtual channel
  - Perhaps parallel PCI for CE?
- Streaming I/O?
  - MPEG-TS 10-signal I/F?
  - I2S 3-signal for audio?
  - DVI for uncompressed video?

Example EnetAV bridge

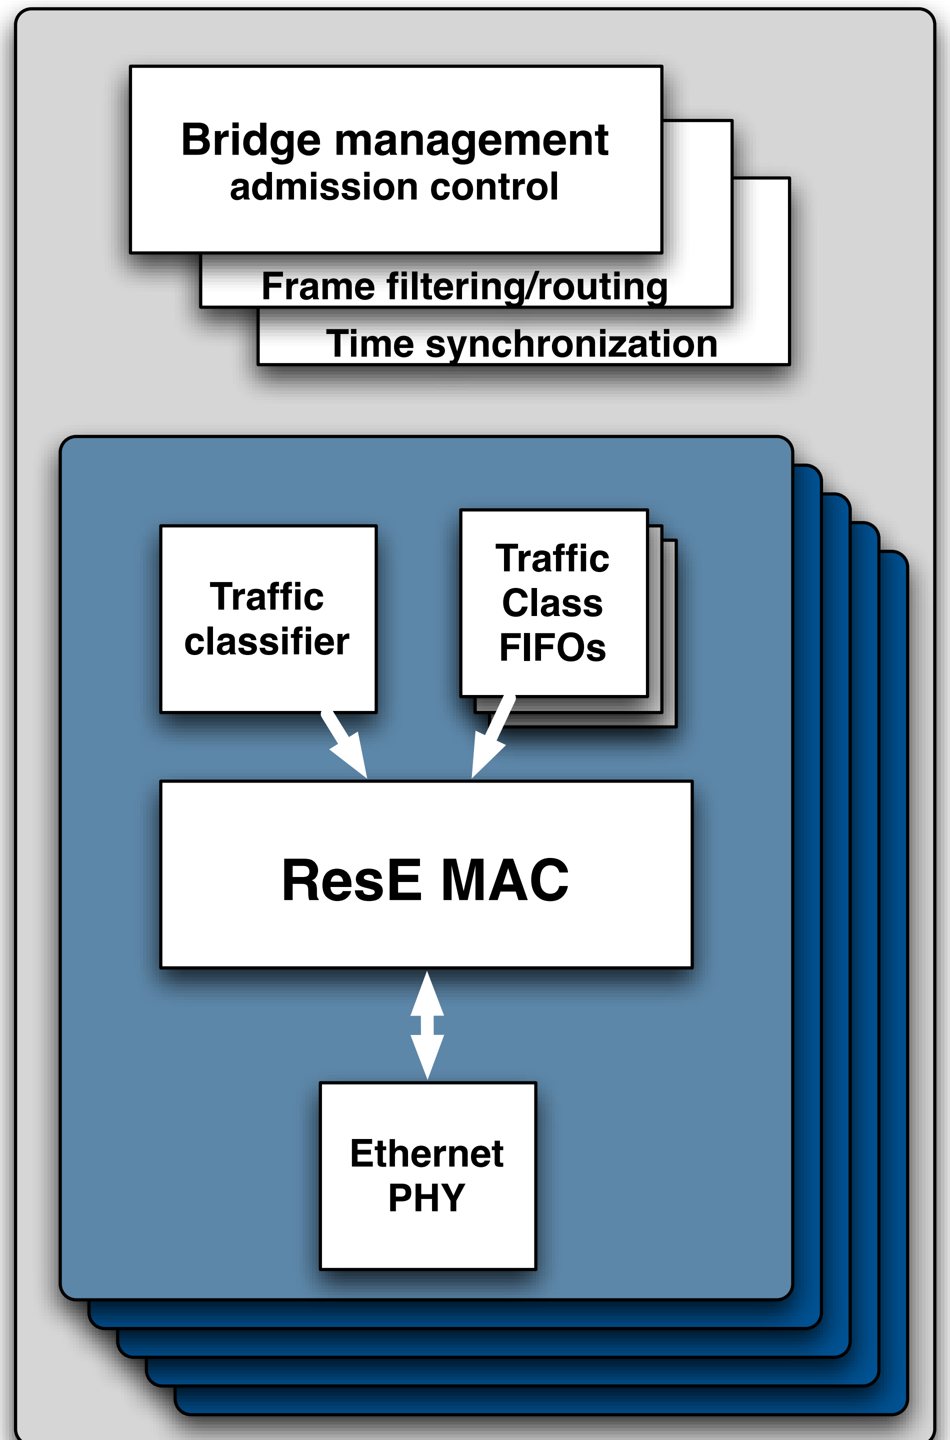

- PHYs are Fast Ethernet or better
  - 1G/5 port for first versions
- Separate CPU at first, but moving to integrated processor ASAP

Example multiport Enet AV NIC

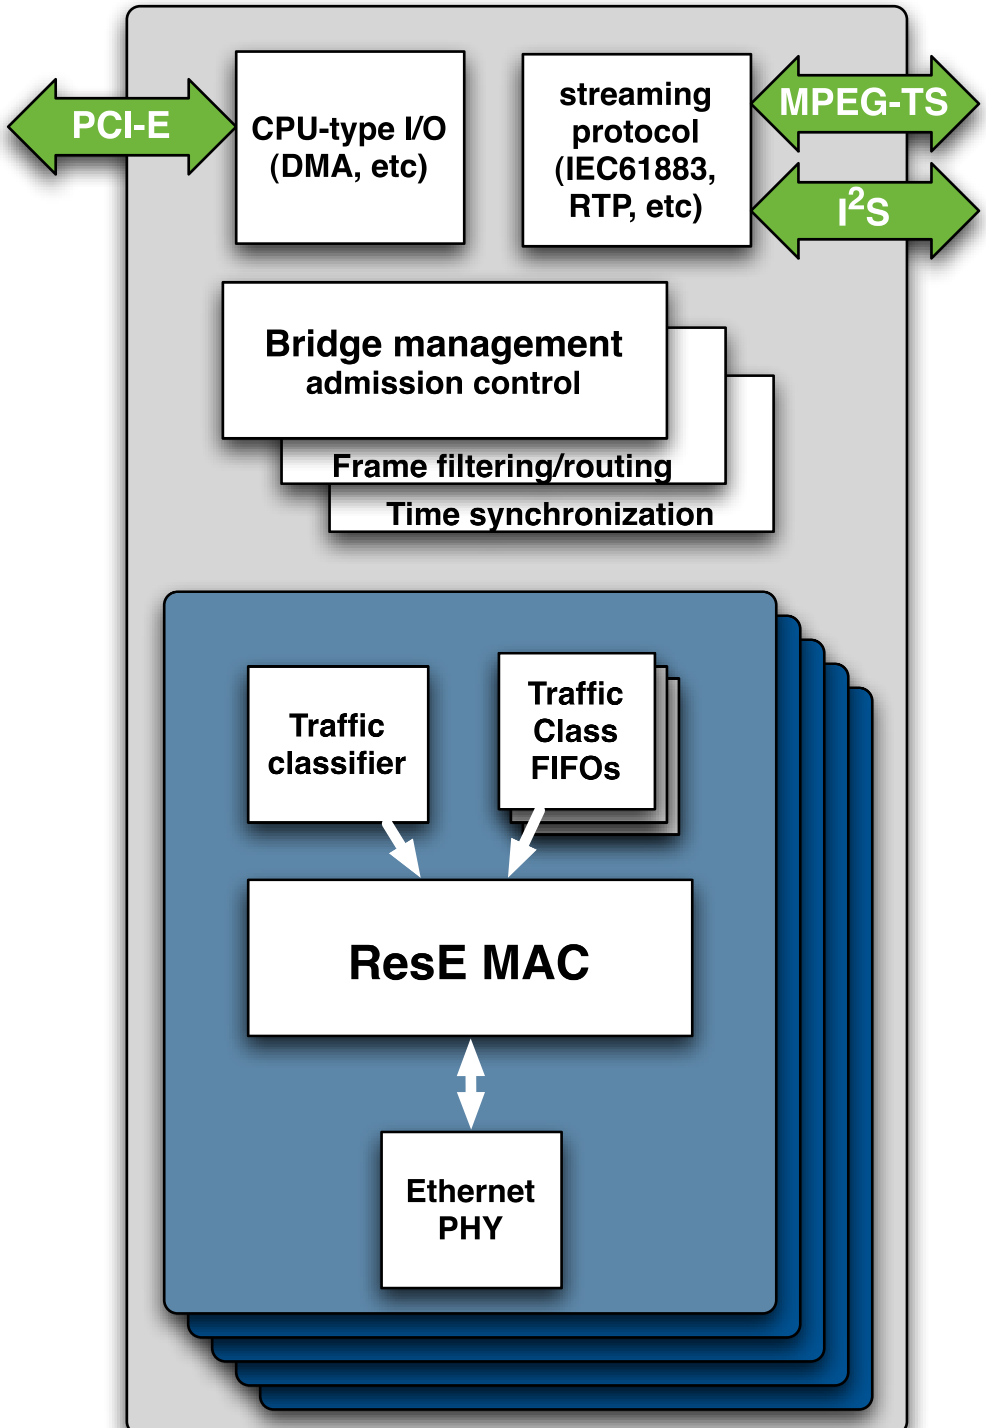

- Best product for TVs, STBs, home gateways, media PCs
- Anything that is a "hub" in the cluster

Beyond Ethernet AV

- Ethernet HD™
  - Multi-gigabit Ethernet with AV QoS
- Multi-gigabit?
  - Supports uncompressed HD video at 1.4+Gbit/sec plus multiple compressed streams and regular best-effort traffic
- Higher layer protocols
  - Streaming protocol adaption layer using IEC 61883 (for non-routable streams) or RTP (for IP-routable streams)
  - DTCP, DTCP-IP, and/or HDCP for content protection
  - UPnP, CEA-2027 or simplified RSVP for stream establishment
-

Why uncompressed HD?

- “Video Rich Navigation”
  - User interface information generated locally (STB/ gateway)
  - Compression adds significant cost
- Professional usage
  - Studios/performance
- Games
  - Compression adds significant cost …
  - And much more importantly …

Latency!

Summary

- Ethernet AV will be the standard interconnect for uncompromised quality of service
  - soon!
- There will be growth in both technology (speeds and feeds) and infrastructure (switches, ICs, intellectual property)
  - The first providers set the real standards, the interoperability requirements.

Thank you