AV Bridging and Ethernet AV™

Michael Johas Teener
Mikejt@broadcom.com

Agenda

- What is AV Bridging?
  - and Ethernet AV?
- Why is it needed?
- Where will it be used?
- How does it work?
- How can it be used in carrier networks?

BROADCOM®
Connecting everything®

What is AV bridging?

- Add Precise Timing in an 802 network (~100 nS native, ~100 pS achievable)
- Add E2E Bandwidth 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

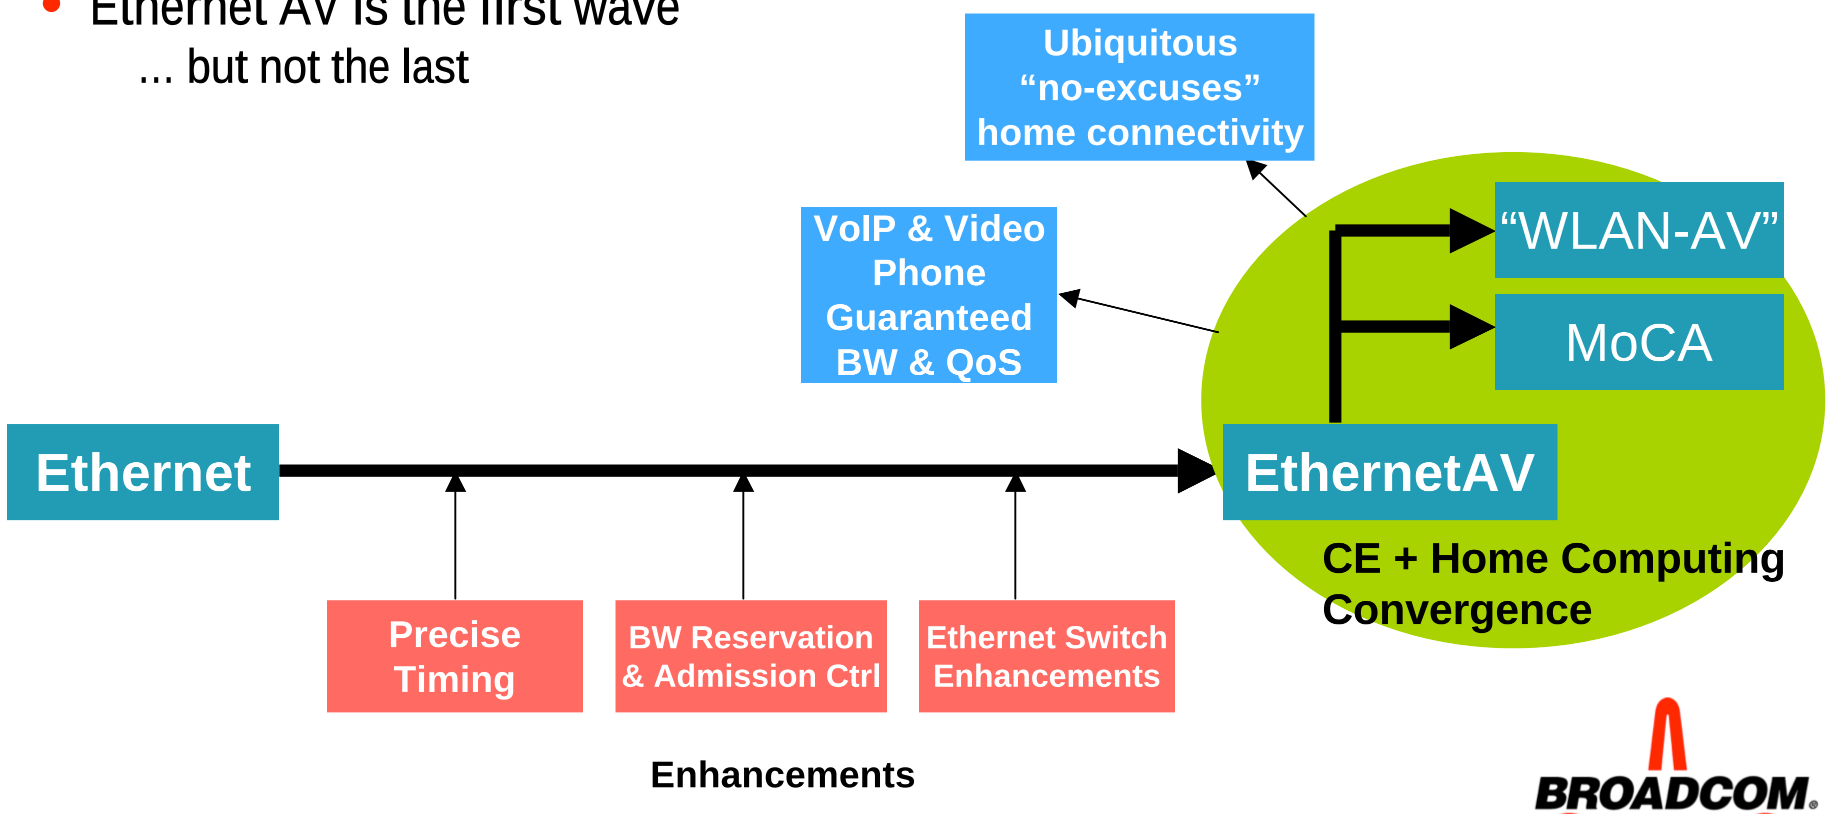

BROADCOM
Connecting everything

Gaining Support and Momentum

- Names of the companies who has been supporting the IEEE standard activities.
- Stronger support comes from companies who had had triple play trials.

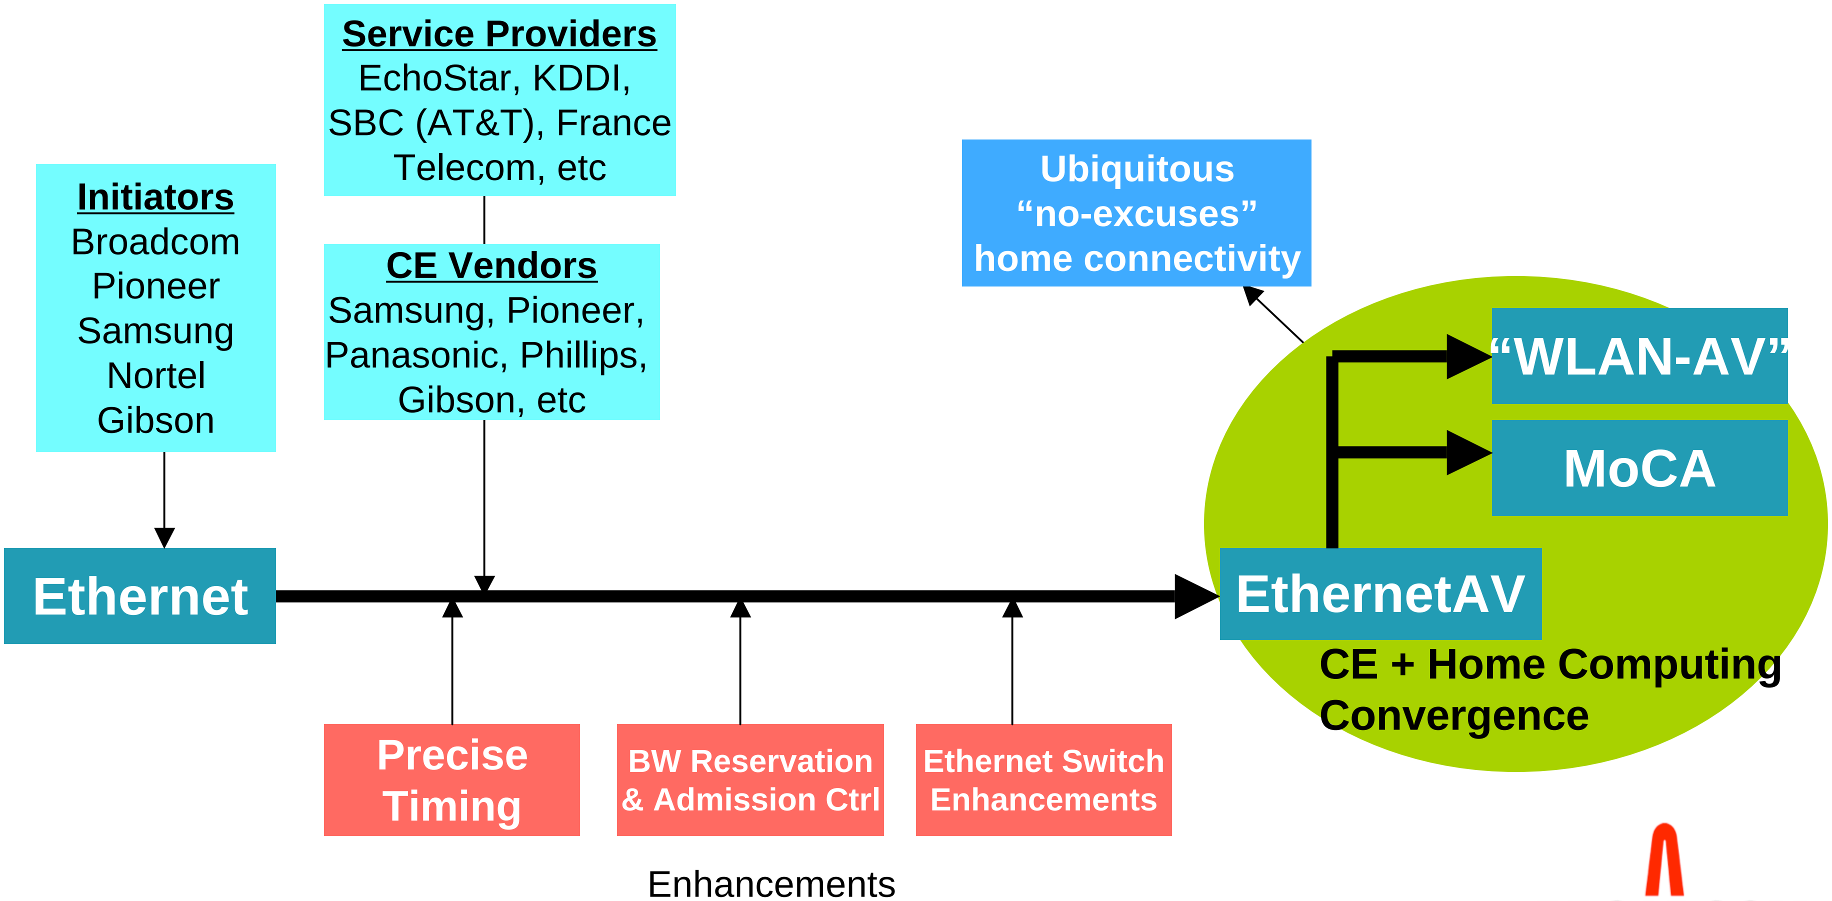

Why is it needed? (1)

- Common IT-oriented networks have inadequate QoS controls
  - all use 802.1 “priority” (actually, “traffic class”)
  - can work in controlled environments (same higher layer QoS)
  - no guarantees, timing synchronization difficult
- Ethernet is the best
  - but it’s easy for the customer to misconfigure or overload
  - no guarantees, timing synchronization difficult
- Wireless has superior convenience, but inadequate bandwidth and excessive delays for whole-home coverage
  - 802.11n will work for many applications, but …
  - latencies through multiple A/Ps may be too much for interactive applications
  - difficult environment (interference, variable S/N), so 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 (it's almost everywhere in NA)
  - 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 (it IS everywhere!)
  - 1394b/c long distance has limited developer base & infrastructure, is not cheap

    … but even this is part of the solution (has best QoS guarantees)

Digital Home Media Distribution

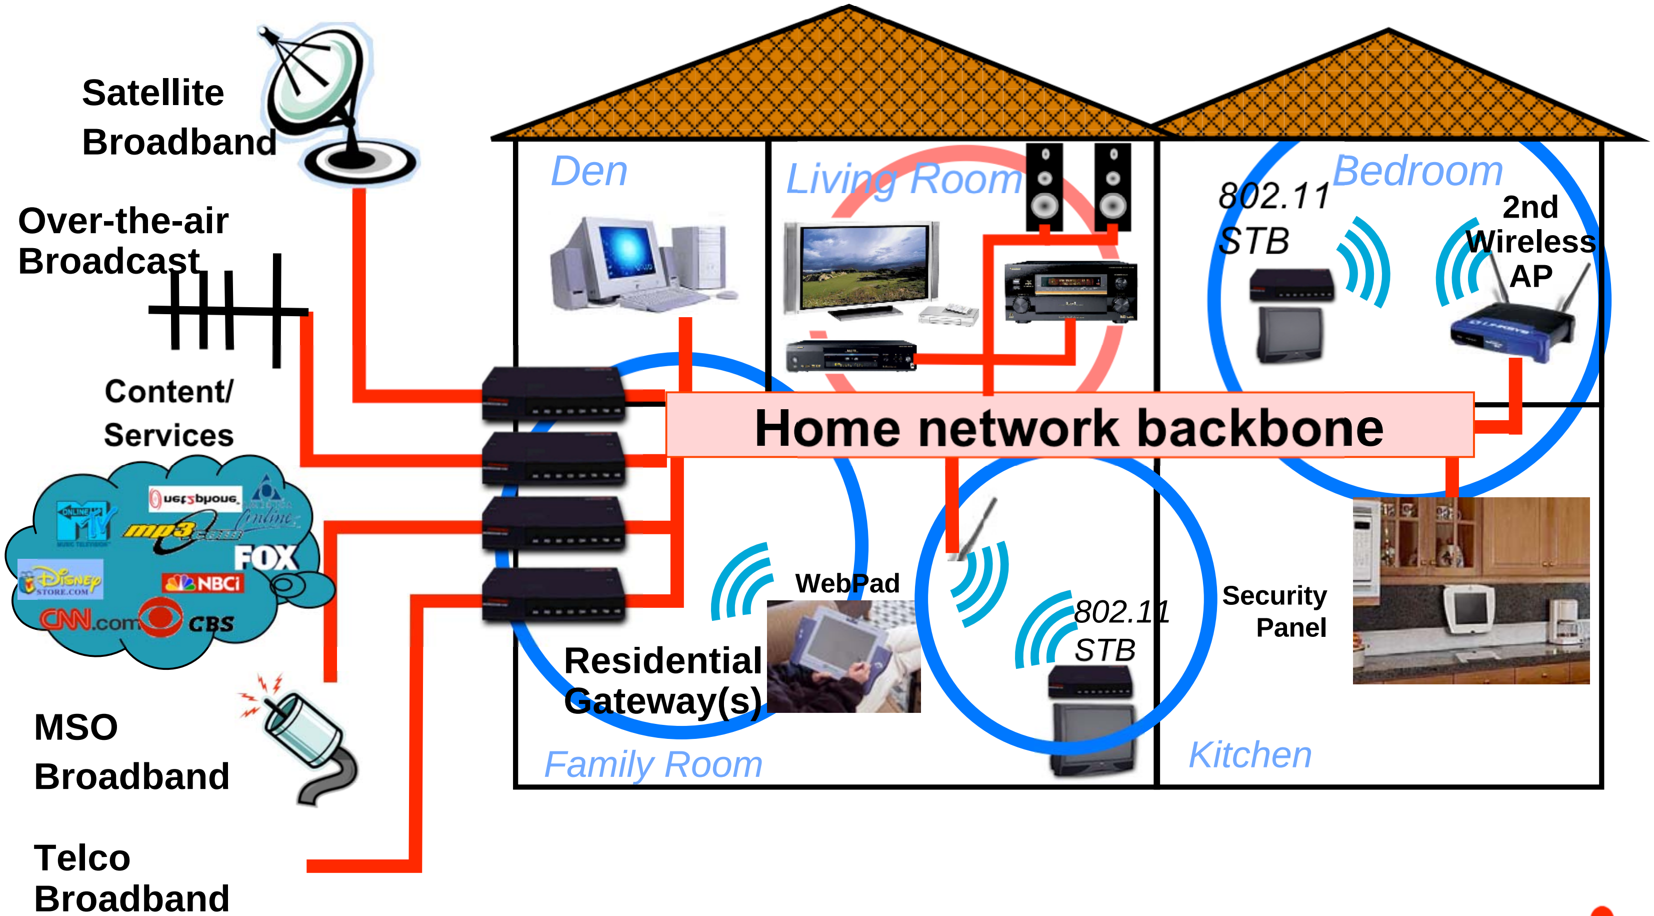

Ethernet AV:
the Gold Standard

- Backbone for home
  - Highest quality/lowest cost way to interconnect networked CE devices
  - “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, ***carrier backbone***

But it's only part of the solution!

- Ethernet AV is the best backbone for QoS
  - But NOT the most convenient in much of the existing market (e.g., requires Cat5)
- Ethernet AV is the most cost-effective high QoS network for endpoints
  - But only if the endpoints are fixed (e.g., not useful for mobile)
- So we need to enable the heterogeneous network
  - Provide QoS services universally, as well as can be done for the particular layer 2 technology

So …

BROADCOM. Connecting everything®

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

BROADCOM®
Connecting everything®

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 (can meet the MTIE mask for professional uncompressed video)
- Trade group to provide trademark "enforcement" of otherwise optional features
  - Require useful bridge performance, network management, PoE management, auto-configuration features

BROADCOM® Connecting everything®

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 - WiFi - more work, but basic tools in place
    - Not discussed in this presentation, but watch this space!
- Three basic additions to 802.1/802.3
  - Traffic shaping and prioritizing,
  - Admission controls, and
  - Precise synchronization

Topology & connectivity

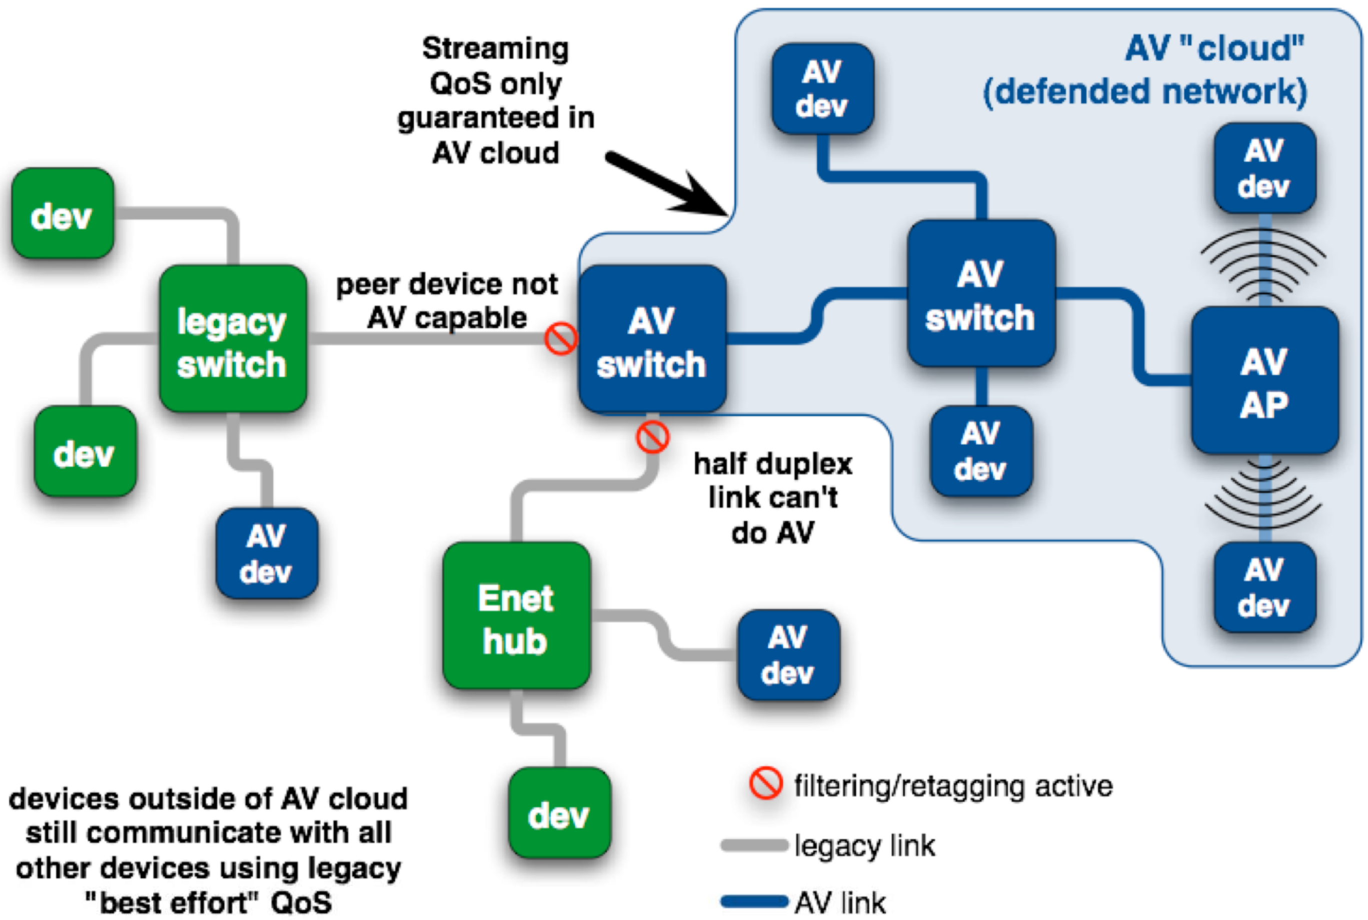

BROADCOM®
Connecting everything®

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 for 100baseTX, near 125 usec for 1G or better: interactive audio/video
  - Class 4 is for moderate latency streaming
    - Perhaps 1ms per bridge hop, longer for wireless, MoCA, power line: voice over IP, movies

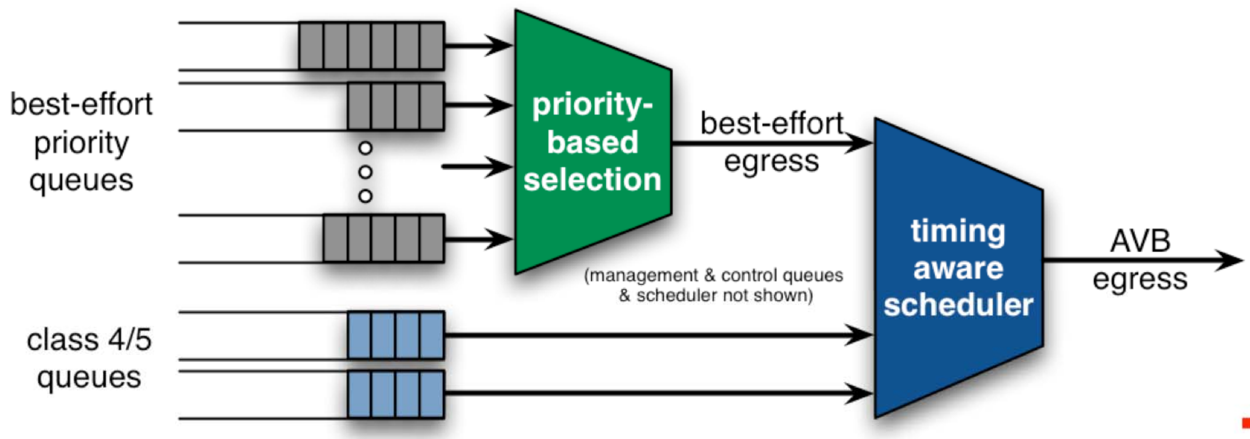

BROADCOM.
Connecting everything®

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 802.1ak "Multiple Multicast Registration Protocol" and the new SRP ("Stream Registration Protocol")
  - Registers streams as multicast addresses using MMRP
  - Reserves resources for streams as bandwidth/traffic class
    - Specifics for "bandwidth" include at least max packet size and number of packets per traffic class measurement intervals
  - Perhaps other information useful for stream management such as path availability, egress port MAC addresses, etc.

Connecting everything®

Admission Control (1)
(registration)

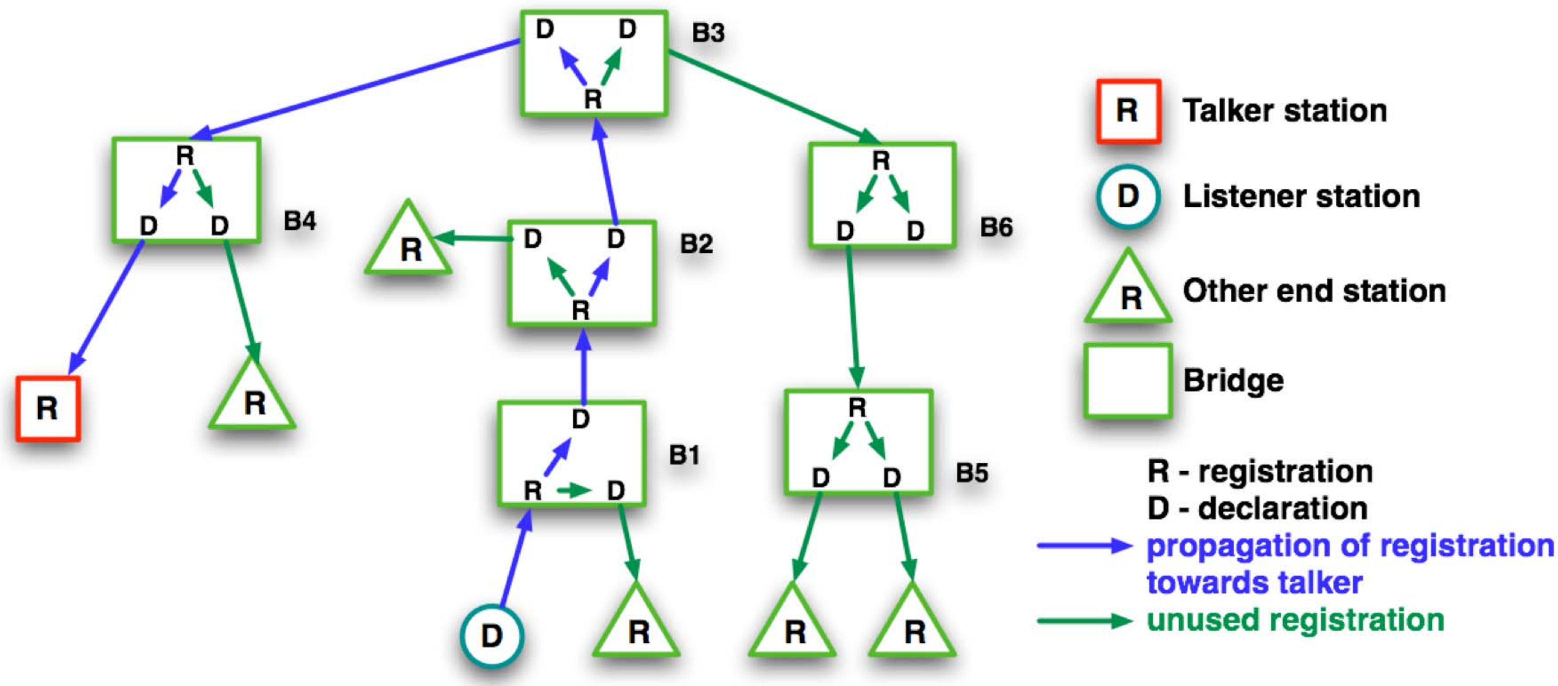

- With MMRP registration, the talker and intermediate bridges know where are potential listeners and how to get to them
- MMRP floods the registration

Admission Control (2) (successful reservation)

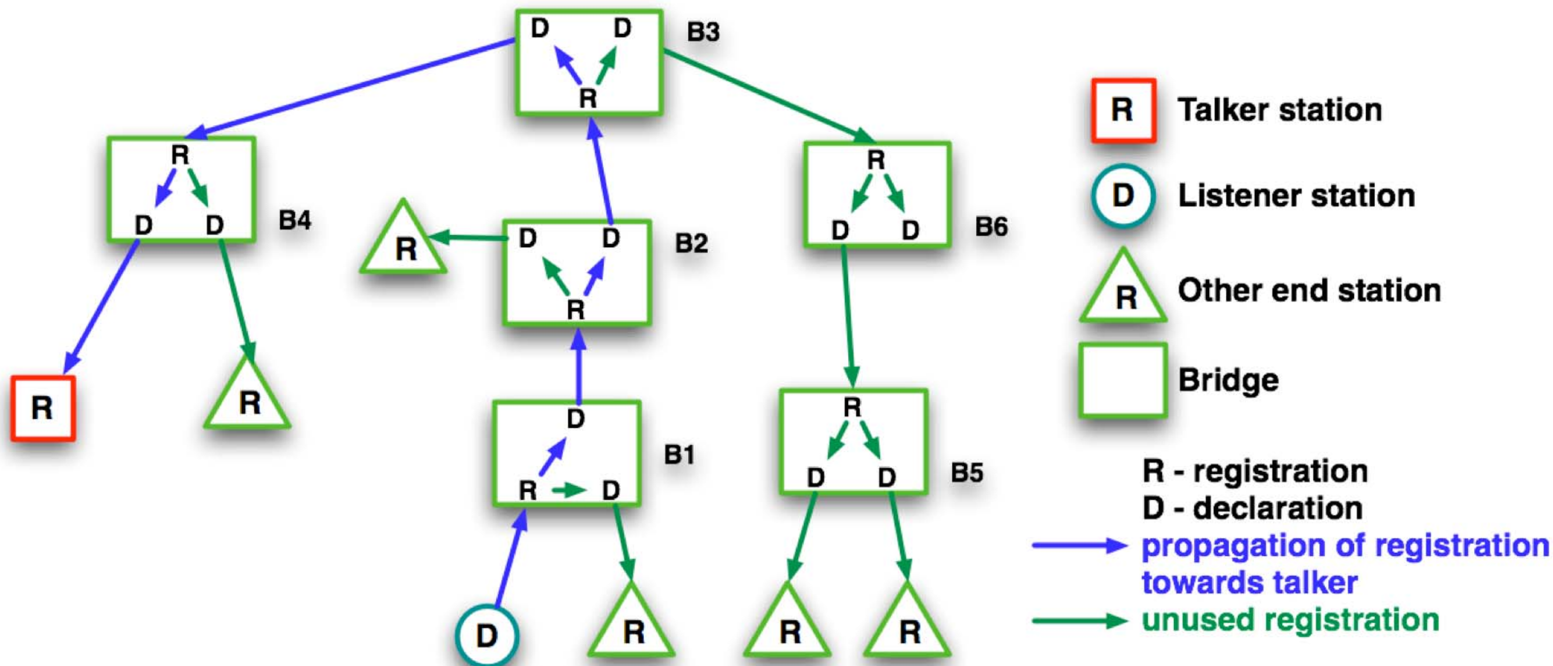

- **Reservation signaling triggers admission control operations in intermediate** bridges. It also locks resources and updates filtering database if the admission control is successful.
- In this example, admission control is successful along the whole path. Reservation signaling serves as the end-to-end explicit ACK signaling to listener.

Admission Control (3)
(failed reservation)

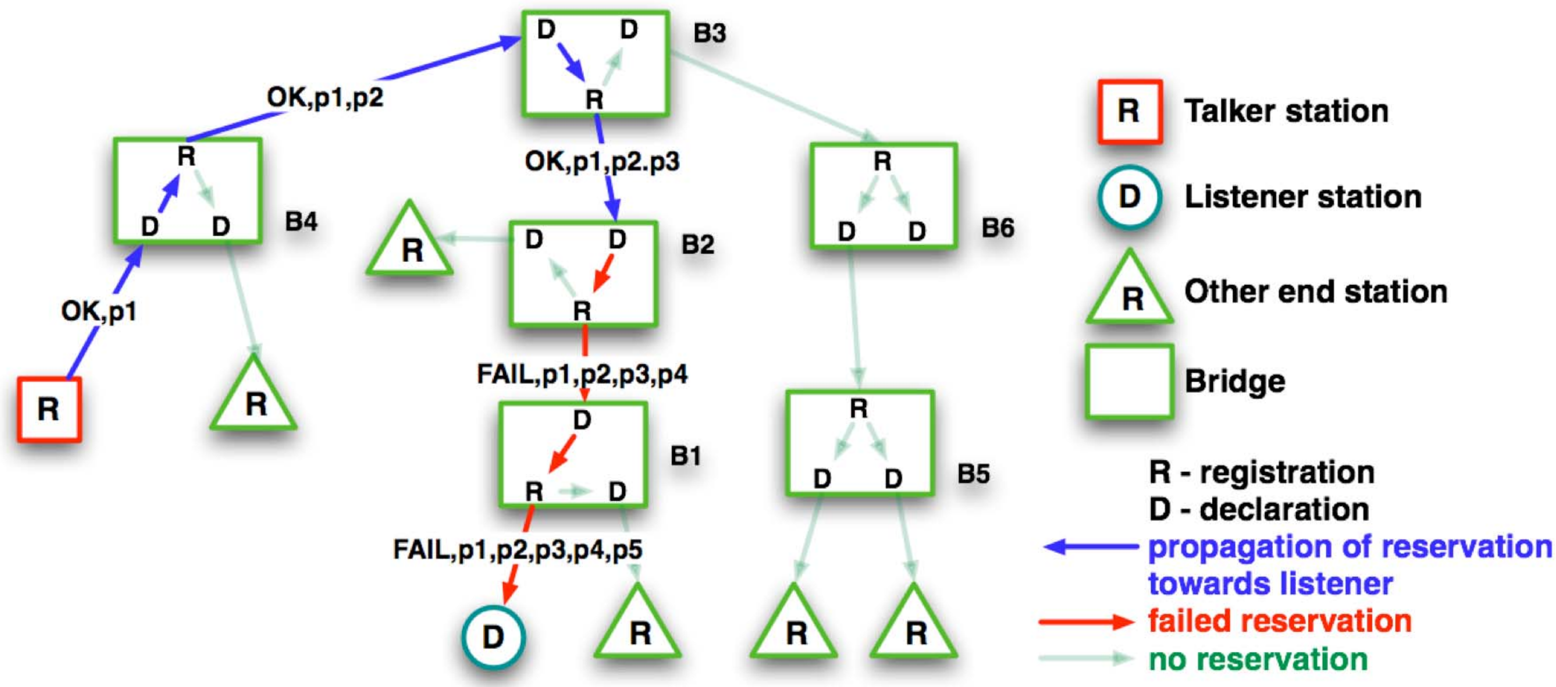

- In this example, admission control fails at B2. The SI (Status Indication) flag of the reservation signaling will be set to FAIL.
- The reservation is still forwarded to the listener. However, down-stream bridges (i.e., B1, B2) will not lock resources for the reservation whose SI is set to FAIL.
- Listener is noted of the failure since a reservation with FAIL SI serves as an end-to-end explicit NACK

BROADCOM
Connecting everything®

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 an explicit "native IEEE 802 layer 2 profile" of new IEEE 1588v2

AV Grand Master clock

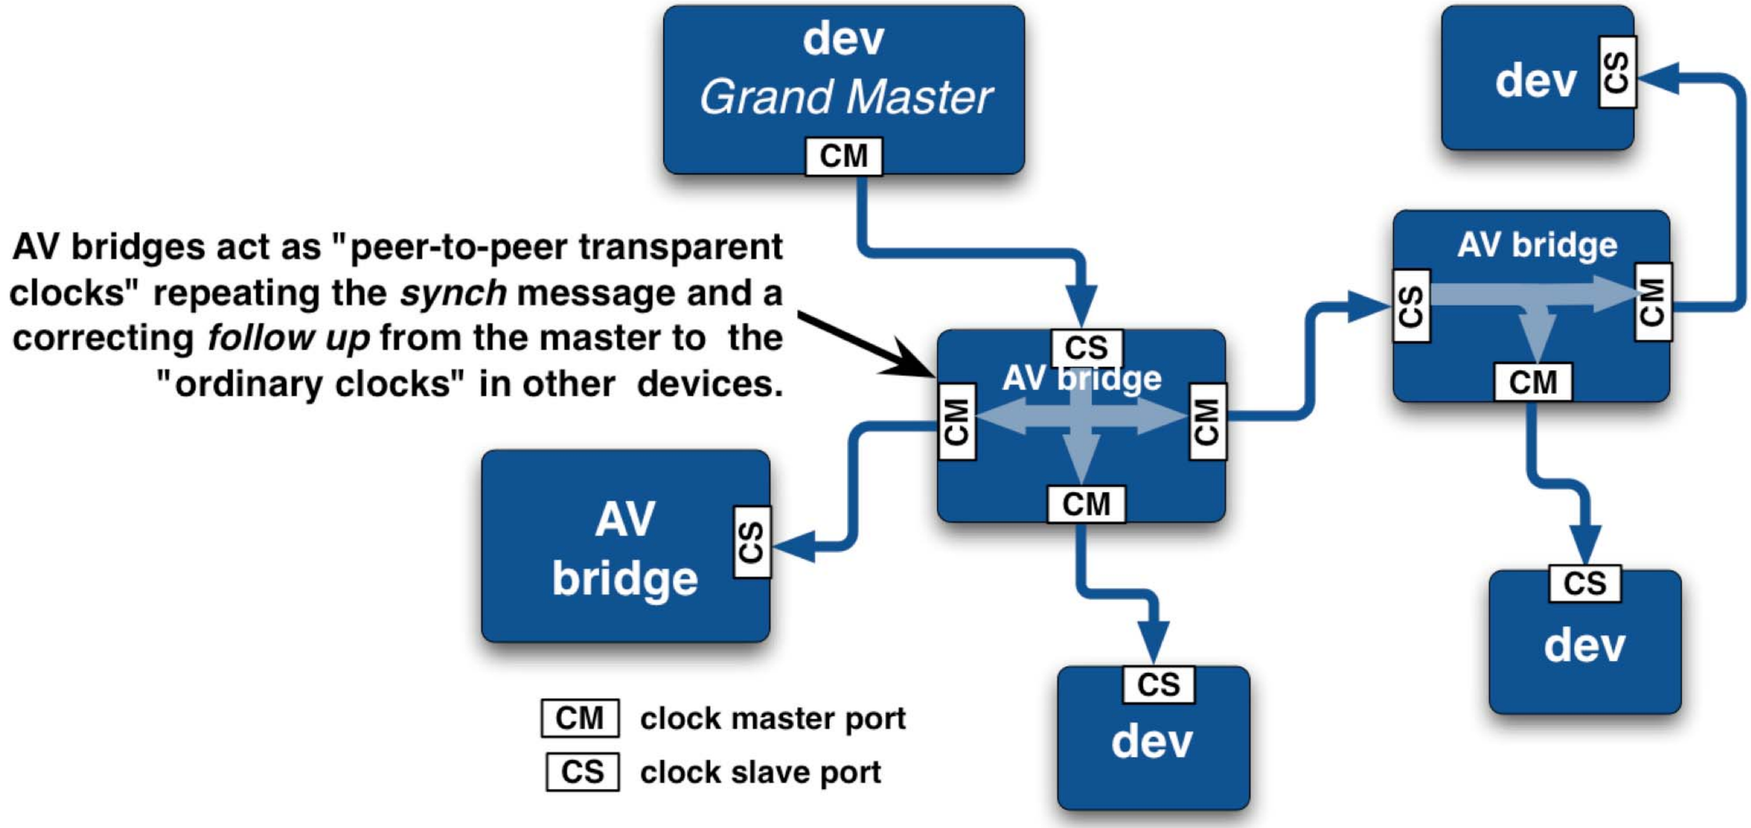

- There is a single device within an Ethernet AV cloud that provides a master timing signal.
  - All other devices ("ordinary clocks") synchronize their clocks with this master.

Master clock selection

- Selection of the master is largely arbitrary (all AV devices that can source data 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
- Selection algorithm and clock attributes are the same as IEEE 1588
  - Typically, fully automatic and transparent to the end user

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.

BROADCOM® Connecting everything®

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”
  - Early drafts already available
  - Expect technical closure in 2007, final draft standards in 2008
- First hardware/software soon after stabilization
  - Possibly a number of “pre standard” interations
- Will follow Ethernet-type product curve
  - 100M/1G/10G NIC/Switch all have markets for Ethernet AV

Example Enet AV NIC

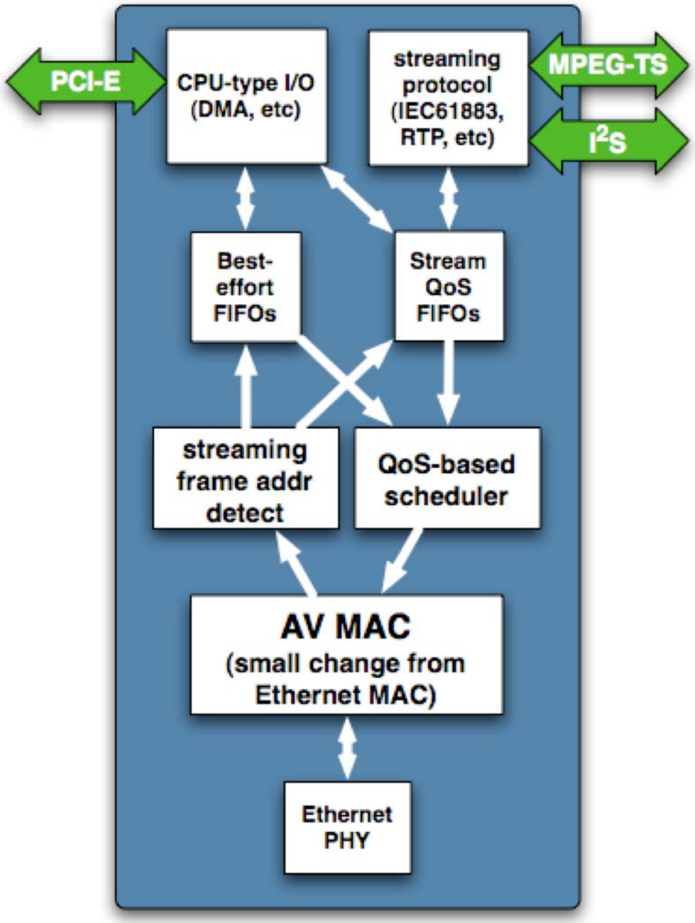

- PHY is Fast Ethernet or better
  - 1G for backbone/professional, 10G for uncompressed video
- CPU interface is PCI-E
  - Streaming frames on PCI-E use virtual channel
  - Perhaps parallel PCI for CE
- Streaming I/O acceleration
  - MPEG transport stream
  - I^2S serial audio
  - DVI for uncompressed video

BROADCOM
Connecting everything

Example Enet AV bridge

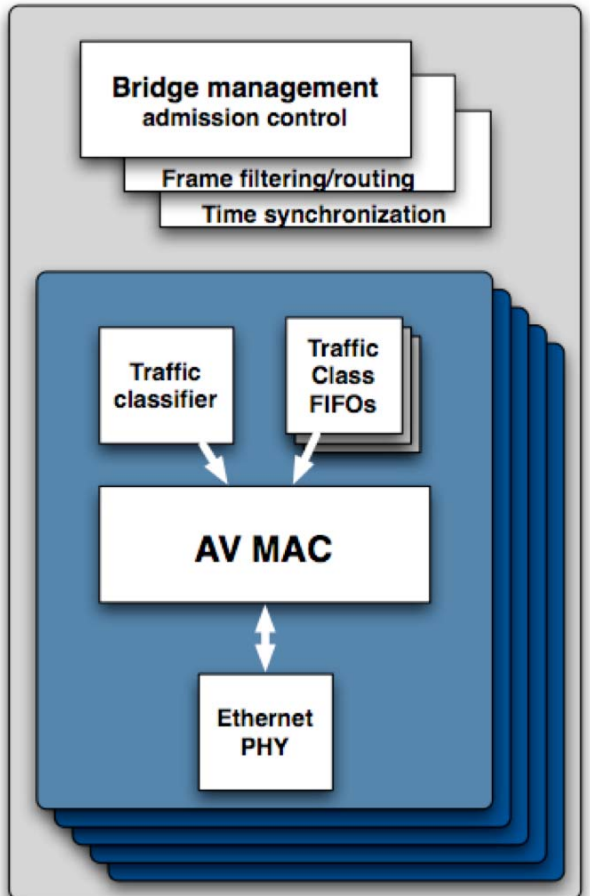

- PHYs are Fast Ethernet or better
  - 1G/5 port for first versions
- Separate CPU at first, but moving to integrated processor ASAP

Example multiport Enet AV NIC

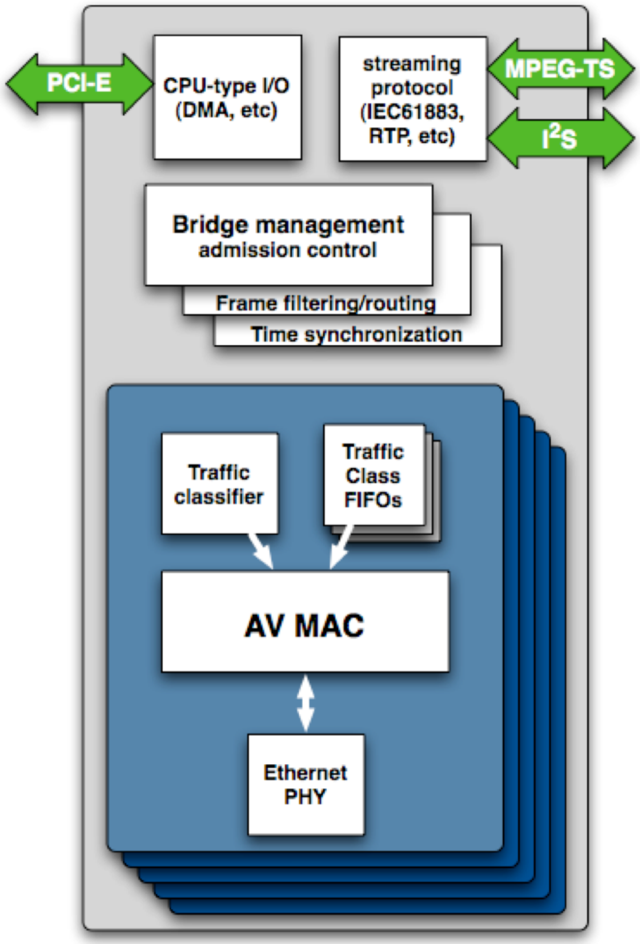

- Best product for TVs, STBs, home gateways, media PCs
- Anything that is a “hub” in the cluster

Key Take Away

- 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

G.8261 (G.pactiming) and 802.1 AV bridging

Michael Johas Teener
mikejt@broadcom.com

G.8261 Timing recovery

- Bridging between SDH/SONET (and similar) networks requires carrying ***both*** the data with adequate QoS ***and*** timing information
  - E.g., an Ethernet network must carry the timing information of one edge SDH/SONET to another.
  - It must also emulate the connection-oriented model for data transport (not lose data from established connections in spite of interfering traffic)

G.8261 Differential Method

- Common reference clock available throughout packet network
- Service clock (at packet network ingress) is encoded using a timestamp with respect to the reference clock and included with data
- Service clock at packet network egress is recovered using reference clock and timestamp information in stream

Using the Differential Method

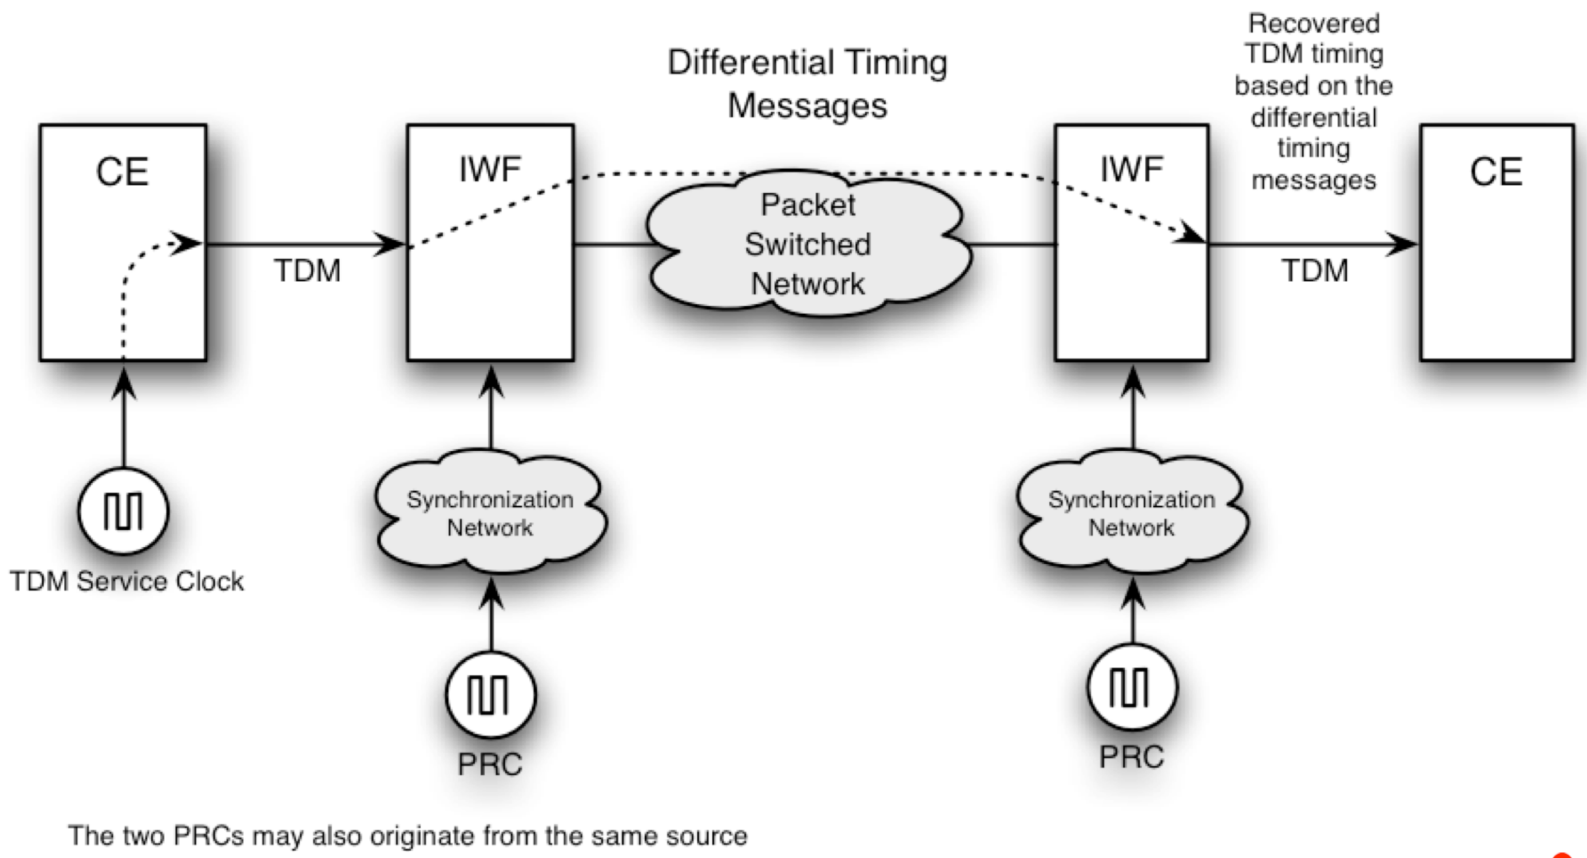

802.1 Audio/Video Bridging

- Provides precise synchronization services
  - 802.1AS - IEEE 1588v2 as applied to 802.1 bridged networks.
  - Probably better MTIE than needed for SDH/SONET
    - (Geoff Garner is currently comparing with TDEV for SDH-1 and -2)
- Provides connection-oriented services
  - 802.1Qat - Stream Reservation Protocol to manage streams
  - 802.1Qav - Guaranteed latency and bandwidth for established streams
- No need to manage PHYs, no need for external PHY-level synchronization
  - Any 802.3 PHY will work, nothing special needed, full configuration flexibility
  - Support for other full-duplex point-to-point PHYs trivial
  - Support for shared media MACs allowed (802.11/16/17) via sublayer definition

Using 802.1 AVB

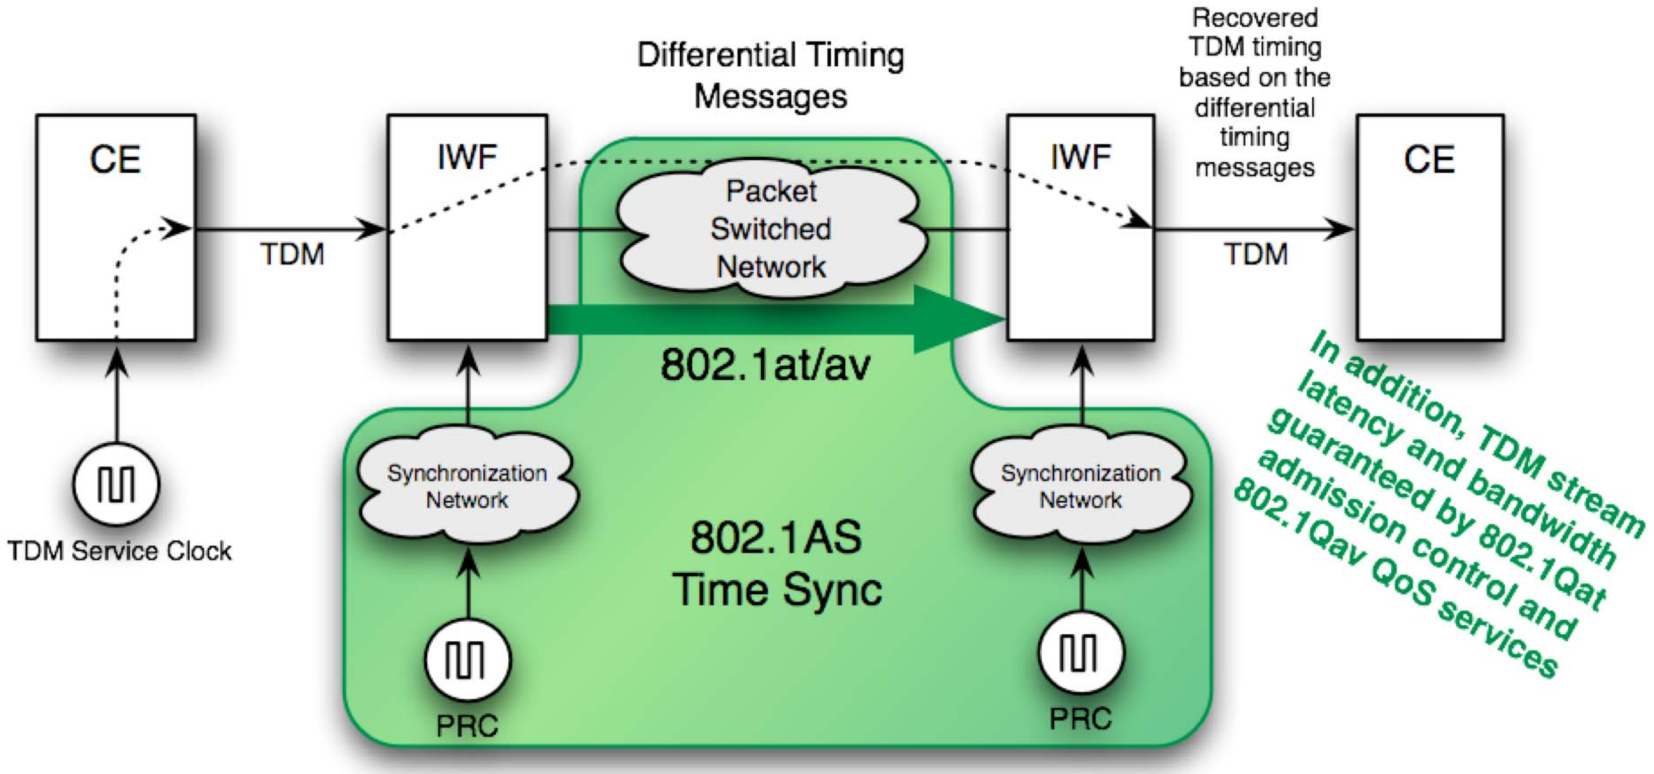

The two PRCs may also originate from the same source

Using 802.1 AVB for multiple independent clocked streams

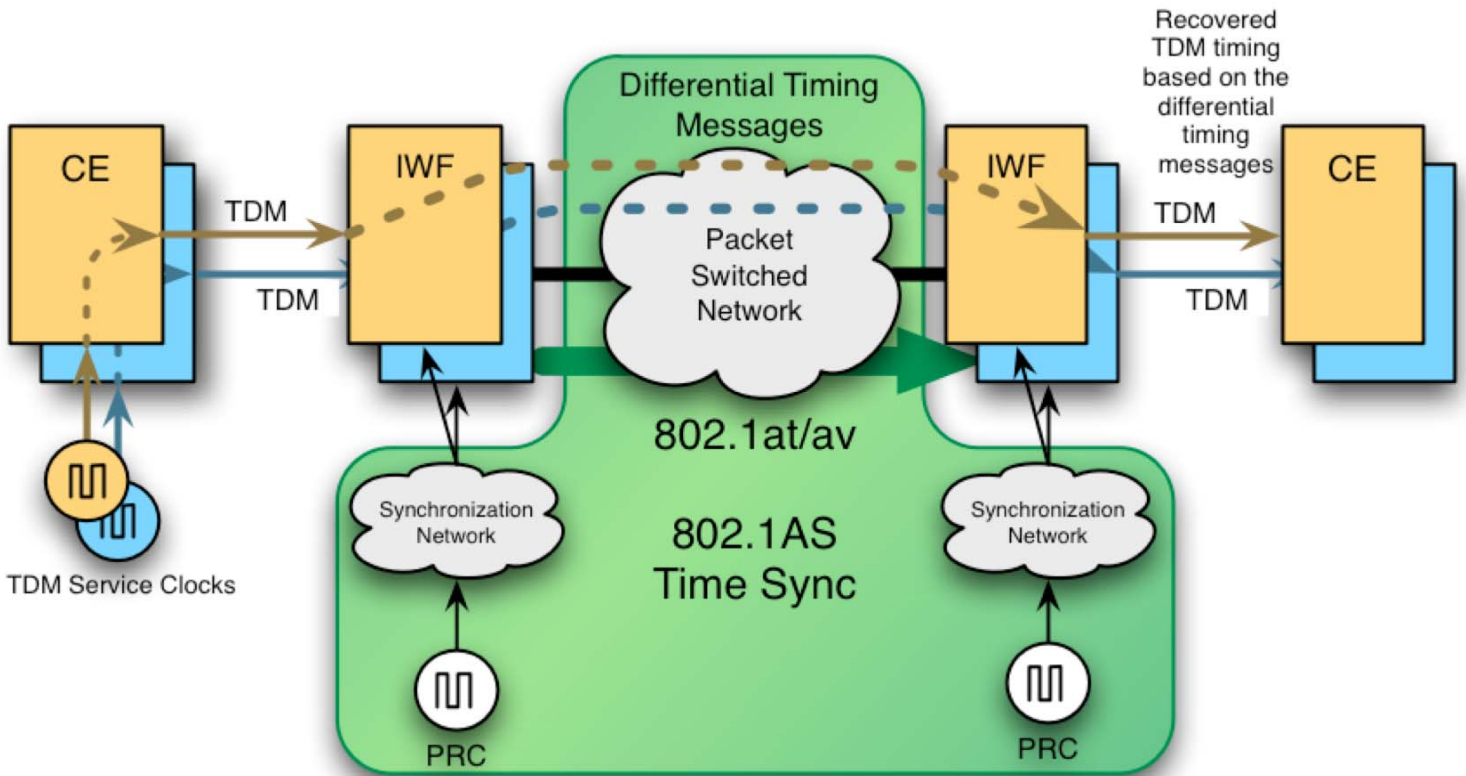

Connecting everything

Advantages of 802.1 AVB

- Will be heavily used in consumer electronics and professional AV networks
  - Driven to be a simple and low cost as possible
- Easily scaled to much larger networks
  - Architects of the 802.1 understand and require scaling
- Supports multiple simultaneous TDM streams with different clocks

New work needed

- Telecom core networks must be managed at a high level
  - Limiting automatic switchover of clock sources and routing of packets
  - “protected environment” for management vs. “plug and play” usage models
- Standardized packet format for TDM emulation
  - Mapping/demapping for timing recovery (timestamp usage)
- 802.1 AVB protocols have all the lower layer tools
  - Higher level management interfaces and overrides need to be defined