Residential Ethernet Tutorial

Michael Johas Teener, Plumblinks

Jim Battaglia, Pioneer

Alexei Beliaev, Gibson Labs

Eric HS Ryu, Samsung Electronics

Yongbum Kim, Broadcom

Agenda

- Introduction
- Digital entertainment networking
- Market potential
- Latency sensitive application examples
- Possible architecture
- Layering possibilities
- ResE standard dependencies to existing standard and their complexities

Residential Ethernet: Introduction

Michael Johas Teener
mike@plumblinks.com

Introduction

- What - Where - Who - Why - When
- Background
  - Why not use Ethernet as is?
  - Why not use 1394?
  - Why not use wireless?
- Performance requirements
- Objectives

What - Where - Who - Why - When

- What is ResE?
  - Enhancement to 802.3 MAC and 802.1D bridges to support very low latency isochronous services
- Where will it be used?
  - Residential networking for entertainment, communication and monitoring
- Who will implement it?
  - Infrastructure and consumer electronics firms
- Why is ResE important?
  - Provides no-excuses QoS to minimize consumer complaints and reduce CE system complexity
- When will it be available?
  - Two years after start. Beginning of 2007.

Background

- Ethernet is very successful as a computer network
  - but not widely adopted for real-time streaming of audio or video
- There are two primary reasons for this:
  - Ethernet does not provide the kind of isochronous and deterministic low-latency services required by these applications, and
  - the market was not large enough to justify the work necessary.
- As a result, other technologies are used
  - either IEEE 1394 (FireWire, iLink, etc.) for the local cluster, or a variety of point-to-point or proprietary systems for longer distances.

Times have changed!

- There is a growing awareness of the need for a network that can distribute and control high quality digital audio and video.
  - Must be a heterogeneous network that supports both wired and wireless components.
- Wireless networks explicitly support AV streams
  - IEEE 802.15.3a/b and 802.11e have scheduled access
- The missing link is the wired component.

Why not use Ethernet as is?

- There are claims that Ethernet can be used without any changes for A/V networks
- But this requires a number of constraints on the network
  - Limit the amount of traffic offered by each device.
  - Limit the topology of the network (number of devices, number of hops).
  - Require the network to have much higher bandwidth capabilities.
  - Limit the use of priority services to only A/V streams.
- These limits make it less likely that any particular queue on the network will overflow and drop packets or add excessively delay.
  - Queues exist in the output ports of all devices, including network switches.

But it still won't work

- **The "over-provisioned, limited topology, priority-based" network only reduces the likelihood of problems.**
- **There is no way of enforcing the constraints.**
  - **A single PC with unusual but perfectly standards-compliant software will cause unacceptable delays and dropped packets.**
  - **Similar problems are caused by "too many" devices or an unexpected configuration of switches or hubs.**

Perhaps use 1394?

- IEEE 1394 was designed from the ground up as an A/V network
  - 1394b and 1394c extend to 100m hops, just like Ethernet
  - 1394.1 allows network scalability just as 802.1D bridges
  - 1394 already bridges to wireless via 802.15.3 “PAL” (protocol adaptation layer)
- Only one problem: 1394 is not Ethernet.
  - Almost all computer applications use Ethernet for the wired network.
  - Very few 1394 home networks are yet deployed.

Perhaps use wireless?

- The average data rate of the proposed 802.11n and UWB nets is OK (particularly at short ranges)
- 802 wireless networks have some A/V services.
  - “protocol adaption layers" exist for 802.15.3a and 802.11e so they can carry 1394 data!
- But the quality of service for wireless is not adequate for HD-quality video.
  - latency is excessive (10’s of milliseconds) in a single attachment point domain, much worse in a mesh
  - ... and normal home environments can result in momentary packet loss
- Power???

Wireless is an import part of the home A/V network, but it is not the sole solution for the backbone.

History: early attempts

- Cirrus Logic “CobraNet” has been used for audio distribution in auditoriums
  - uses the “limited topology” concept
  - one extra switch, or a legal but "non-CobraNet" configuration and quality is poor
- Gibson Guitar “MaGIC” for live performances
  - uses a non-standard bridge as the key bit of technology
  - cannot connect with existing Ethernet

History: current efforts

- Mandated HDTV switchover in USA started in 2004 - will complete in 2008
  - recent success of the digital audio market (e.g., “iPod”)
- July 2004 meeting of the IEEE 802 had a “Call For Interest” for a Residential Ethernet project
  - supported by Pioneer, Nortel, Gibson Guitar, Samsung, Broadcom, NEC and others.
  - The Residential Ethernet Study Group was approved and given the task of setting the objectives for the project and providing the justifications for future work

ResE objectives for 802.3 and 802.1

- **All plug and play features required, not optional**
  - CE industry requirement for minimal options and configuration
- **Admission controls to guarantee path bandwidth**
  - If a stream is started, it must continue to work
- **Isochronous and best-effort traffic will be carried together, with some bandwidth reserved for best-effort**
  - Always need some bandwidth for control
- **Links will be 100Mb/s full duplex or greater**
  - Standard frame on 10Mb is too long (1.2ms) and adds to latency
  - Most common CE stream will be HD video @ 20Mbit/s
  - Full duplex allows bridge-based QoS services without compromise

ResE objectives for 802.3 and 802.1 (2)

- **Isochronous traffic will have less than 2ms end-to-end latency through the entire network and only 250µs through one hop**
  - **Worst case for CE application is musical instrument (see following presentation)**
  - **Seems to be “free” for implementations**
- **Delivered isochronous traffic will have very low jitter and wander approaching zero**
  - **Minimizes buffer and filtering requirements for applications**

ResE objectives for 802.3 and 802.1(3)

- High quality synchronization services will provide all stations with a low jitter "house clock"
  - Applications need a good time stamp source
- Support for all 802.1 services, in particular 801.1Q VLANs
  - ResE will be used in shared housing (e.g., apartment buildings)
- Support arbitrary topologies within reasonable limits (802.1D RSTP)
  - … minimize configurations that "don't work"

ResE objectives for "other groups"

- Isochronous bridging to IEEE 1394
- backbone for IEEE 802.11 and IEEE 802.15.3 such that all QoS parameters are respected
- perhaps even bridging to USB, MOCA, and higher level protocols such as RTP

Compatibility requirements

- Must carry existing Ethernet traffic transparently
  - Any current Ethernet device will connect and provide existing QoS
  - Fully support all IP-based higher level services
- Must carry 1394-type isochronous streams with minimal changes
  - Support IEC 61883 payloads
  - "DCAM" machine vision cameras
- "Good citizen" in 802.1 networks
  - 802.1Q VLAN, standard management information base
- Explicit support for power over Ethernet
  - PoE reporting required

Thank you!

Digital Entertainment Networking

Media-centric Applications on the Home Network

Jim Battaglia

Pioneer

What is it?

- Sharing of digital audio, video, and images across a home network
- Accessing and controlling digital devices and content through multiple points within the home
- Connecting in-home entertainment devices with traditional and non-traditional content delivery services
  - On-line Music Services (download, subscription, Internet/Satellite radio)
  - Cable, Satellite, Terrestrial Broadcast TV
  - IPTV
  - Meta-data services (e.g. EPG)
- Protecting the rights of consumers and content owners in the digital age (CP/DRM)

Media Distribution in the Digital Home

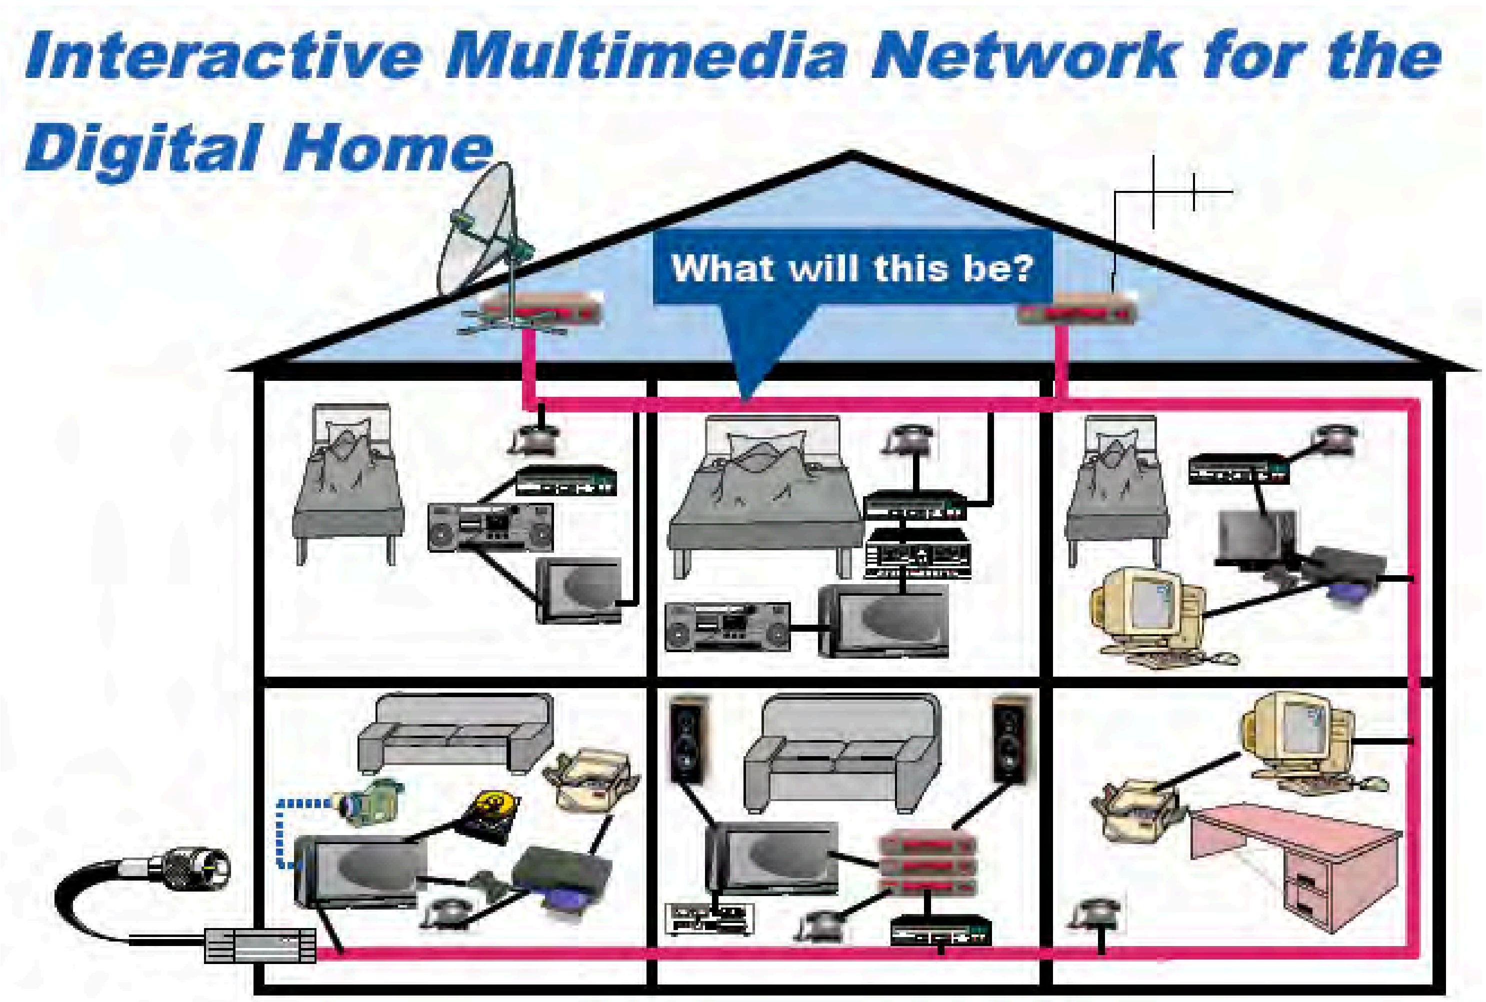

DEN Applications

- Sharing of content anytime, anywhere
  - Personal content
    - photos, home movies, personal music recordings
  - Commercial content
    - Live TV delivered through cable, satellite, terrestrial broadcast
    - VOD and PPV
    - Internet audio services: radio, subscription, downloads
    - Internet video services: streaming, downloads
    - IPTV
  - EPG and other meta-data
- Examples
  - Photos stored on my PC are viewed on my DTV
  - Movie recorded on my Living Room DVR is played back on my bedroom DTV (Shared PVR could be the "killer app")

DEN Typical Implementation Scenarios

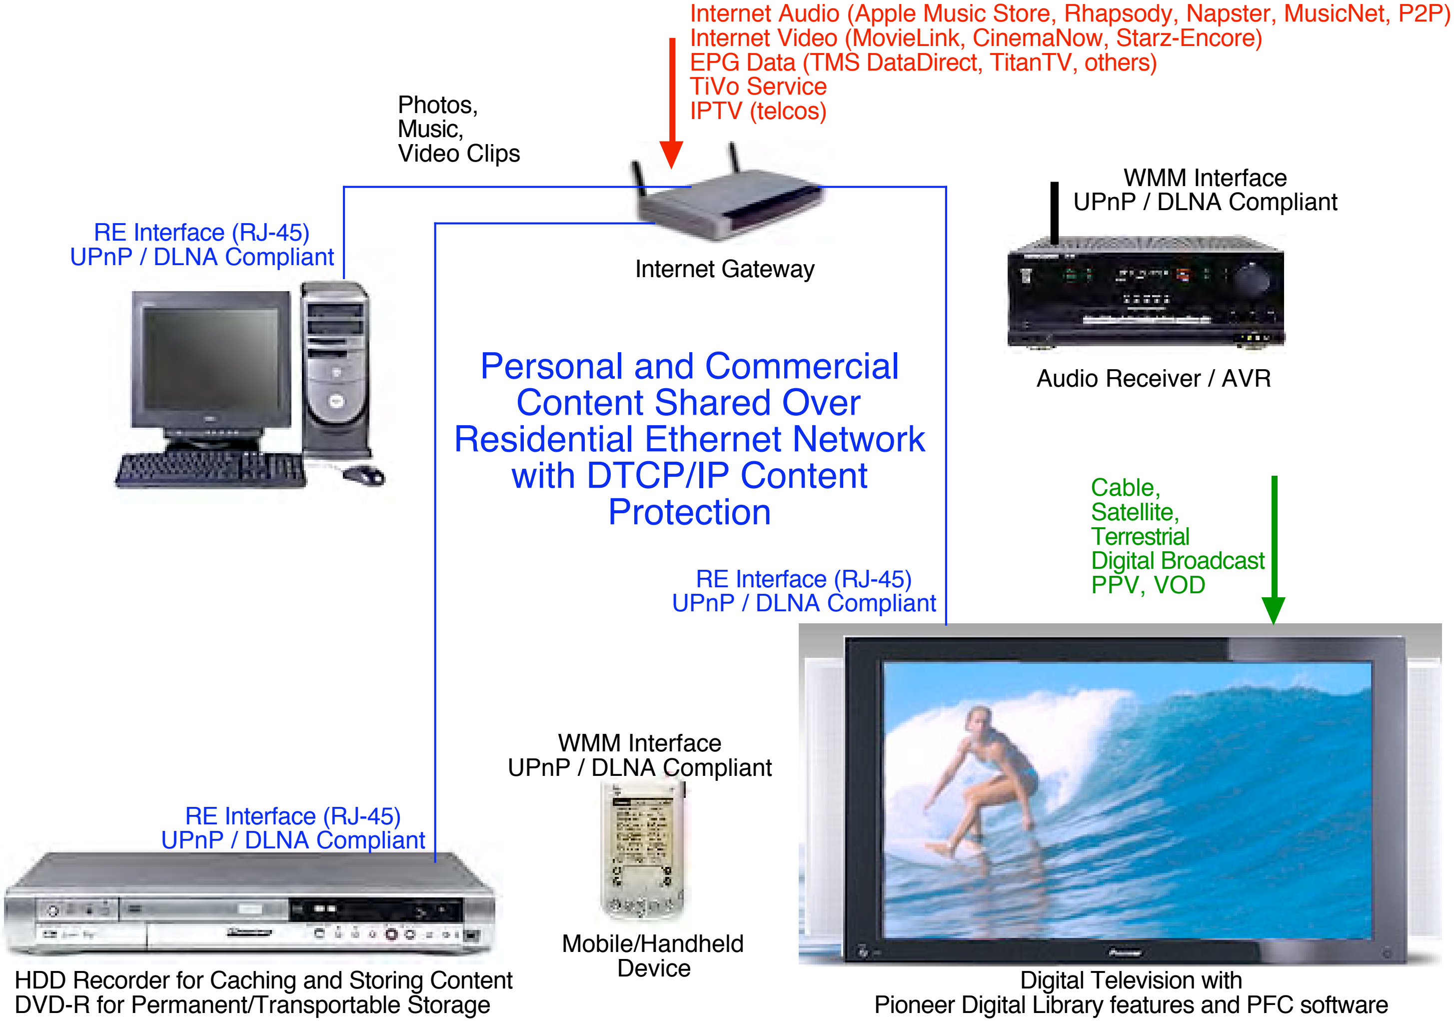

U.S. Market Trends

- Over 10 million households already have home networks installed - mostly for “data applications”
- Over 30 million homes have broadband Internet access
- Almost half of all new homes built include CAT-5 wiring
- Over 3 million subscribers to PVR services
- Over 40% of homes now receive digital television content
- Major ISP’s/telcos rolling out IPTV services

Growth in Digital TV, Broadband, VOD, and DVR Will Facilitate DEN Deployment and Enhance its Benefits

Source: JupiterResearch analysis, 4/04 (US only)

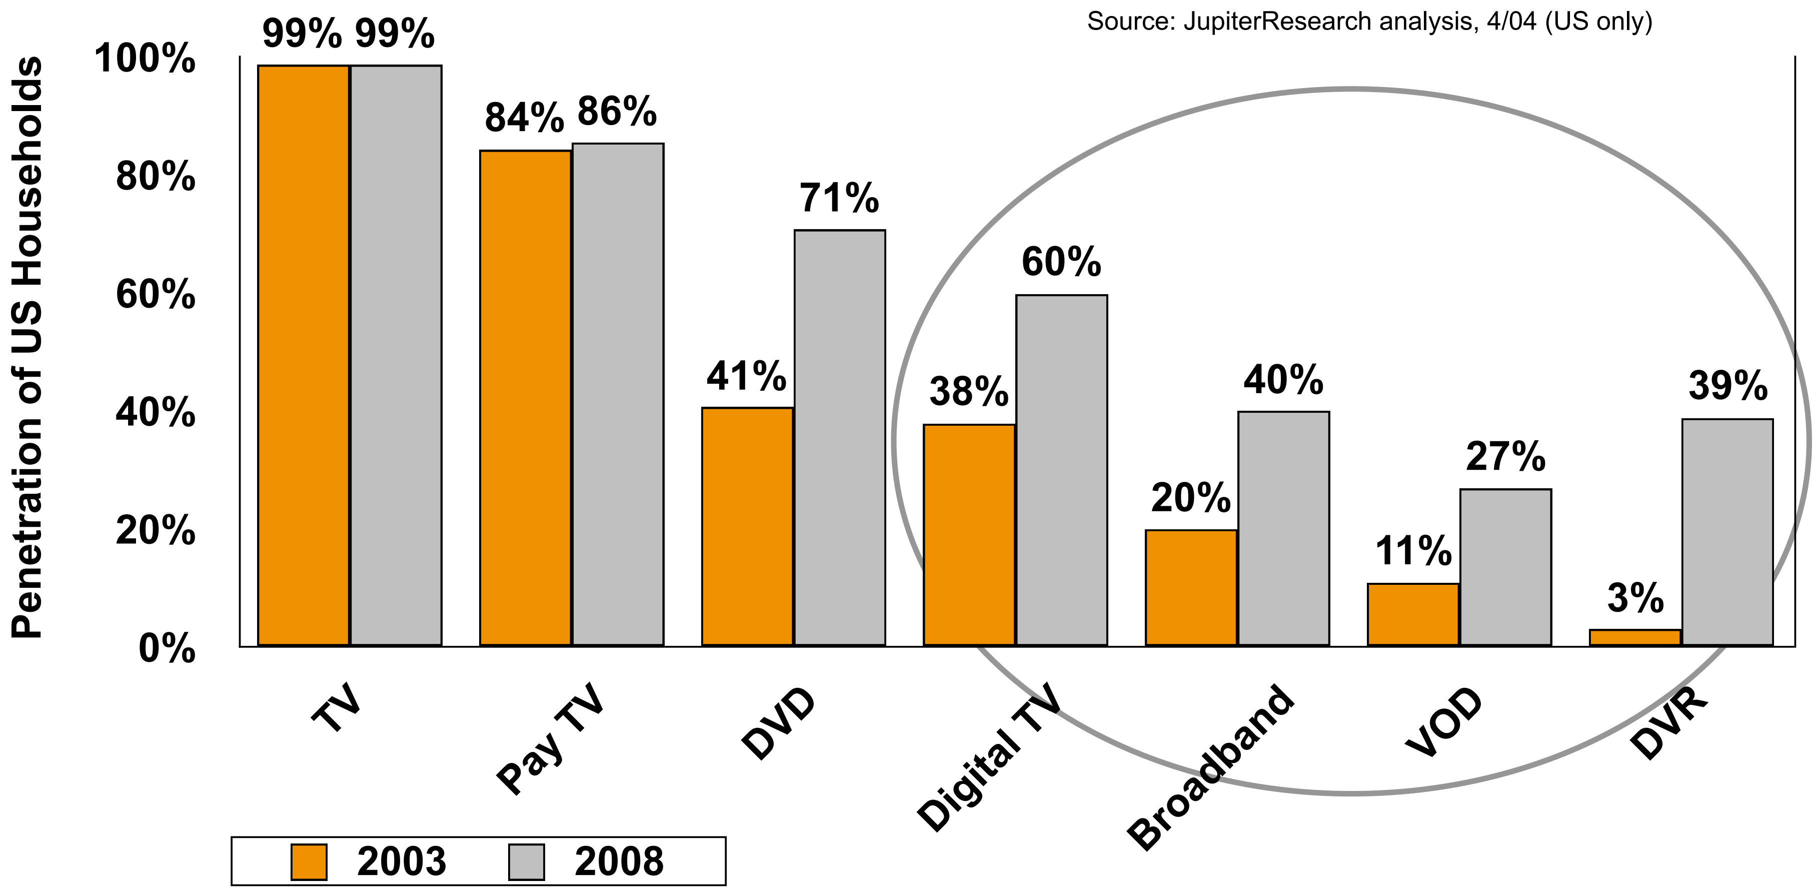

So Why Do We Need *Residential* Ethernet?

- New media-centric applications for digital entertainment impose additional performance requirements on the home network
- Examples
  - Very high quality of service (QoS)
  - Low and reliable latency
  - Little or no system administration by the end-user
  - "Seamless" support for multiple physical media (TP, wireless, coax)

CE Goals for Residential Ethernet

- **Add the following capabilities to Ethernet home networks**
  - Scheduled access
  - Guaranteed bandwidth reservation
  - Low, deterministic latency sufficient for high quality audio/video distribution
  - Network responsiveness sufficient to support compelling user experience
  - Quality of service at least as good as that provided by the content service provider
- **Deliver the above within the following constraints**
  - no changes to PHY silicon
  - minimal changes to the MAC silicon
  - support of 100BaseT (and higher speed) networks
  - complete backwards compatibility
  - virtually zero cost increment for product manufacturers

Why is Ethernet / UTP-5 an important part of the home networking puzzle?

- **General**
  - Hard-wired solutions are less vulnerable to interference, security, and coverage problems
  - Ethernet/UTP-5 has good bandwidth / cost performance
  - UTP-5 is inexpensive and easy to install
  - Ethernet is ubiquitous in other environments
  - Most other physical media used in a home networking environment have Ethernet/UTP-5 connections, e.g. wireless access points, cable/DSL modems
- **US Market Example**
  - Over 10 million homes already have home networking deployed - most of these have Ethernet/UTP-5 backbones
  - Almost half of all new homes being built are deploying “structured wiring” which includes “home-run” UTP-5 wiring (star topology)
- **Some regions outside of the US have even greater UTP-5 penetration**
  - E.g. Korea
- **But, we should assume they’ll be other physical media present in the home networking environment**
  - Wireless, coax, fiber, others

Other Industry Initiatives Advancing Digital Entertainment Networking

- **DLNA: Digital Living Network Alliance**
  - Product interoperability: communication standards, media formats, certification, content protection
- **UPnP: Universal Plug-N-Play**
  - Device discovery and control, media management
- **CEA (R7)**
  - Remote User Interfaces, standardized EPG access over IP networks
- **WiFi**
  - WMM
- **MoCA**
  - Media networking over coaxial cabling
- **DTLA, Coral Consortium, Others**
  - Content Protection, DRM Interoperability
- **CableHome, DSLHome, DBS IPTV Forum**
  - Deployment of commercial media services to a home networking environment
- **All these initiatives could benefit from the services provided by Residential Ethernet**

Residential Ethernet Works Within the DLNA Interoperability Framework

Content Sharing Framework

- Media Formats (Images, Audio, AV)
- Media Transport (HTTP)
- Media Management (UPnP AV)
- Device Discovery & Control (UPnP Device Arch)
- Networking & Connectivity (IPv4, Ethernet, 802.11)

- **Complete set of components to deliver user experience for sharing content**
- **How media content is encoded and identified for interoperability**
- **How media content is transferred**
- **How media content is identified, managed, and distributed**
- **How devices discover and control each other**
- **How devices physically connect together and communicate**

Conclusions

- **Digital entertainment networking is here and now for the CE industry and consumers.**
- **As this new technology is deployed, it is imperative that we meet consumer's expectation for quality of service and ease-of-use.**
  - **Right from the beginning!**
  - **No excuses!**

Residential Ethernet can help us fully realize the benefits of home entertainment in the digital age.

Thank you!

Market Potential

Eric HS Ryu
Samsung Electronics
eric_ryu@samsung.com

New market opportunity

~ 2004	2005~2006	2007~
HA, Internet access	Interactive multimedia service	Intelligent converged service

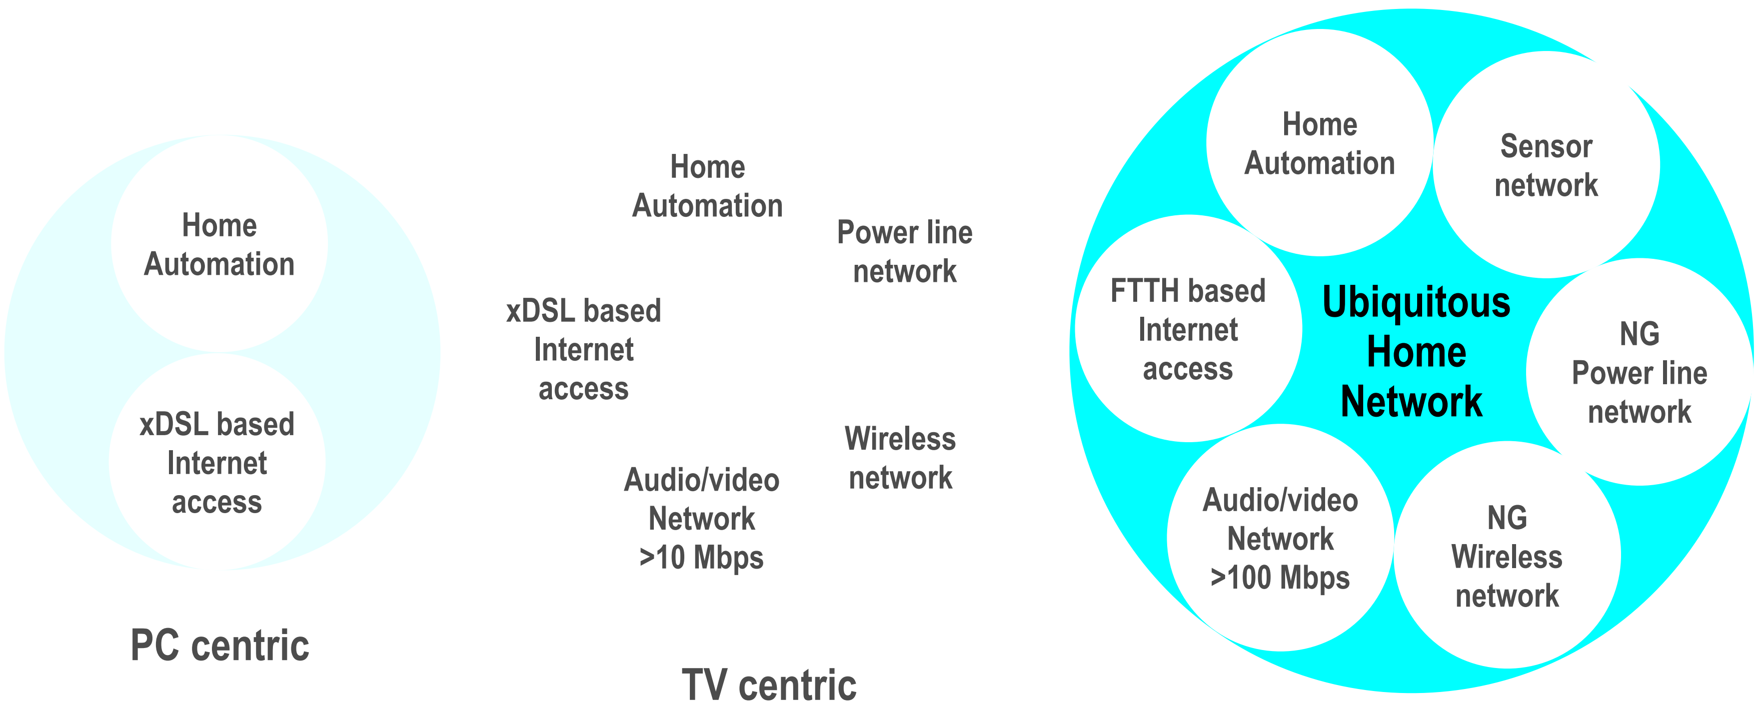

PC centric	TV centric	Converged platform centric

Network based home applications are now creating new market segments.

The home network is at the heart of the digital home

- **Enabling home entertainment networks is key to opening up several large potential markets.**

Residential Ethernet enables a converged connectivity solution

	Access network	**Home gateway**	**Home network**	**Internet information terminal**	**Home network service**
Concept	Connectivity with internet outside home	Interface between access and home network	Connectivity between terminals	Terminal with function to share information with external network	Various kinds of service through home network terminal
Related tech. And product	-xDSL -Cable -Power line -Wireless -Satellite -FTTH	-HaVi -OSGi -UPnP -Jini -VESA-HN -LonWorks Cebus -DLNA	-Phone line -Ethernet -Power line -WLAN -UWB -Bluetooth -RF, IrDA -1394	-Web phone -Communication device -Entertainment device -Computing device	-Home automation -Entertainment service -Telephony service -Health, welfare service

(from 3rd party report)

Why Residential Ethernet ?

- ***Residential Ethernet*** **adds ability to legacy Ethernet to provide isochronous delivery of A/V signals at the attractive cost/ performance ratio of Ethernet**
  - **Most different type of connectivity solutions in residential area can be serviced through Residential Ethernet**
  - **Such as analog/digital video link, analog audio links, conventional Ethernet, UWB, 802.11, USB, IEEE 1394, etc**
- ***Residential Ethernet*** **scales from medium (FE) to high bandwidth (10GE) PHYs**
  - **Display, MP3, PC, Audio, Digital TV, Set-top box, access point, media player (DVD, ng-DVD), game-box, etc**
- ***Residential Ethernet*** **offers cost effective, interference free, backbone for hybrid home network supporting asynchronous data and isochronous A/V signals**
  - **Wired connections between stationary devices in different rooms**
  - **Wireless connections to mobile devices**

Potential applications area of Residential Ethernet

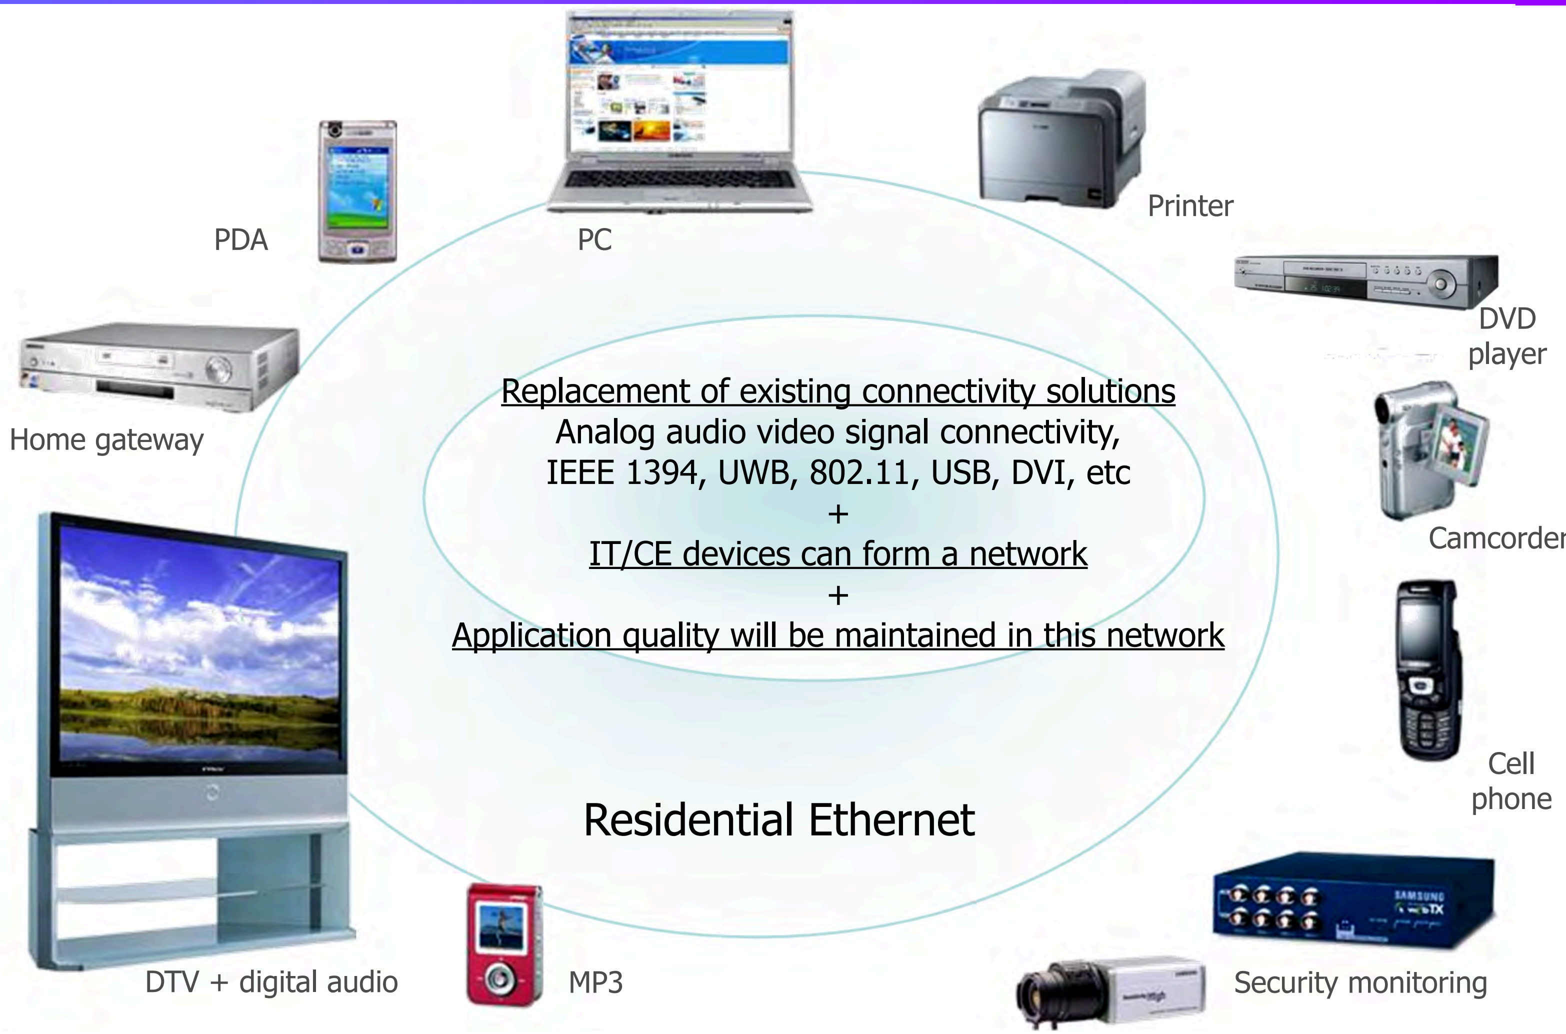

A future digital home networking model

- ❑ Multiple services in the home
- ❑ Home appliances and information terminals are connected by wired/wireless networks.
- ❑ The home network is connected to the access network through a home gateway

❑Enhanced functionality of ResE will replace the connectivity solution of home applications

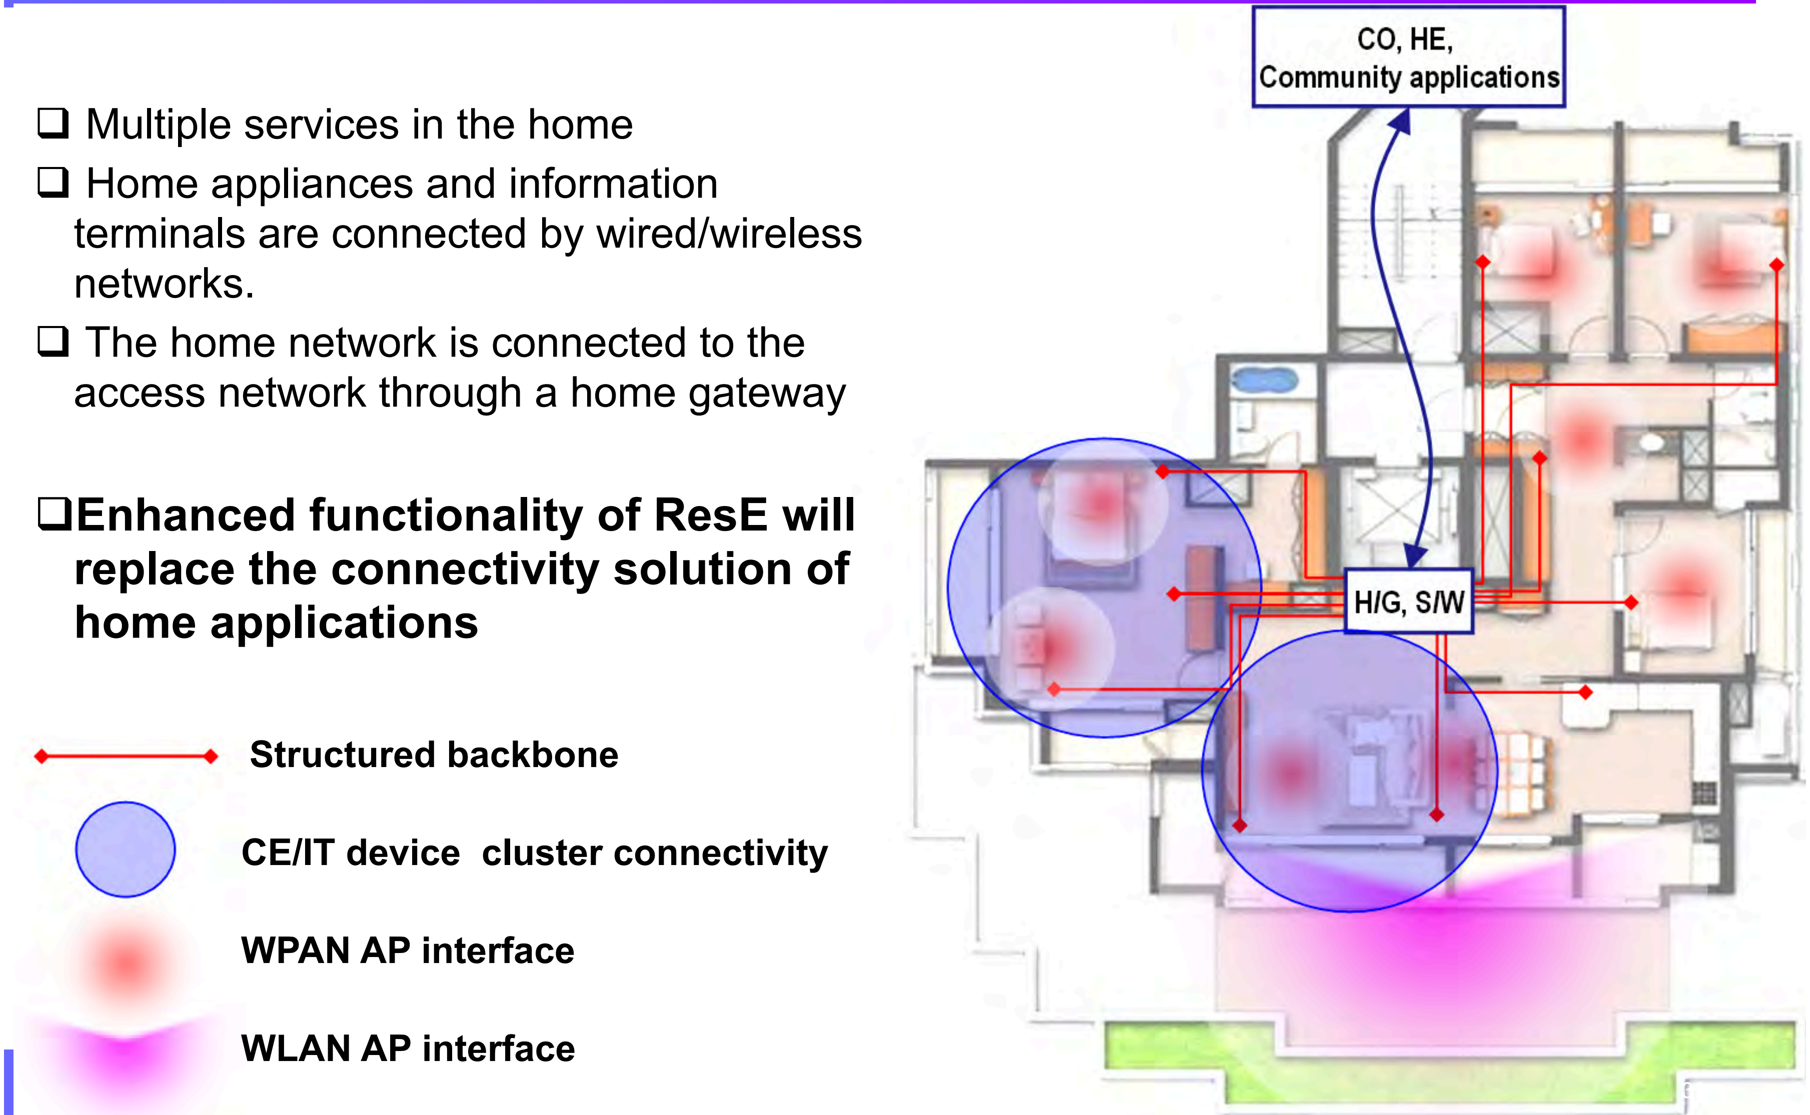

Global home network market potential

- 22 million houses are equipped the home network on year 2003 (Global)
- That will increased to 130 million until year 2008. (43% CAGR)
- This network installation will shift the PC based network to entertainment oriented network.
- Many homes already have structured CAT5 wiring in place;
  - New homes come with structured wiring installed
    - 10 million houses at 2008 in US according (2004 research data)*
- The market size of the digital home (wired/ wireless network, home gateway, portable/ fixed terminal, consumer electronics) is growing rapidly at 18% a year (2002 research data)*

**: The numbers based on 3rd Party Research. Available the reference information by request*

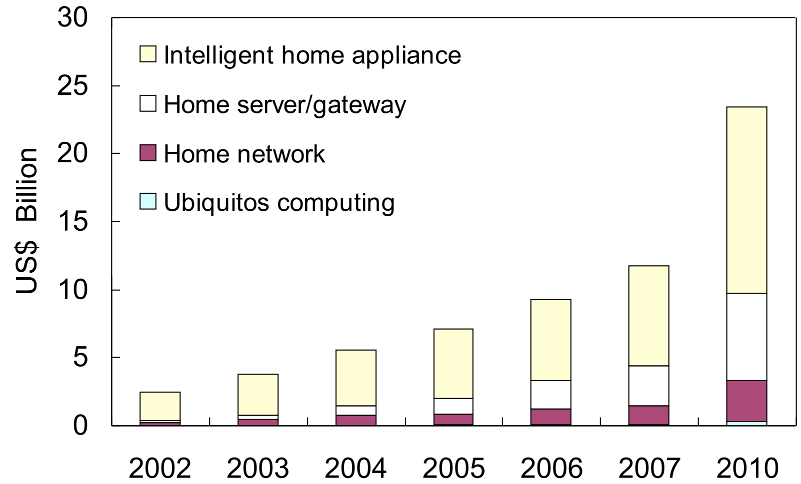

Reference: 2002~2003 research data*

	2002	2003	2004	2005	2006	2007	CAGR
Global	40.7	51.8	63.8	76.8	89.4	102.6	20%
Korea	2.5	3.8	5.6	7.1	9.3	11.8	36%

Global intelligent home networking market forecasting
Reference: 2002~2003 research data*

Broad Market Potential

- **Broad set(s) of applications**
- **Multiple vendors, multiple users**
- **Balanced cost (LAN vs. attached stations)**

"Residential Ethernet" networks represent a new and very broad application space for Ethernet. The digital networking port* on consumer electronics (96 billion USD in 2003) equipment has not yet been decided, and 802.3/Ethernet has a strong possibility of being the dominant, long-term solution of choice if it also provides isochronous services.

At the RE Study Group meetings, individuals from companies representing component suppliers, equipment vendors and users expressed their support for the project. Ethernet equipment vendors and customers are able to achieve an optimal cost balance between the network infrastructure components and the attached stations.

***NOTE: 174 million ports in 2004; 2008 - 458 million; growth rate 21%, 50/50 wireless/ wireline (3rd. Party Research)**

Digital home networking in Korea

- **Korea government has set goals for the future digital home enhancement.**
  - **Digital home project (started July 2003) : more than 10 million digital homes (61%) by 2007**
  - **Smart home project (started July 2003)**
  - **Cyber building regulations Broadband Certification (Started Nov. 2000)**
- **Communication, CE, construction industry, broadcasting companies have formed several consortiums to meet the goal of the projects**
  - **KT consortium : KT, KTF, Samsung, KBS, Banks, etc.**
  - **SKT consortium : SKT, Hanaro telecom, LG, Daewoo, SBS, etc.**

Recent Ethernet Infra

Cyber Building Grade	Specification				Main Use Case
	2003.11 Version			Early version	
	Wire type/Topology	Ethernet Wall Plug*	External Connection spec		
Super grade	Better than Cat5e/ Star	2 x 2 pairs in each room	SMF 2 cores MMF 2 cores Total 4 cores	Better than 1st grade	Multimedia networking (+ wireless infra)
1st grade	Better than Cat5e/ Upgradeable to STAR	4 pairs in living room 2 paris in other rooms	Cat5e 4 pair x 2	100Mbps 2 ports per room	Data communication, control, and sensing
2nd grade	Better than Cat5e/ Upgradeable to STAR	2 pairs in every room	Cat5e 4 pairs	100Mbps 1 port per room	Data communication
3rd grade	Better than Cat3/ Upgradeable to STAR	2 pairs in living room 1 pair in other rooms	Cat3 4 pairs	10Mbps 1 port per room	Data communication

Note: * RJ45 or optical connector

More than 1 million houses are certified as a cyber building by 2003 (more than 10% of all house)

Summary

- *Residential Ethernet* networks represent a new and very broad application space for Ethernet.
  - Combines data networking, CE device A/V connectivity and Wireless AP interconnection in a single network
  - Increasing availability of structured wiring in homes will help to make *Residential Ethernet* a mainstream home networking solution; therefore the related market will grow rapidly
  - With isochronous traffic support Residential Ethernet will become the unified connectivity solution for CE A/V and IT devices
  - The technology will expand its application area to similar environments such as car, transportation, small office, etc.; therefore the market will grow even faster

Thank you!

Latency sensitive application examples

Alexei Beliaev, Gibson Labs

Playing music requires timely feedback

Comfort music playing requires delay to be no more than 10 - 15 millisecond*

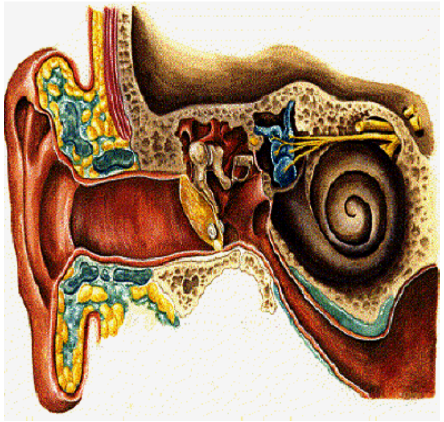

Home recording

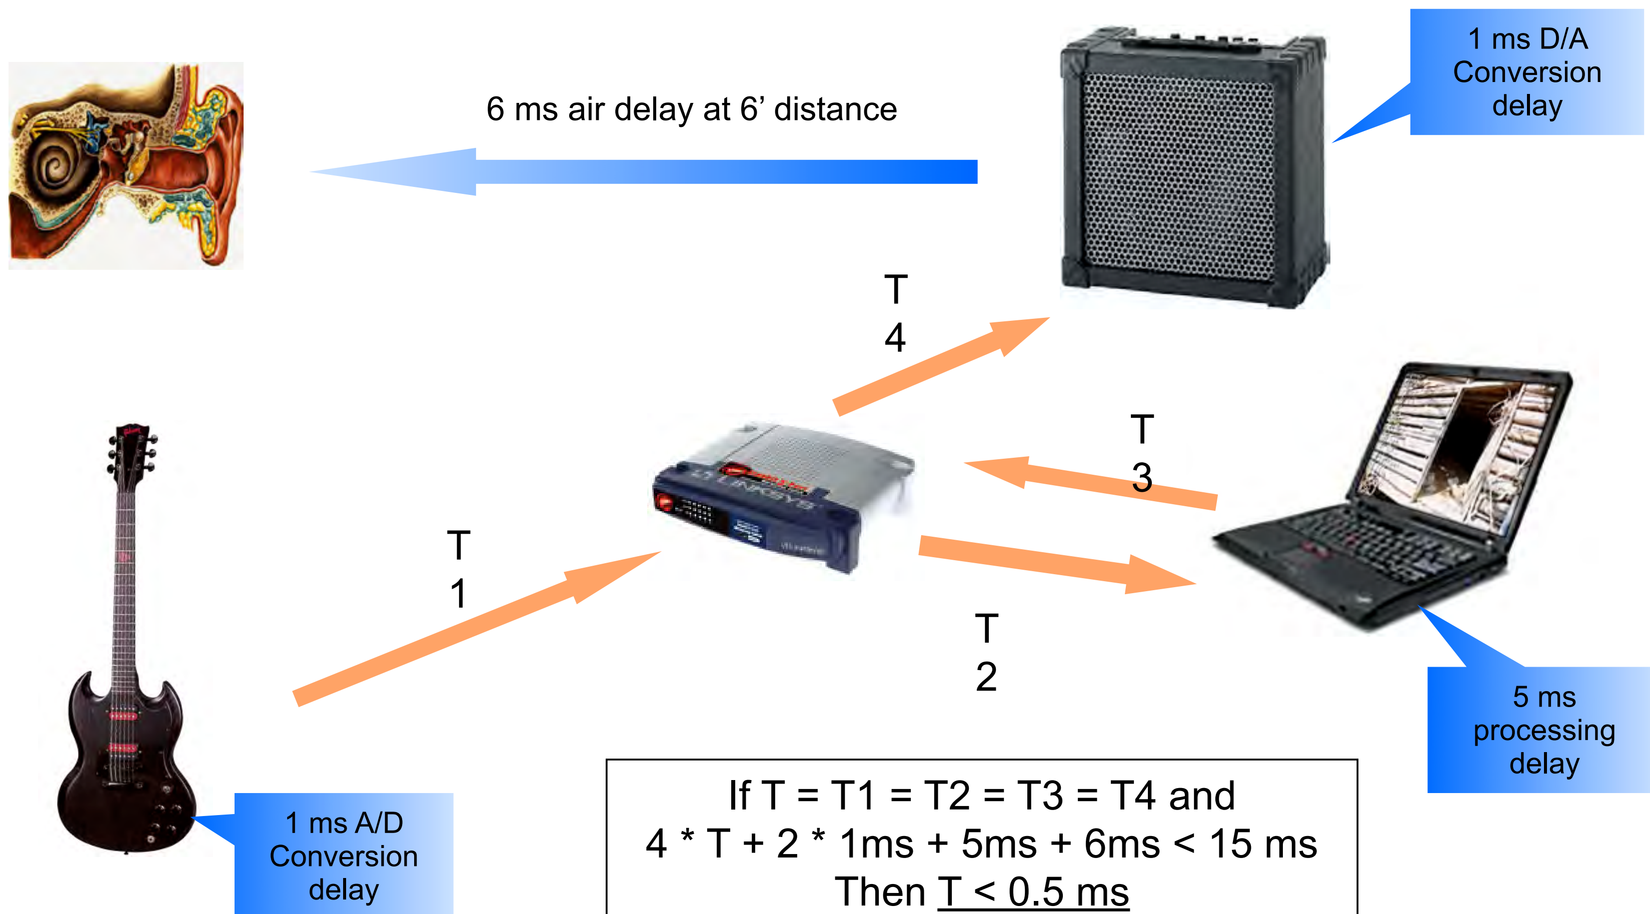

If T = T1 = T2 = T3 = T4 and
4 * T + 2 * 1ms + 5ms + 6ms < 15 ms
Then <u>T < 0.5 ms</u>

Garage jam session

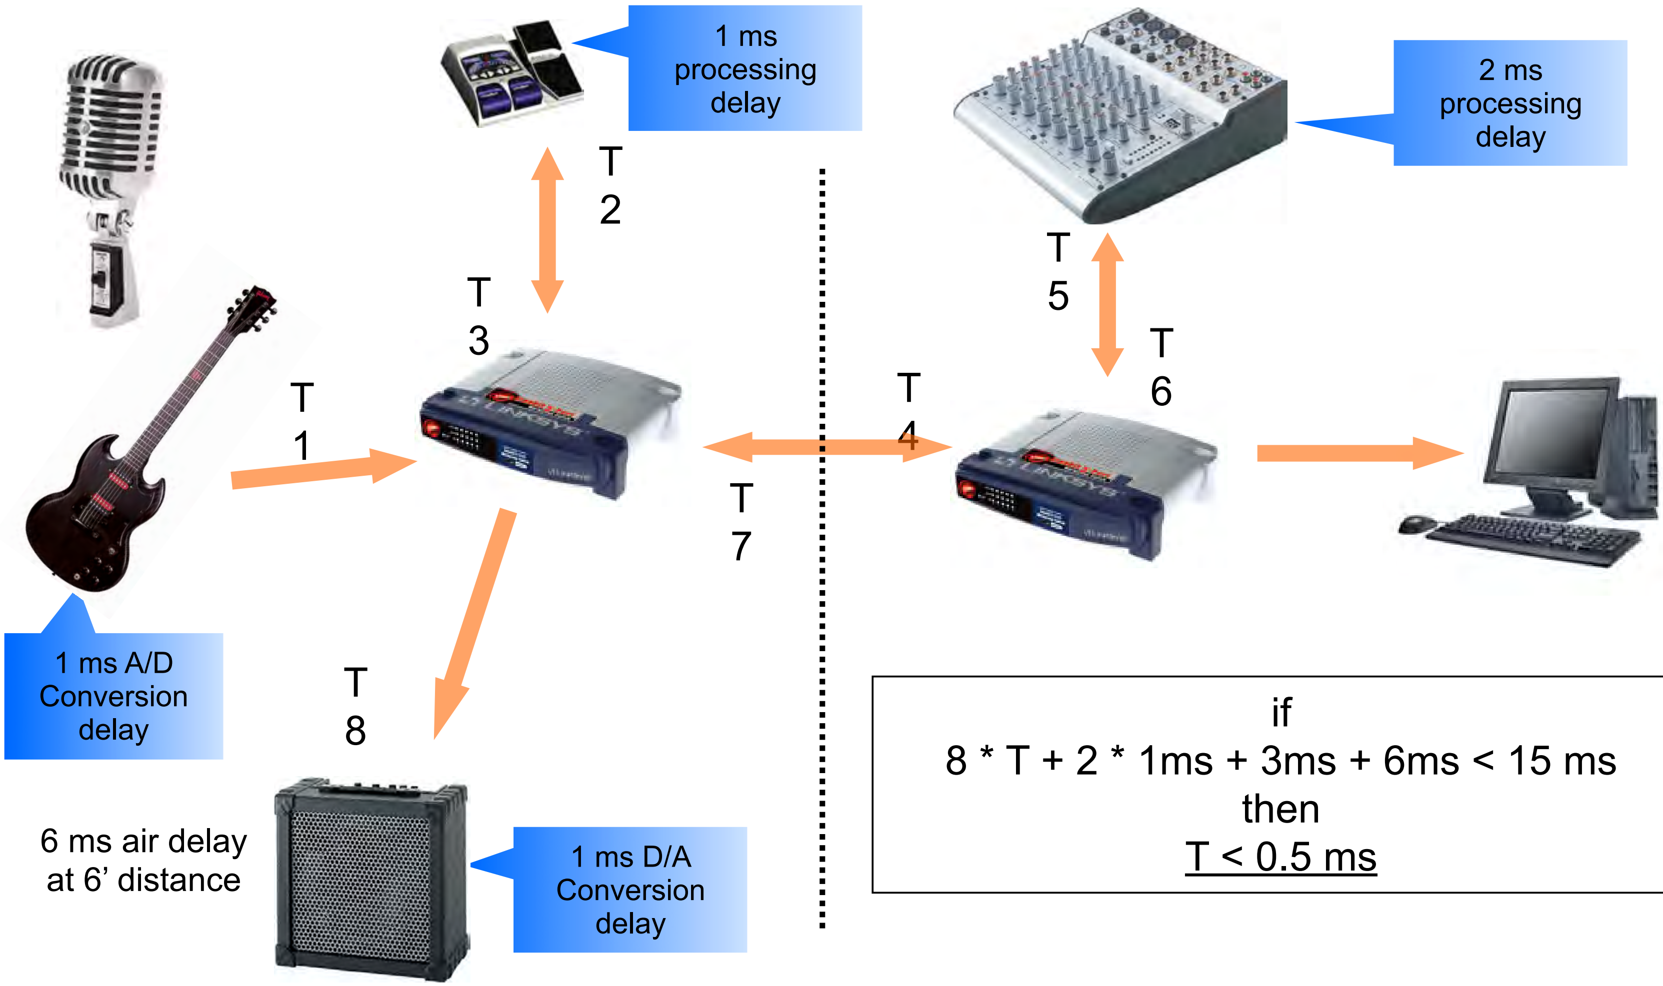

if
8 * T + 2 * 1ms + 3ms + 6ms < 15 ms
then
T < 0.5 ms

Glitch-free connection of nodes

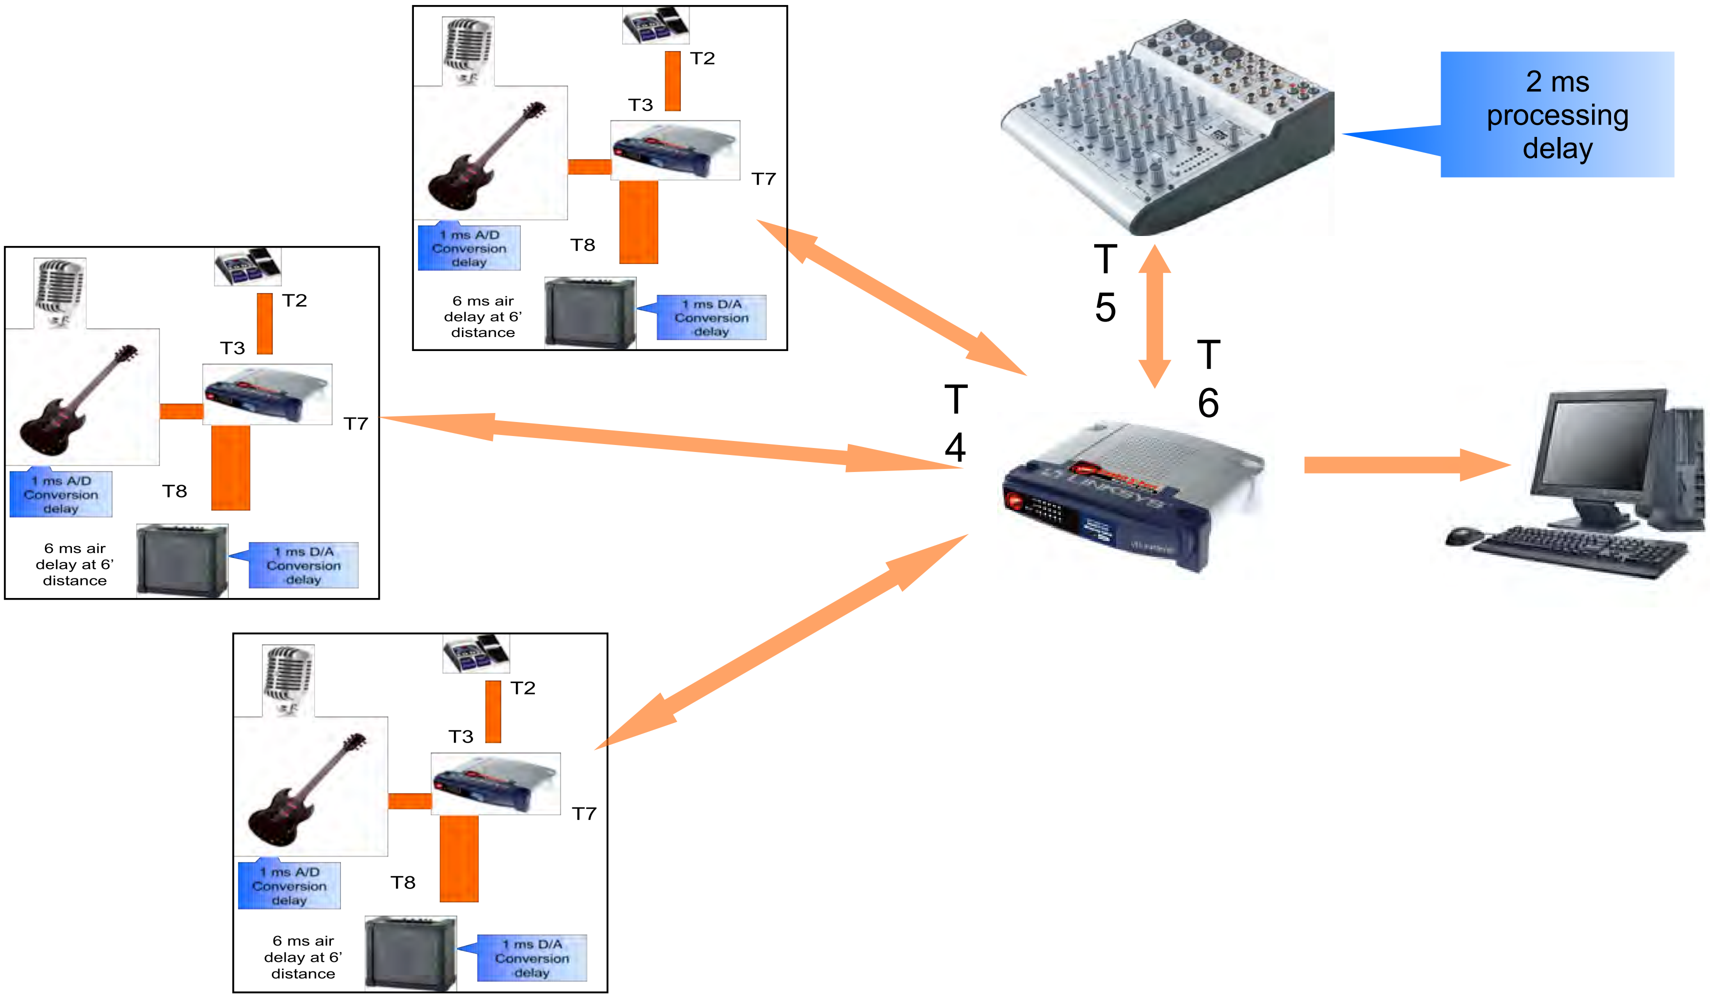

Bandwidth reservation
(no interference between applications)

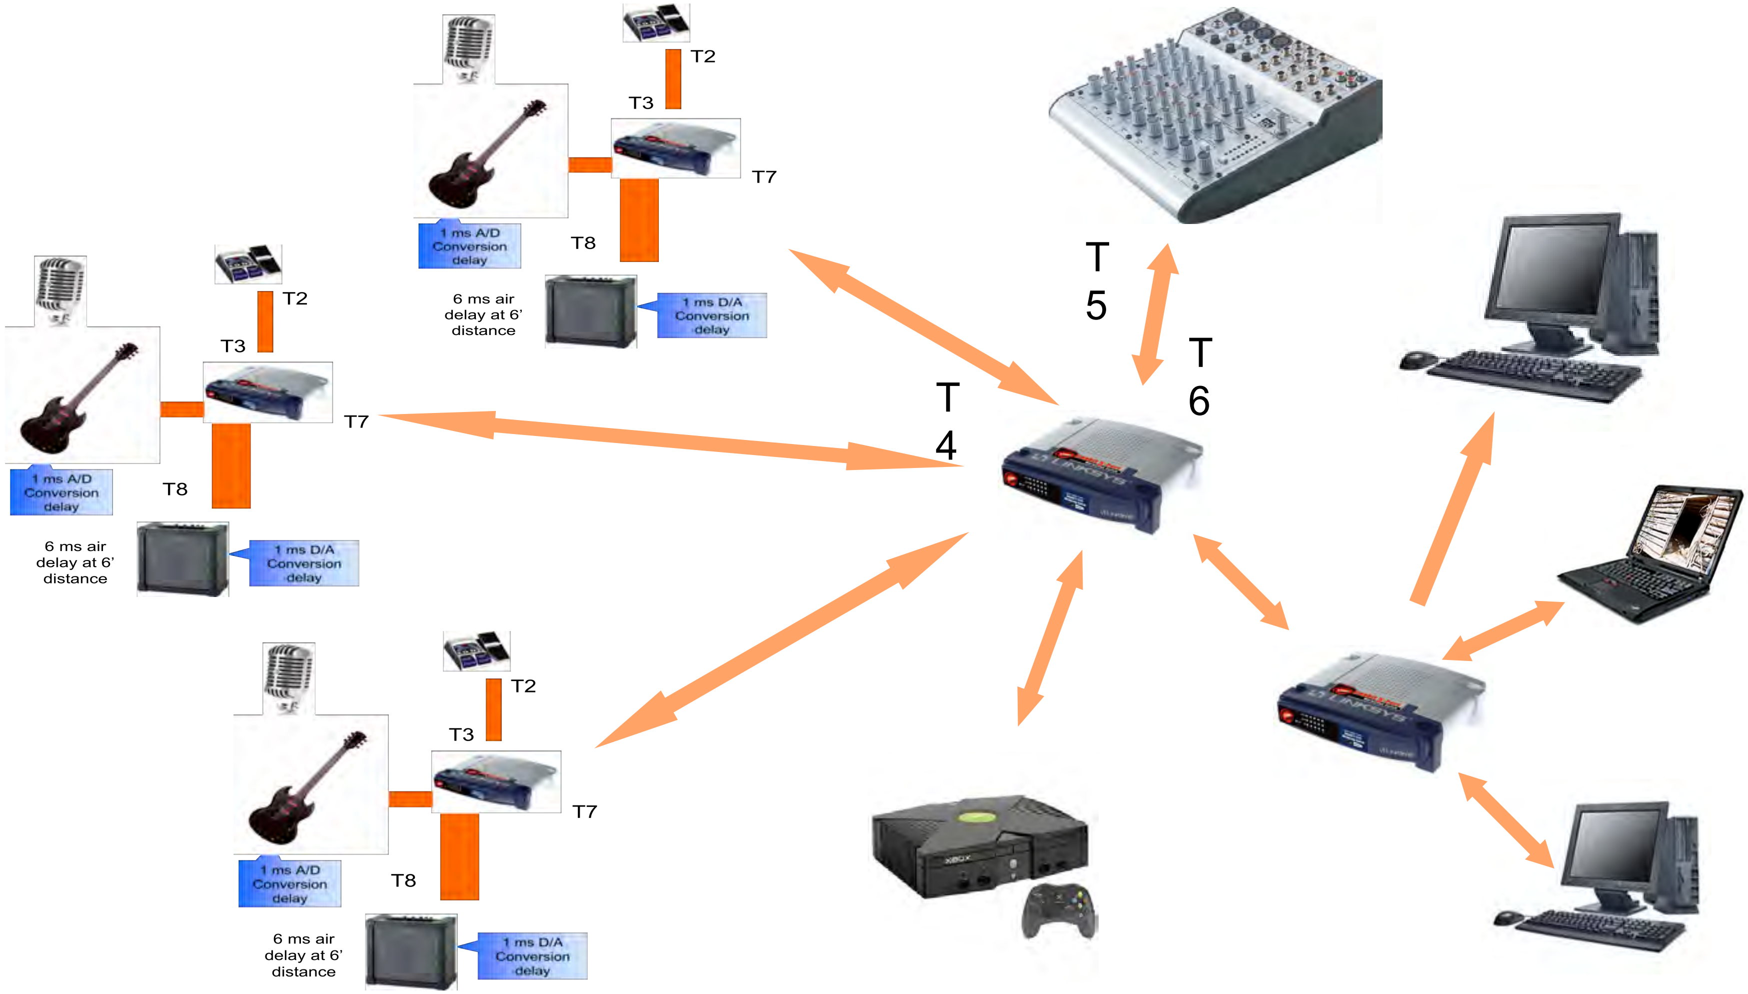

Other important characteristics

- Low jitter, simple deterministic timing model
- Low cost, reliable and simple in use
- Operates over 802.3 standard existing PHY's and the associated media
- Support for multiple channels of advanced fidelity audio and high-quality video simultaneously
- Power over Ethernet

Conclusion

Existing and future applications require Residential Ethernet to provide non-interfering data transmissions with guaranteed bandwidth and predictable low latency and jitter.

Thank you!

Residential Ethernet: a possible architecture

Michael Johas Teener
mike@plumblinks.com

Agenda

- Background
- Architecture overview
- Possible implementations
- Compatibility

Background

- There are many ways to meet the objectives for ResE
- This presentation outlines one approach that
  - meets all ResE objectives
  - reuses existing technology
  - scales well
  - can support integration of wireless services
- There are other proposals!

Proposed architecture

- Propose changes to both IEEE 802.3 (Ethernet) and IEEE 802.1D (bridges/ switches)
- Three basic additions to 802.3/802.1
  - Precise synchronization,
  - Scheduled priorities, and
  - Admission controls
- Some of this has been published previously, and is available at
  - http://www.teener.com/ResidentialEthernet/ Residential%20Ethernet.htm

Topology & connectivity

- ResE services only available in ResE cloud
- Normal best-effort services available everywhere

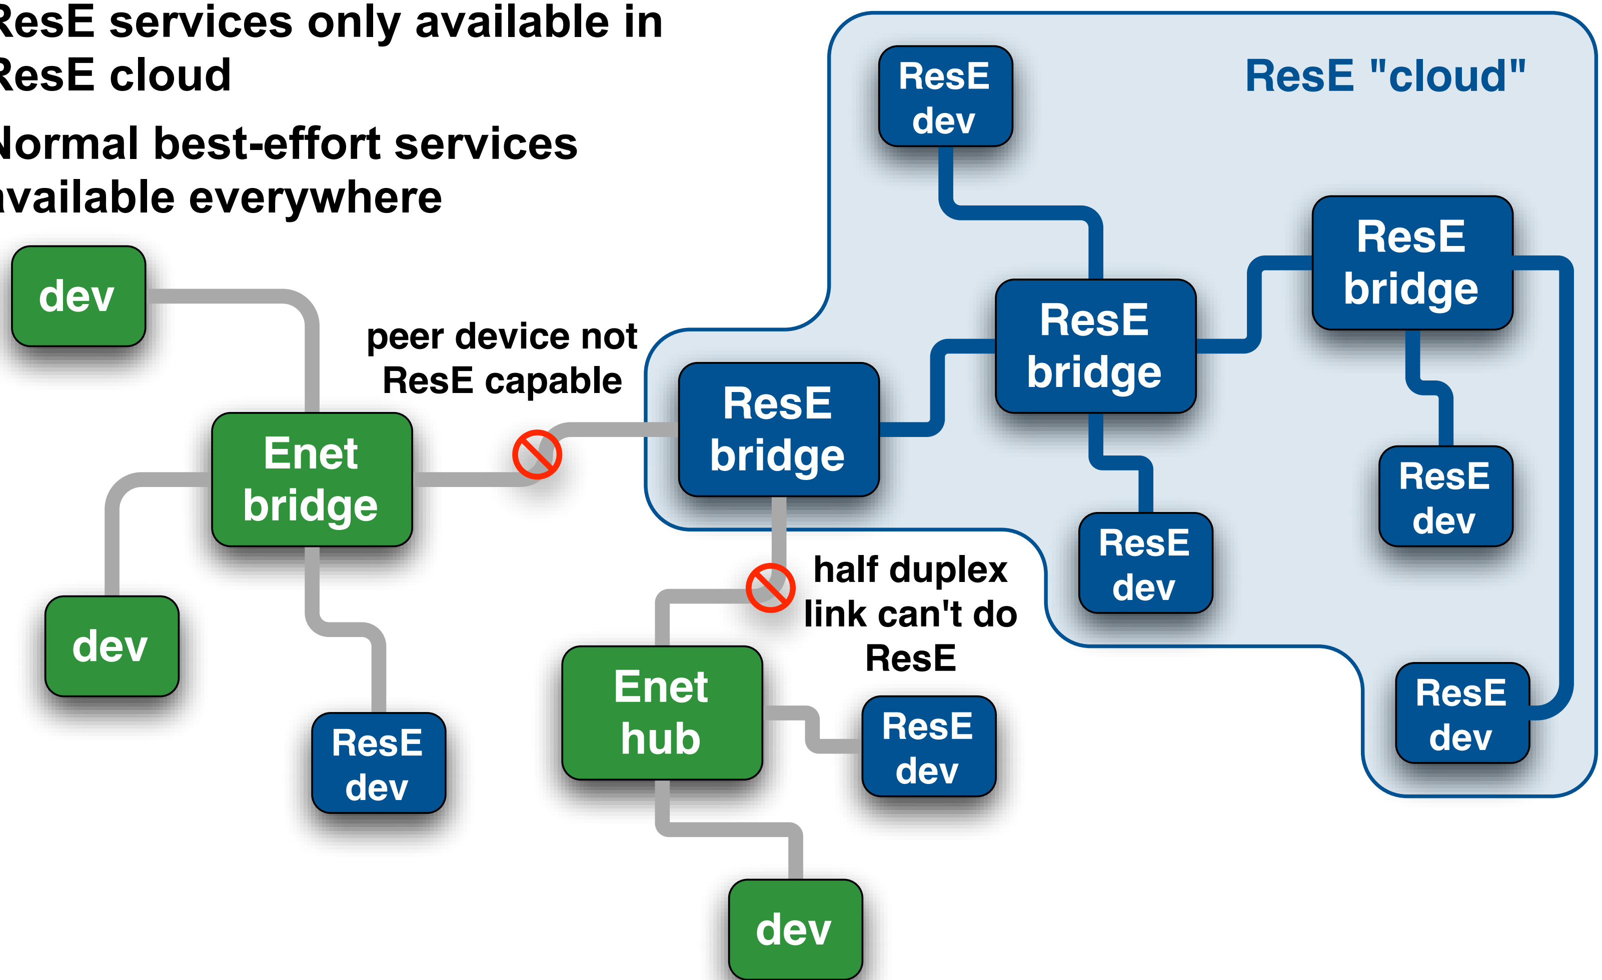

Precise synchronization

- ResE devices will periodically exchange timing information
  - both devices synchronize their time-of-day clock very precisely
  - algorithms and messages derived from IEEE 1588
- This precise synchronization has two purposes:
  - to allow isochronous packet scheduling (see below) 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.

ResE clock

- Provides both cycle timing and common time base
- Isochronous packet scheduling will use 8 kHz "cycles"
  - commonly used in most current isochronous transports
    - the wide-area network, IEEE 1394, or even USB
  - guarantees that all stations agree on when cycles start and how they are numbered
- The common time base similar to the 24.576 MHz version provided by IEEE 1394
  - perhaps more precise

“Grand Master” clock

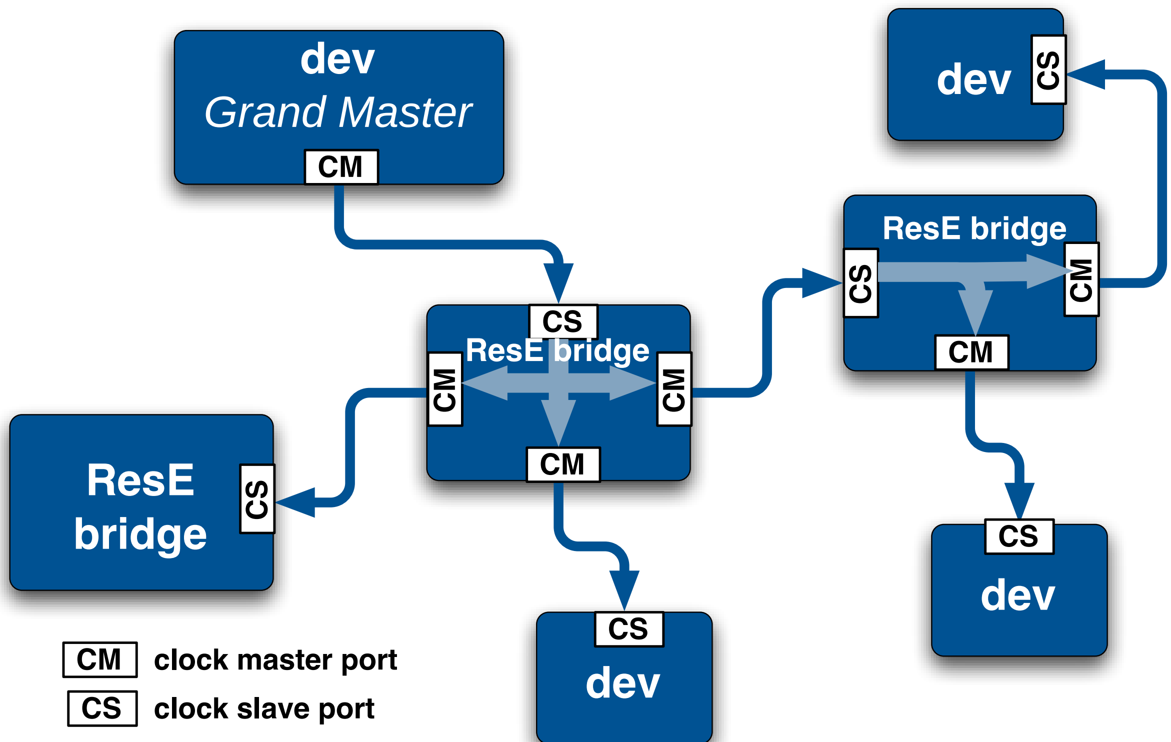

- There is a single device within a ResE “cloud” that provides a master timing signal
  - All other devices synchronize their clocks with this master ... cascaded though bridges

Grand master clock selection

- **Selection of the grand master is largely arbitrary (all ResE devices will be master-capable)**
  - **can be overridden if the network is used in a professional environment that already has a “house clock”**
- **Methodology based on "rapid spanning tree protocol" used in 802.1D bridges**
  - **selection priority is a combination of a managed priority field and bridge ID/MAC address**

ResE link initialization

- **Devices use autonegotiation to discover that they are both ResE capable**

ResE link initialization

- **Devices use autonegotiation to discover that they are both ResE capable**

ResE link initialization

- Devices use autonegotiation to discover that they are both ResE capable

ResE link initialization

- If both devices are ResE, they select which is clock master
  - closest to grand master using current spanning tree
  - new grand master may be selected if devices are bridges

ResE link initialization

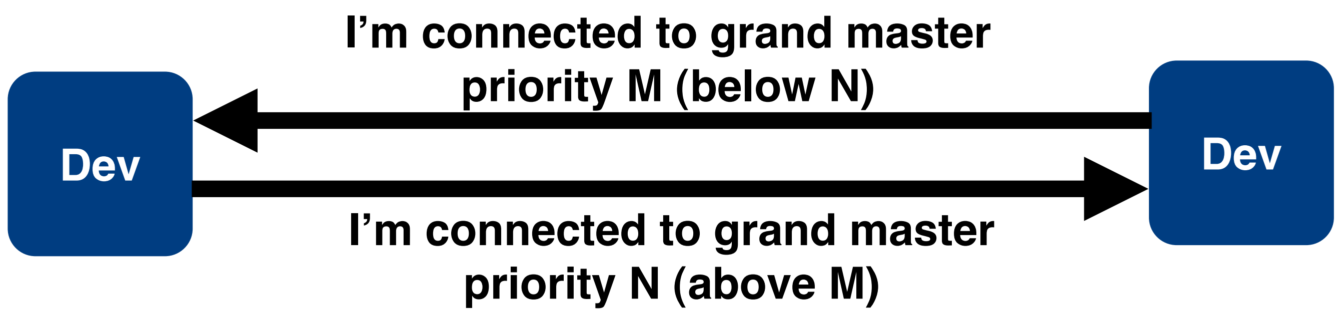

- **If both devices are ResE, they select which is clock master**
  - closest to grand master using current spanning tree
  - new grand master may be selected if devices are bridges

ResE link initialization

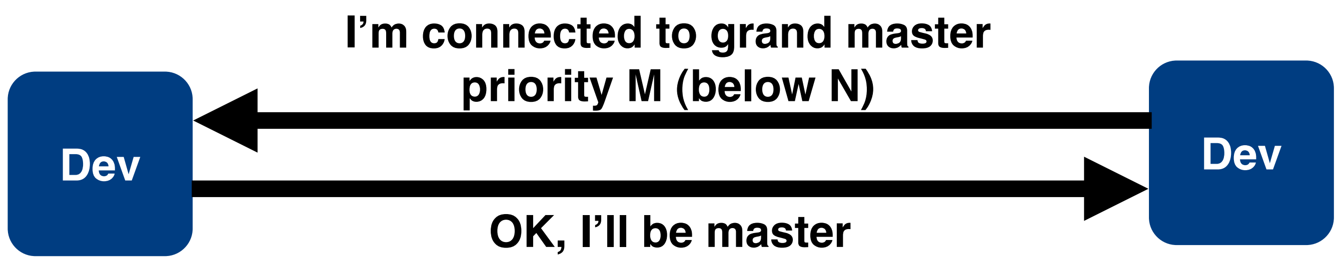

- **If both devices are ResE, they select which is clock master**
  - closest to grand master using current spanning tree
  - new grand master may be selected if devices are bridges

ResE link initialization

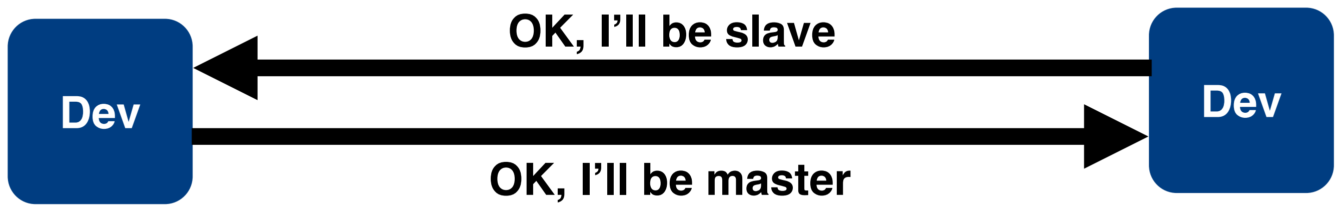

- If both devices are ResE, they select which is clock master
  - closest to grand master using current spanning tree
  - new grand master may be selected if devices are bridges

Synchronization algorithm

- Same as IEEE 1588 methods
  - which run at layer 3 using IP
- ... on layer 2 instead using LLC
  - minimize queuing delays
  - faster stabilization
- Cascaded link synchronization

Isochronous transport (scheduled priorities)

- **Uses the 802.1D routing and priority mechanisms in a modified way to provide isochronous services**
- **Each cycle, all isochronous packets assembled at a *network interface* for transmission are tagged with the next cycle number and placed in a transmit queue.**
- **During cycle "n", packets are transmitted first if they are tagged with cycle "n-1", then "n", then all best effort traffic (normal legacy Ethernet packets). Packets tagged with the current cycle number must wait until the next cycle.**

Example isochronous traffic

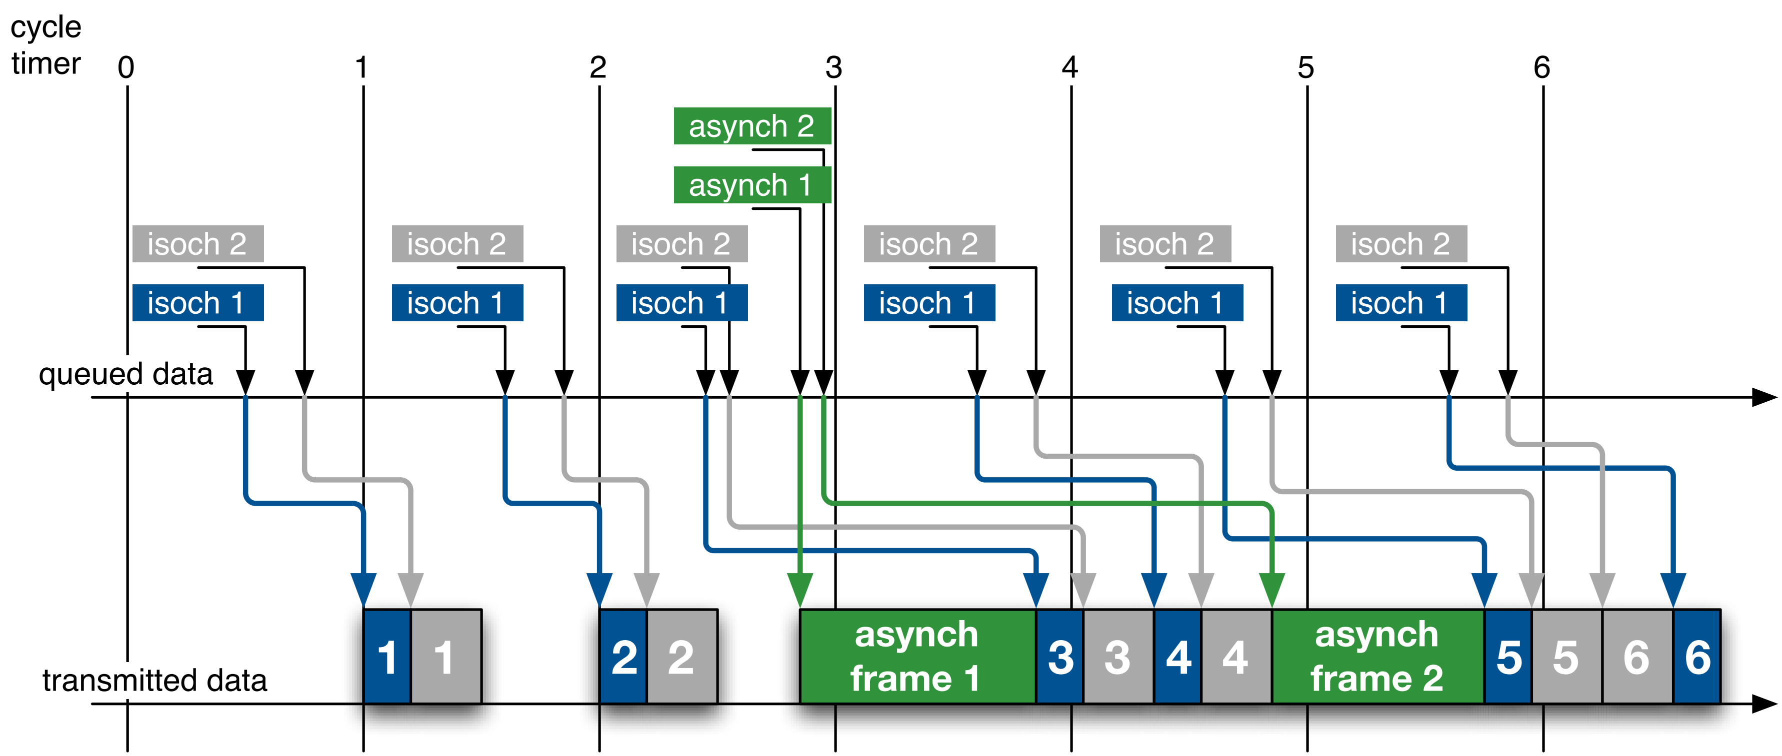

Isochronous packets

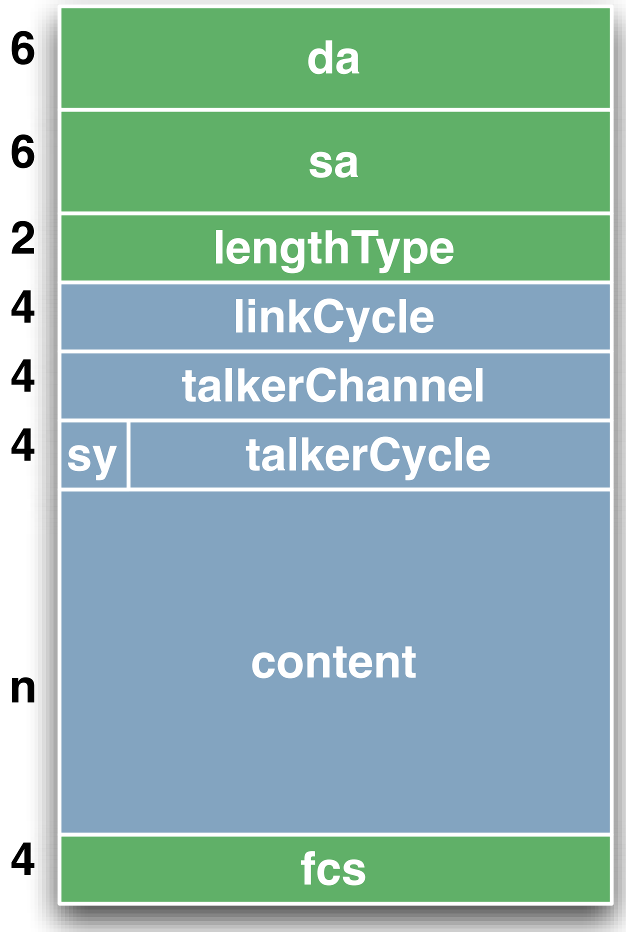

- Isochronous packets are normal 802.3 frames
  - flagged by a special value of the "Length/Type" field

Note: 802.1X/Q/etc effects not shown

Isochronous bridges

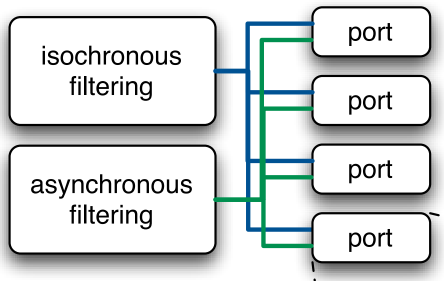

- **Isochronous packets will be routed just like best-effort traffic in existing 802.1D bridges**
  - **linkCycle field will be incremented by two**
  - **placed in the isochronous transmit queue for the appropriate output port**

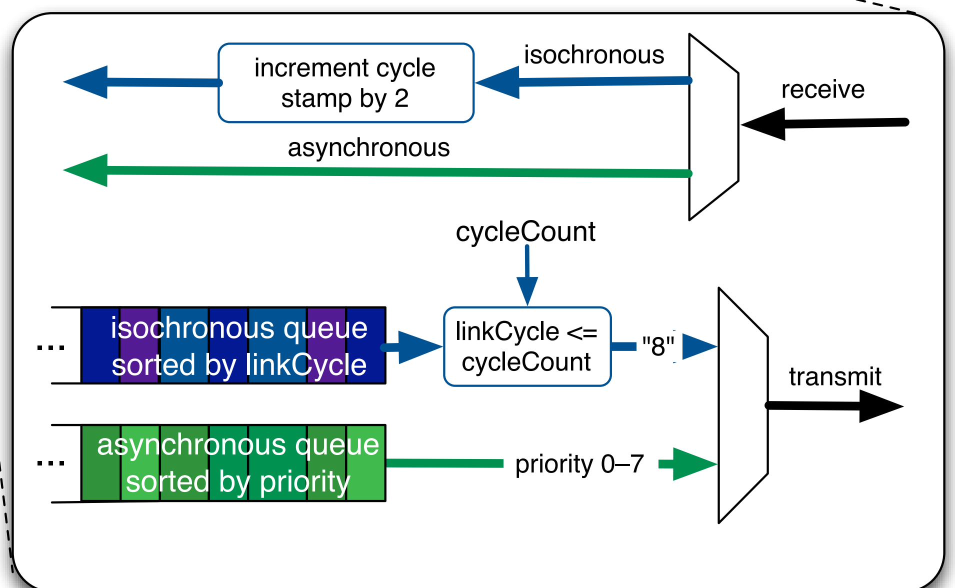

Isochronous bridge example topology

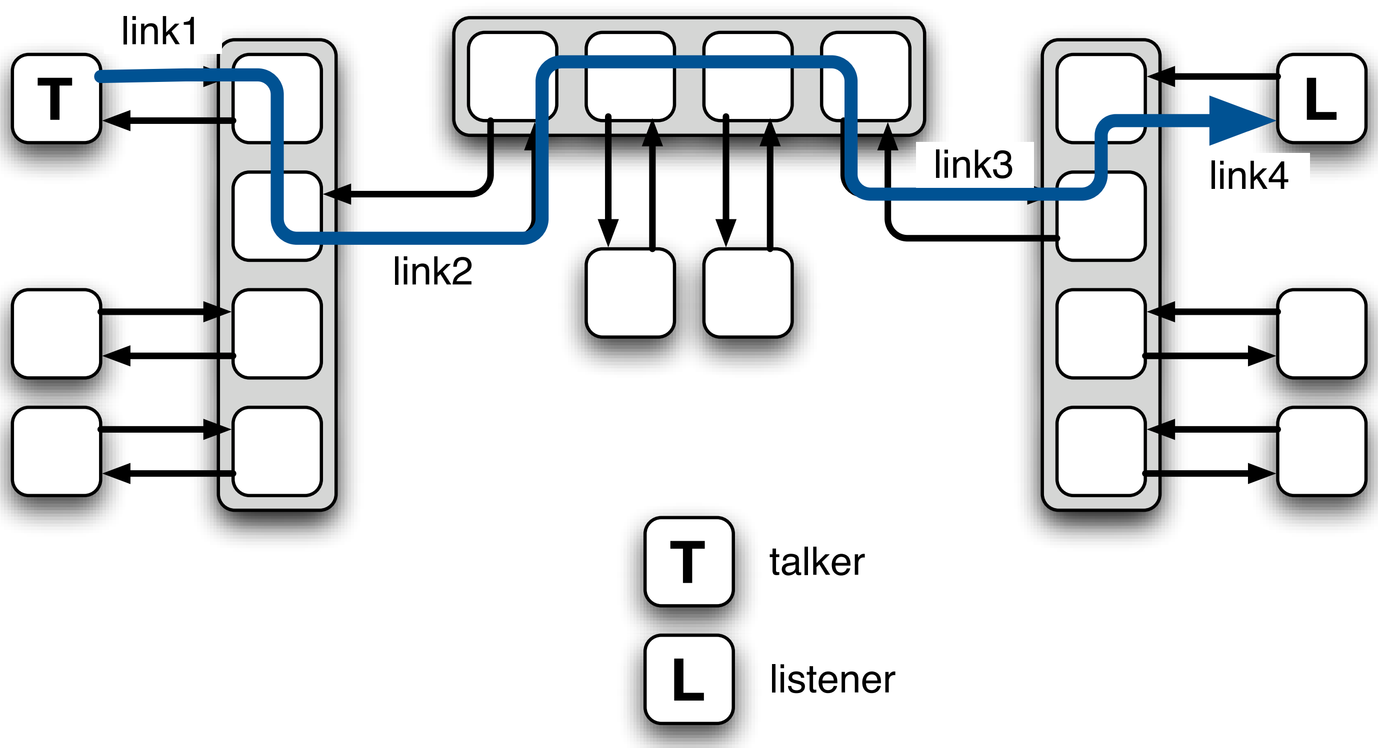

- **In this example, an isochronous stream is being sent from the talker to the listener through three bridges. There are four links in the path as shown above.**

Example traffic through ResE bridges

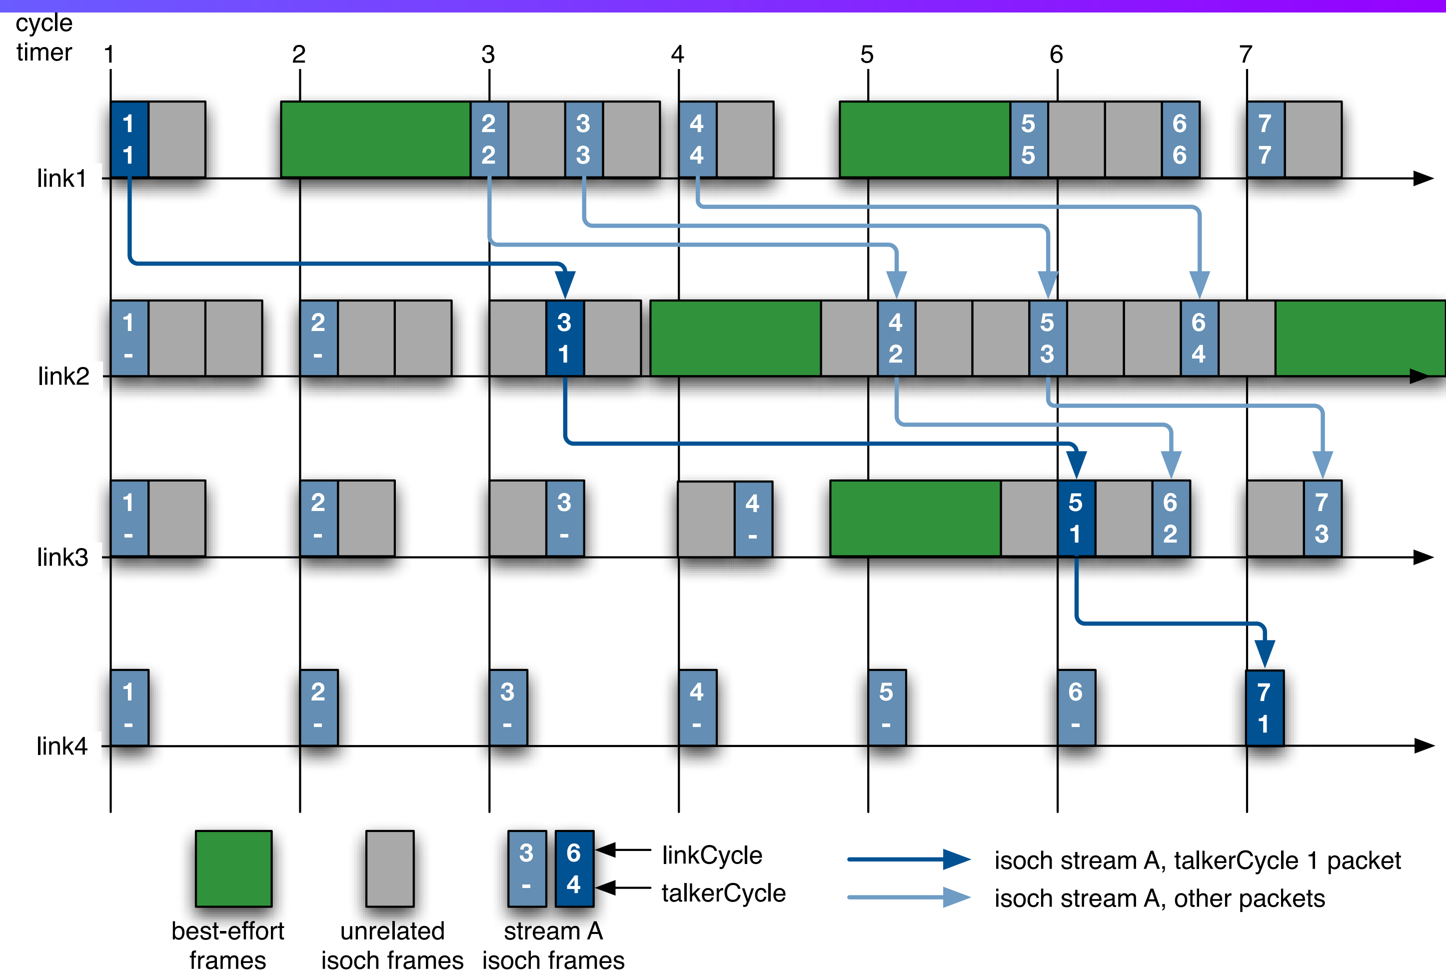

Note how it takes 6 cycles to get through the three bridges, and that there is very little delivery jitter on the last link since there is no interfering traffic.

Admission controls

- Scheduled 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 ResE it is the listener's responsibility to guarantee the path is available and to reserve the resources.
- Done via "join request"

"Join request" procedure (1)

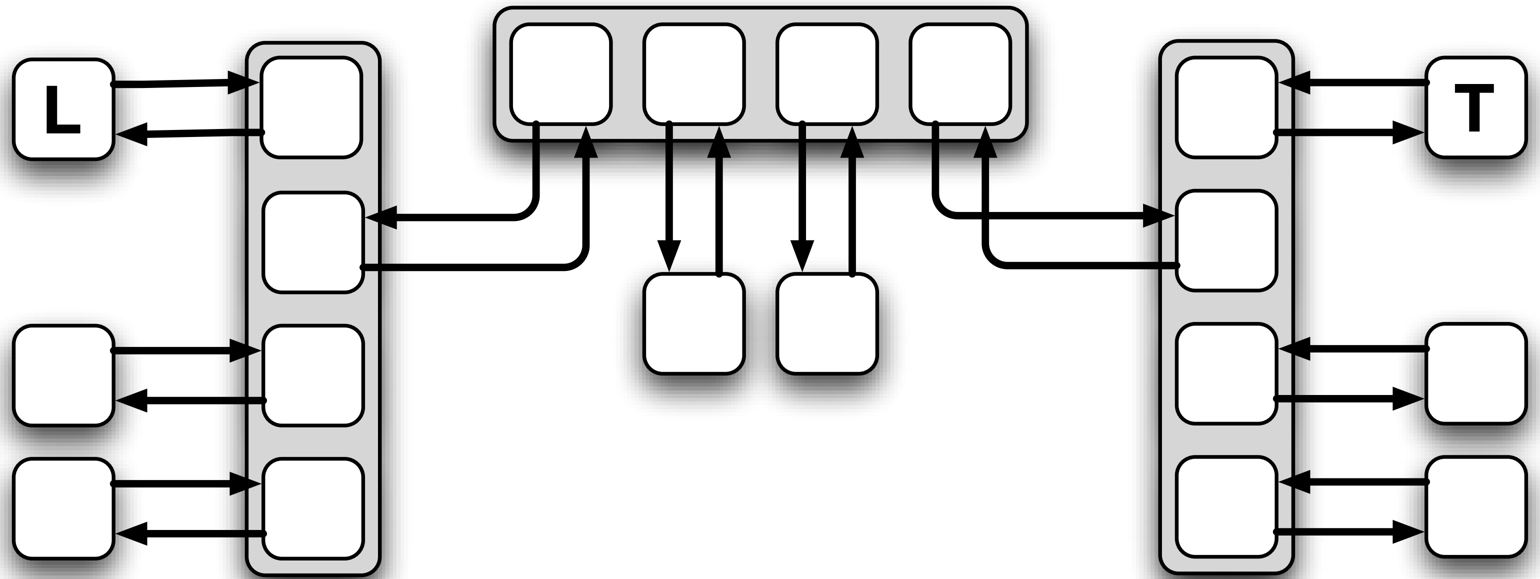

"Join request" procedure (1)

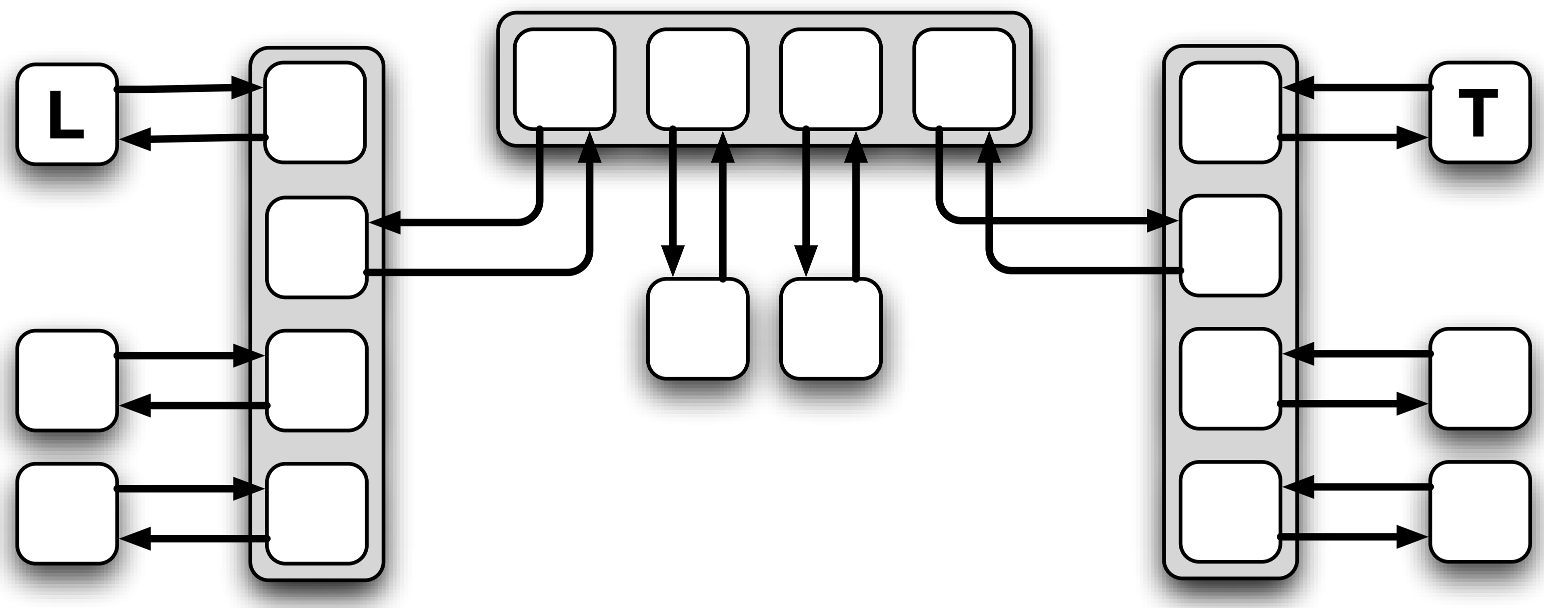

- The listener sends a “join request” control packet to talker with amount of bandwidth needed (in bytes/cycle) as well as a channel number to use.

"Join request" procedure (1)

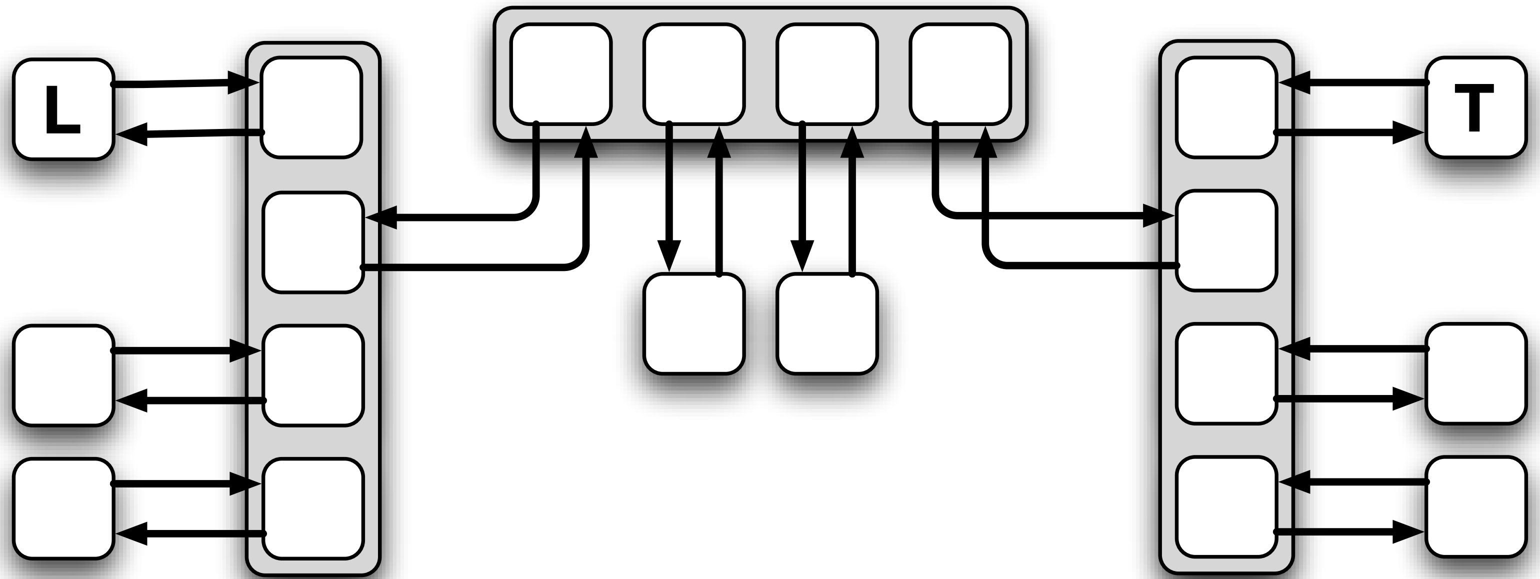

- The listener sends a “join request” control packet to talker with amount of bandwidth needed (in bytes/cycle) as well as a channel number to use.
- The next device in the direction of the talker makes a tentative bandwidth reservation on the port going back to the listener.

"Join request" procedure (1)

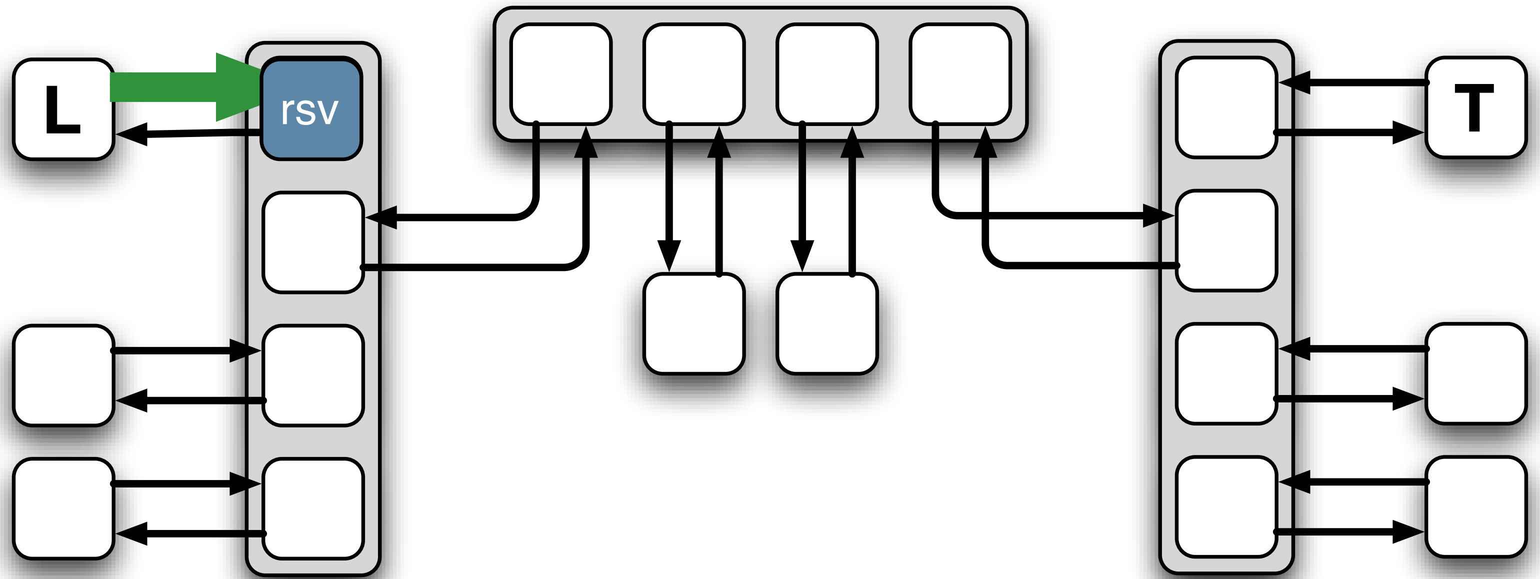

- The listener sends a “join request” control packet to talker with amount of bandwidth needed (in bytes/cycle) as well as a channel number to use.
- The next device in the direction of the talker makes a tentative bandwidth reservation on the port going back to the listener.

"Join request" procedure (2)

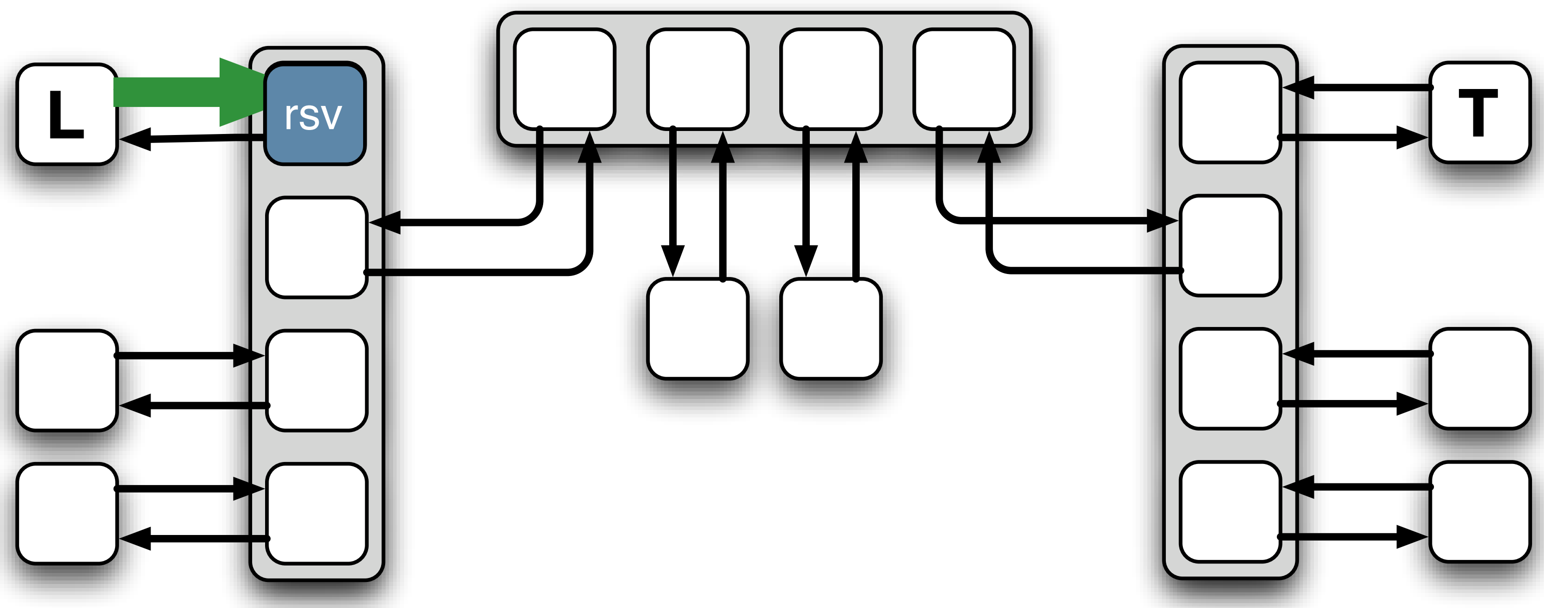

- The bridge then passes the request onto the next device towards the talker. Other intermediate bridges make a tentative bandwidth reservation on the port going back to the listener, update a delay count and pass on control packet toward talker. Note that if a bridge is already routing the stream, it can respond on it's own, acting as a proxy for the talker.

"Join request" procedure (2)

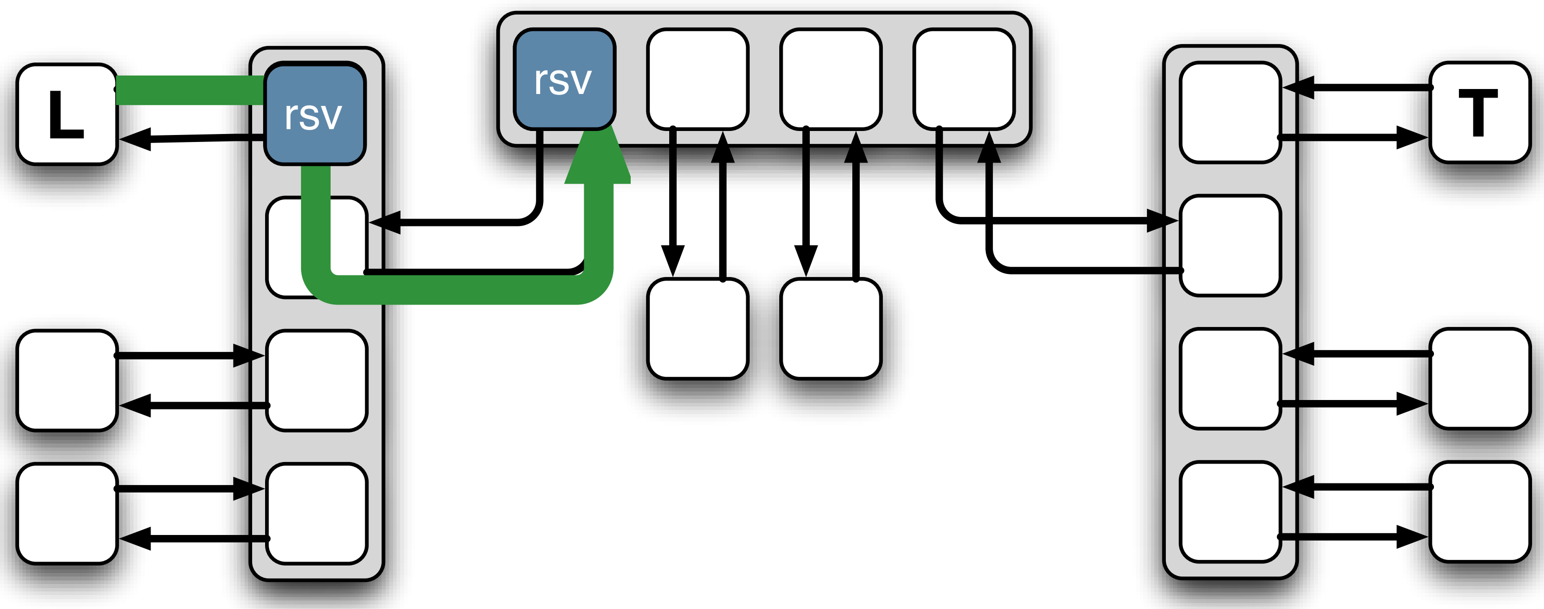

- The bridge then passes the request onto the next device towards the talker. Other intermediate bridges make a tentative bandwidth reservation on the port going back to the listener, update a delay count and pass on control packet toward talker. Note that if a bridge is already routing the stream, it can respond on it's own, acting as a proxy for the talker.

"Join request" procedure (2)

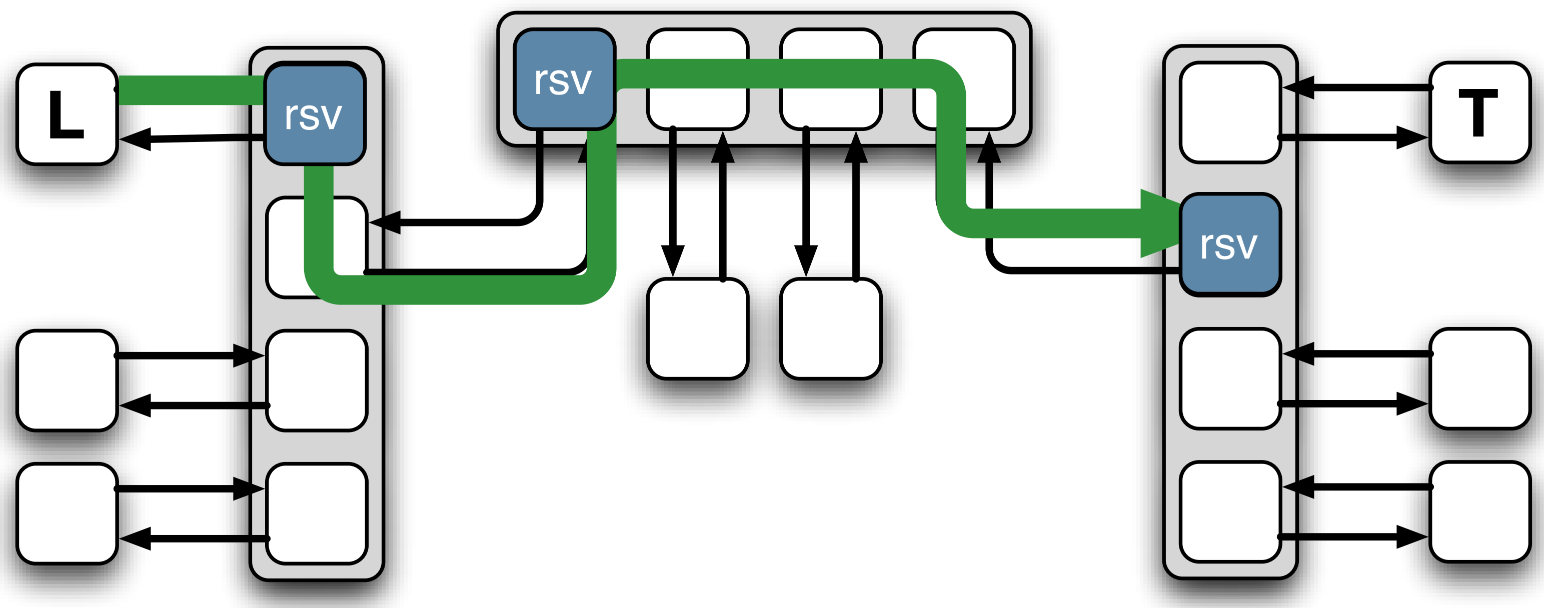

- The bridge then passes the request onto the next device towards the talker. Other intermediate bridges make a tentative bandwidth reservation on the port going back to the listener, update a delay count and pass on control packet toward talker. Note that if a bridge is already routing the stream, it can respond on it's own, acting as a proxy for the talker.

"Join request" procedure (2)

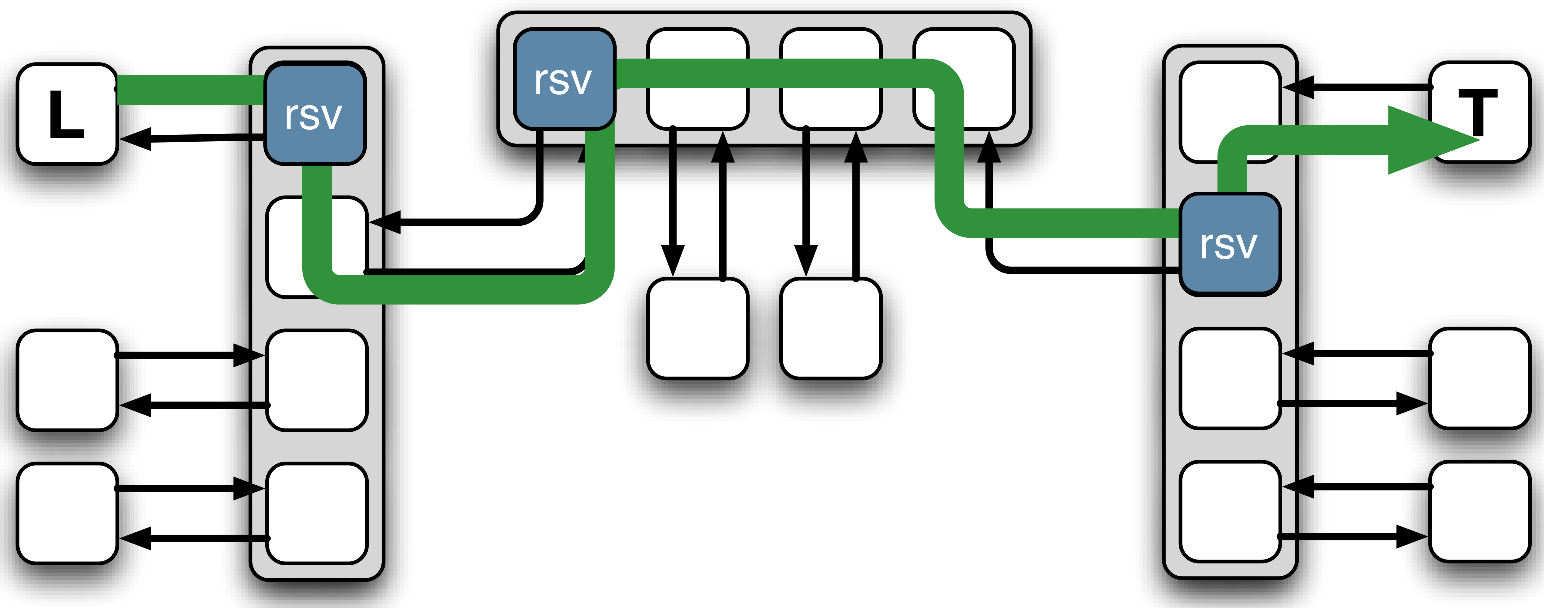

- The bridge then passes the request onto the next device towards the talker. Other intermediate bridges make a tentative bandwidth reservation on the port going back to the listener, update a delay count and pass on control packet toward talker. Note that if a bridge is already routing the stream, it can respond on it's own, acting as a proxy for the talker.

"Join request" procedure (3)

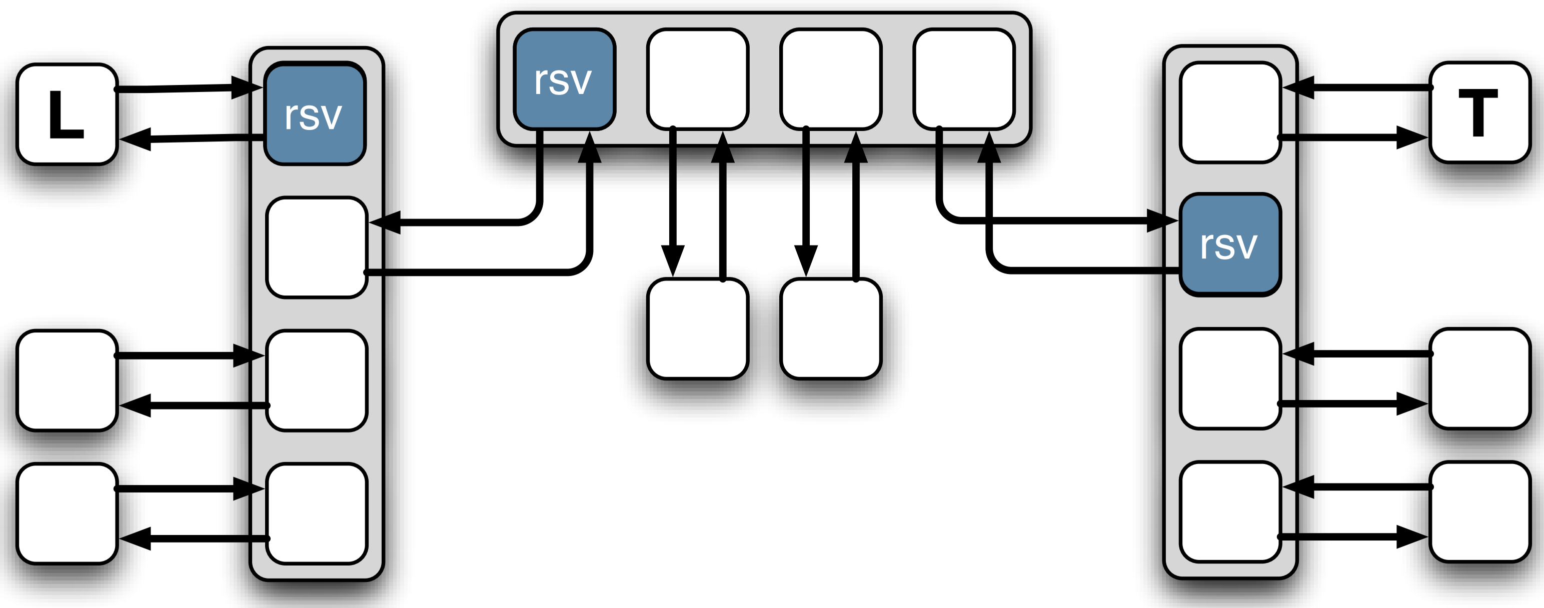

- When the packet reaches the talker, the talker allocates resources and returns a “join response” packet towards the listener with status information (which includes resource available (or not), and delay).

"Join request" procedure (3)

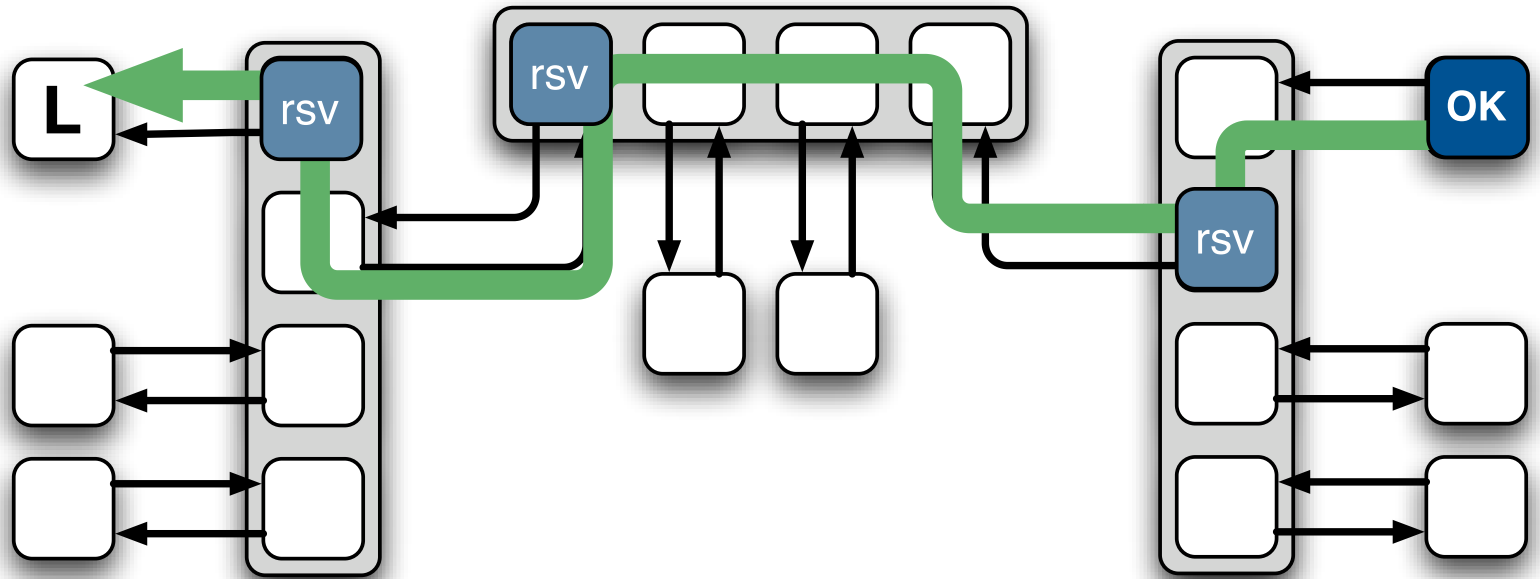

- When the packet reaches the talker, the talker allocates resources and returns a “join response” packet towards the listener with status information (which includes resource available (or not), and delay).

"Join request" procedure (4)

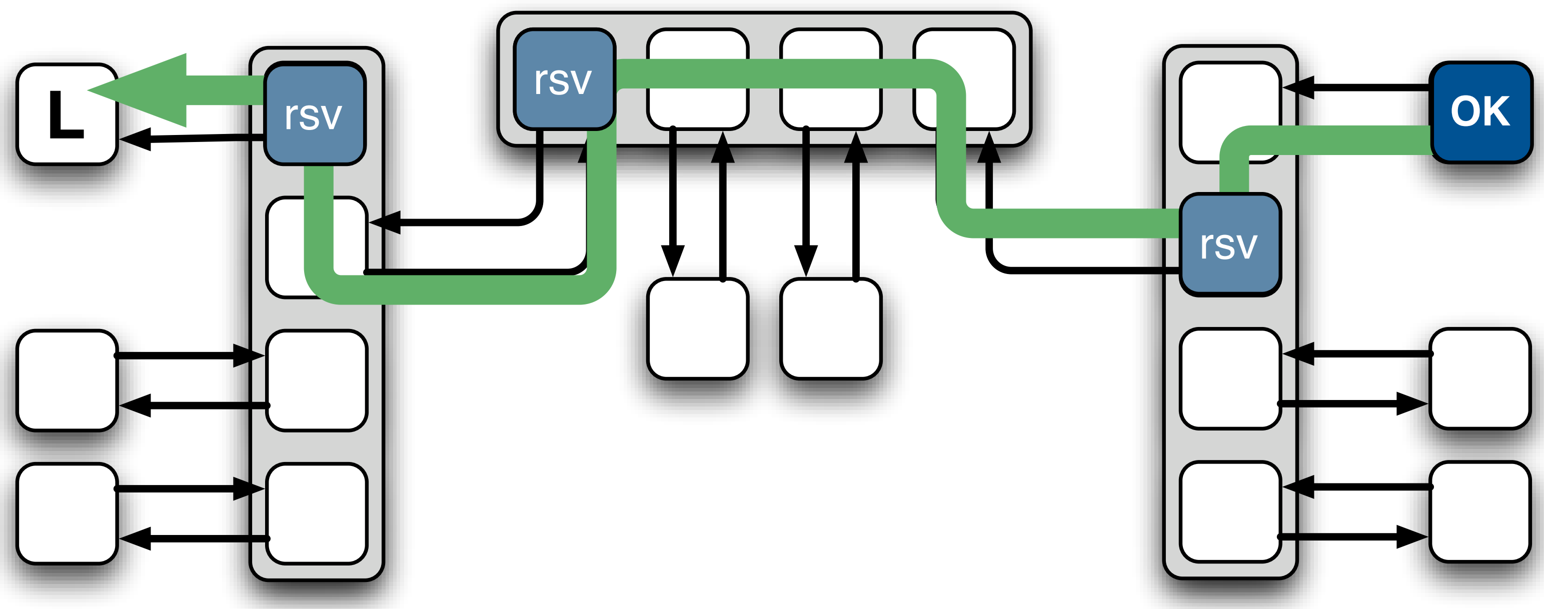

- Intermediate bridges receiving the "join response, resources available" packet confirm the resource allocation. (If the packet is "join response, resources not available", the tentatively assigned resources are released.)
- ... and the talker starts sending its stream imediately

"Join request" procedure (4)

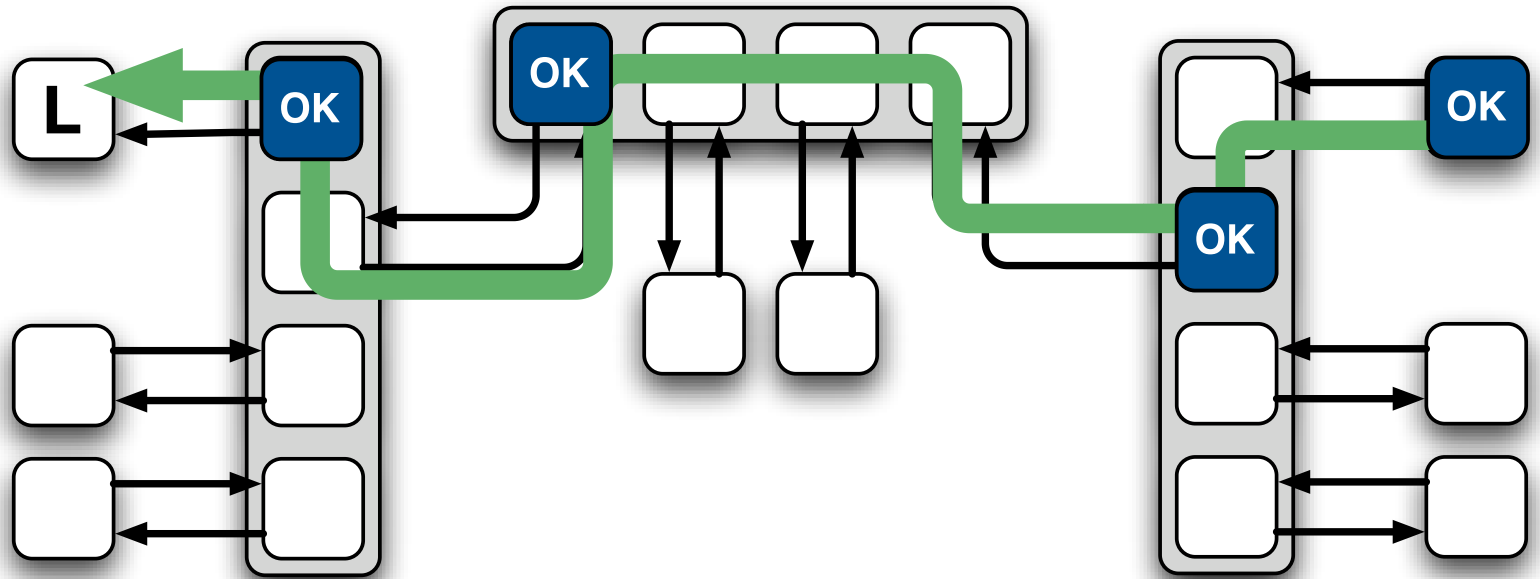

- Intermediate bridges receiving the “join response, resources available” packet confirm the resource allocation. (If the packet is “join response, resources not available”, the tentatively assigned resources are released.)
- ... and the talker starts sending its stream imediately

"Join request" procedure (4)

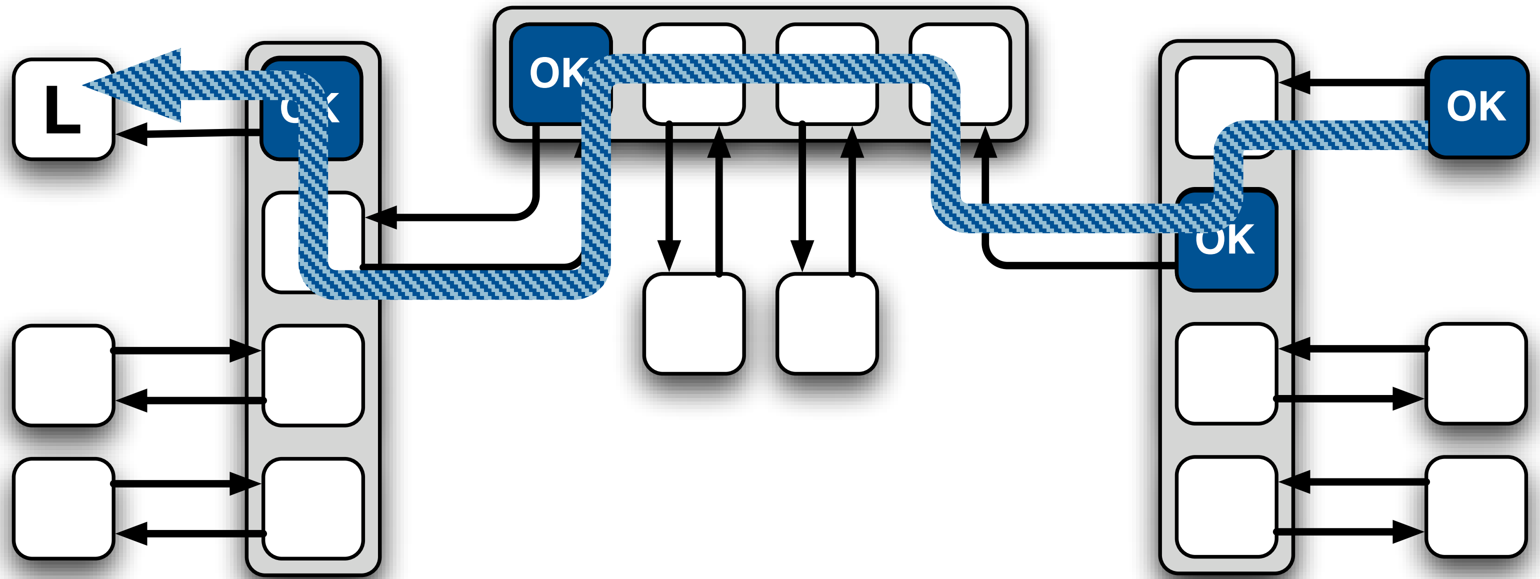

- Intermediate bridges receiving the "join response, resources available" packet confirm the resource allocation. (If the packet is "join response, resources not available", the tentatively assigned resources are released.)
- ... and the talker starts sending its stream imediately

"Teardown request" procedure

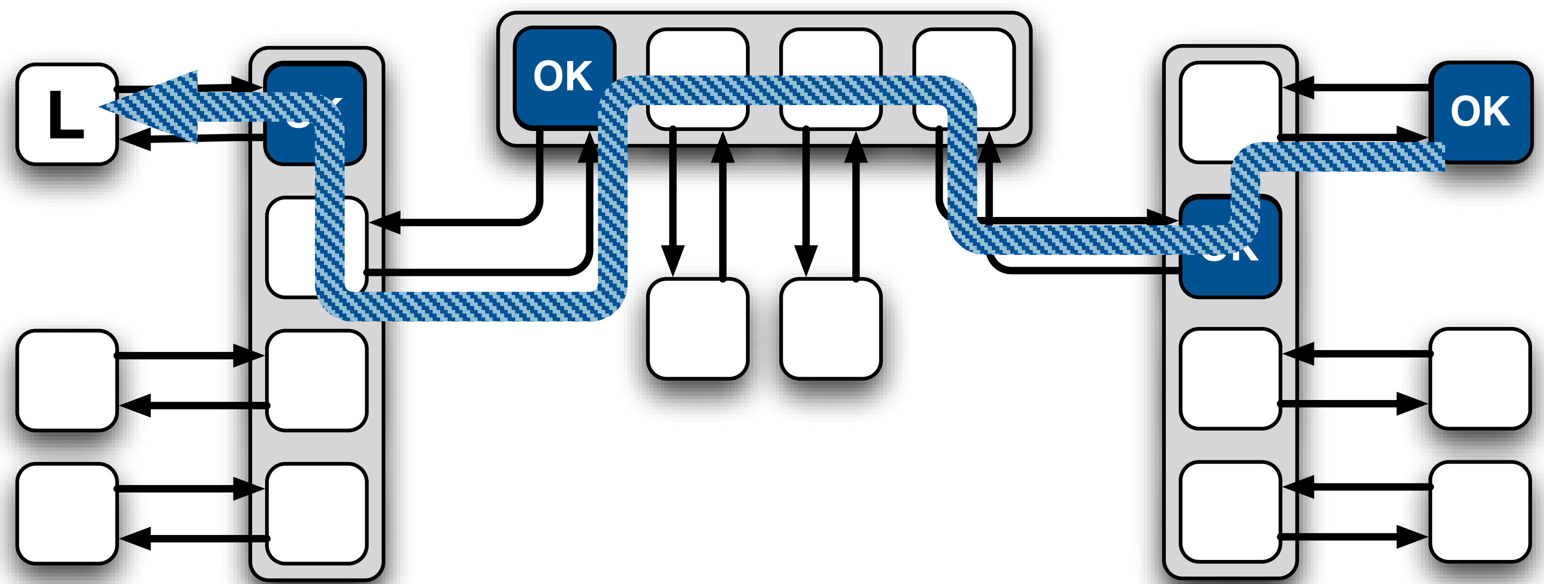

- The disconnecting device sends a “teardown request” control packet towards the talker.
- If an intermediate bridge is not forwarding the stream to another listener, it just forwards the request on towards the talker.
- When a “ teardown request” reaches a talker, it discontinues transmission of the stream (usually notifying a higher layer protocol), deallocates the resources used on its output port, and sends a “ teardown response” control back towards the disconnecting device.
- If an intermediate bridge receives a “ teardown response”, it deallocates the resources assigned to the stream on the port towards the disconnecting device, and forwards the “ teardown response packet”.

"Teardown request" procedure

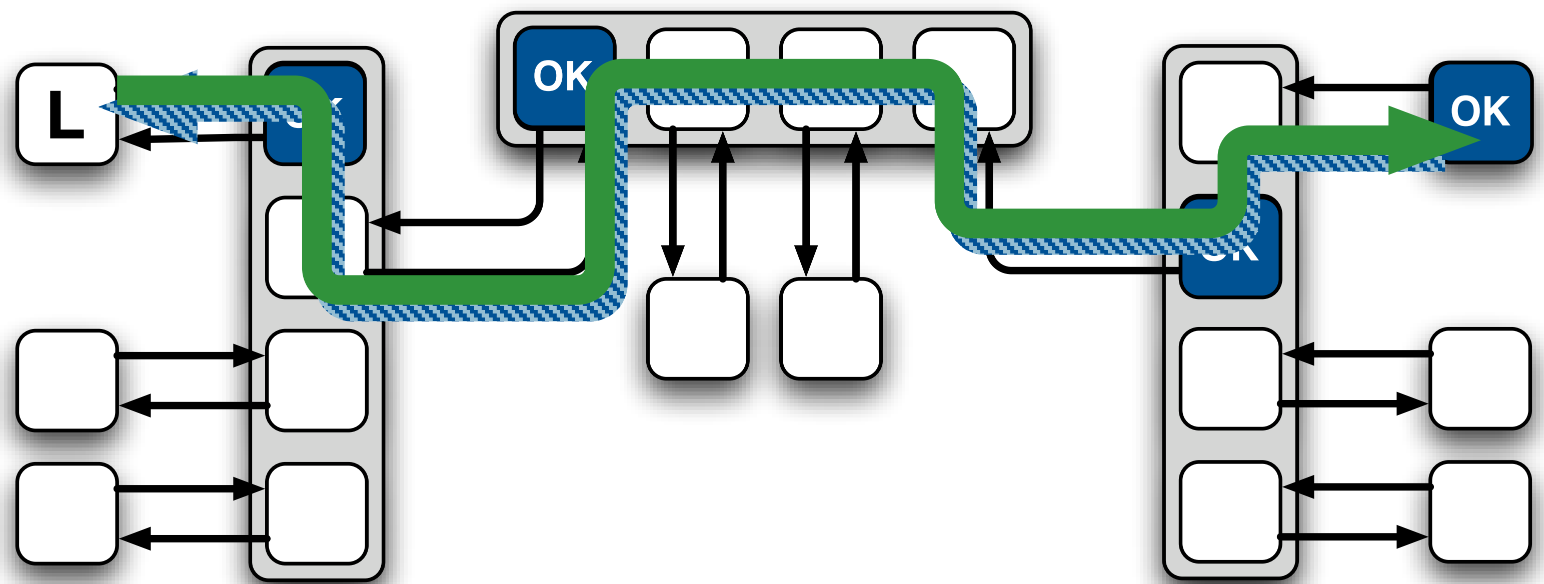

- The disconnecting device sends a “teardown request” control packet towards the talker.
- If an intermediate bridge is not forwarding the stream to another listener, it just forwards the request on towards the talker.
- When a “ teardown request” reaches a talker, it discontinues transmission of the stream (usually notifying a higher layer protocol), deallocates the resources used on its output port, and sends a “ teardown response” control back towards the disconnecting device.
- If an intermediate bridge receives a “ teardown response”, it deallocates the resources assigned to the stream on the port towards the disconnecting device, and forwards the “ teardown response packet”.

"Teardown request" procedure

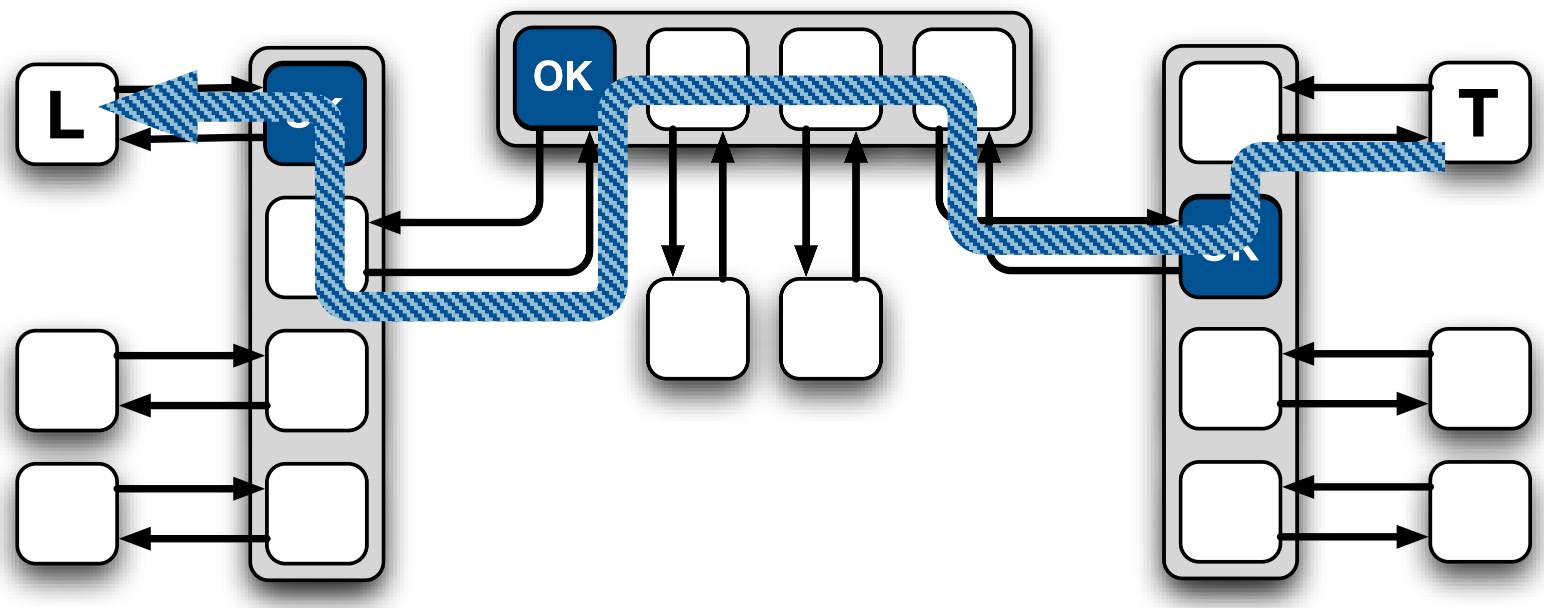

- The disconnecting device sends a “teardown request” control packet towards the talker.
- If an intermediate bridge is not forwarding the stream to another listener, it just forwards the request on towards the talker.
- When a “ teardown request” reaches a talker, it discontinues transmission of the stream (usually notifying a higher layer protocol), deallocates the resources used on its output port, and sends a “ teardown response” control back towards the disconnecting device.
- If an intermediate bridge receives a “ teardown response”, it deallocates the resources assigned to the stream on the port towards the disconnecting device, and forwards the “ teardown response packet”.

"Teardown request" procedure

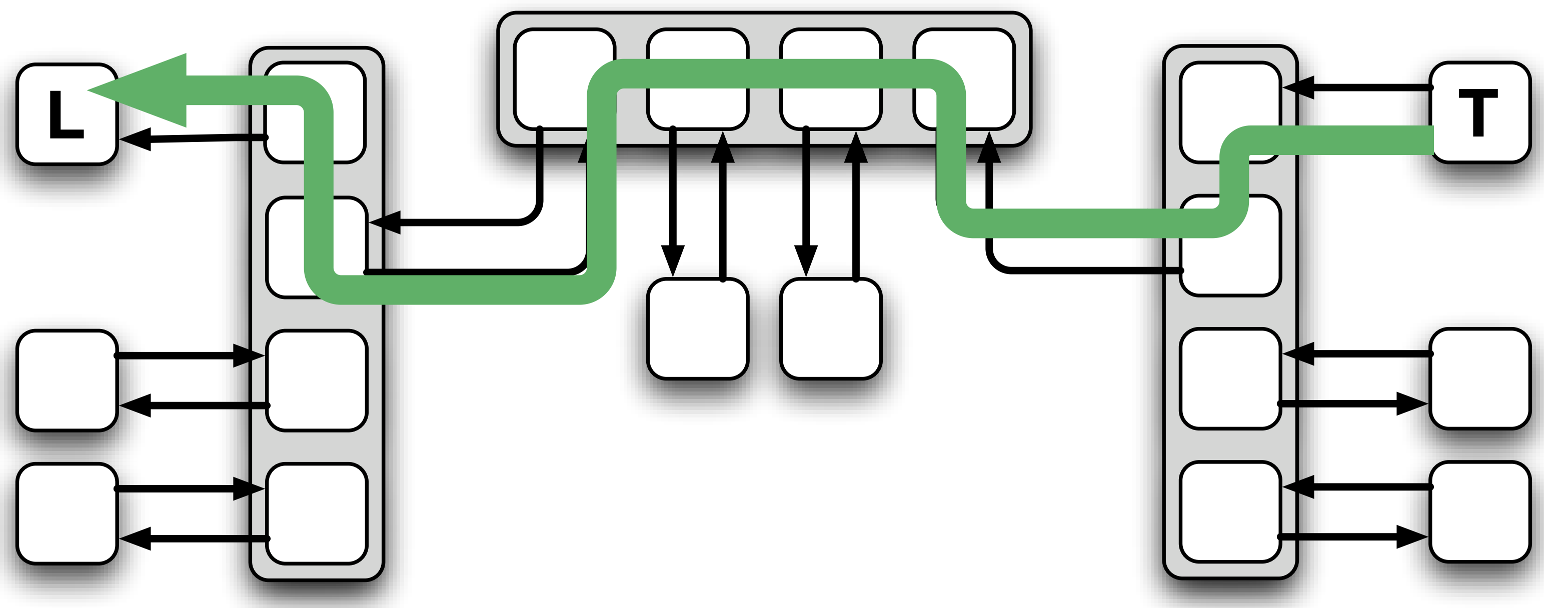

- The disconnecting device sends a “teardown request” control packet towards the talker.
- If an intermediate bridge is not forwarding the stream to another listener, it just forwards the request on towards the talker.
- When a “ teardown request” reaches a talker, it discontinues transmission of the stream (usually notifying a higher layer protocol), deallocates the resources used on its output port, and sends a “ teardown response” control back towards the disconnecting device.
- If an intermediate bridge receives a “ teardown response”, it deallocates the resources assigned to the stream on the port towards the disconnecting device, and forwards the “ teardown response packet”.

Automatic disconnections

- The talker must send one packet every cycle, even if it has no content
  - maybe every "n" cycles, where "n" is small
- Any receiving device (including intermediate bridges) could automatically release assigned resources and notify listener(s) if the appropriate isochronous packet was missing for "m" cycles, where "m" is some small number > "n"

Admission control protocol

- **Use subset of IETF RSVP**
  - **Only “homogeneous” connections**
  - **Only unicast or one-to-many multicast**
  - **Only constant bit rate**
  - **Run at layer 2 instead of layer 3**
- **Obviously, various timeouts and disconnects affect the process**
  - **the basic ideas have been worked out already as part of RSVP and IEEE Std 1394.1-2004, the FireWire bridging protocol**

so there isn’t much invention to do!

Changes needed (1)

- **Endpoint MAC needs**
  - Timer
    - Probably synchronization protocol hardware, but not required
  - Isochronous transmit FIFO
    - (isochronous receive use existing FIFO)
  - Best to have dedicated ports for streaming data
    - MPEG-TS, I2S, etc., like existing 1394 Links

Changes needed (2)

- Bridges
  - ResE MACs
  - Isochronous routing/filtering
    - similar to asynch logic
  - Admission control firmware
    - similar to 802.1 spanning tree and multicast management
    - subset of RSVP
  - Clock mastership
    - subset of IEEE 1588

Complexity comparison

- Industry experts say that ResE bridges will not be significantly more complex than 802.1D learning bridges
  - zero long term cost
- ResE endpoint devices will be simpler than 1394 endpoint devices
  - No distributed arbitration or bandwidth management
- ResE endpoint devices will not be significantly more complex than legacy Ethernet
  - Ethernet costs driven by complexity of PHY

Possible implementations: NICs (network interface controller)

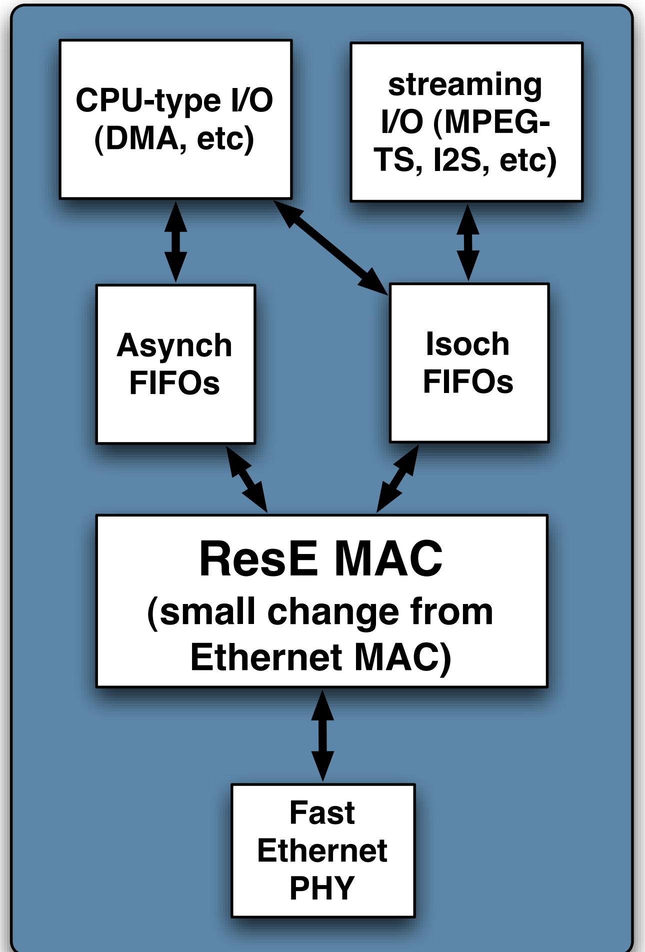

- Enhanced Ethernet NIC
  - could be replacement for 1394 Link/PHY
  - Fast Ethernet is “fast enough” for end points
  - PHY could be Gbit Ethernet as well

Possible implementations: Bridges

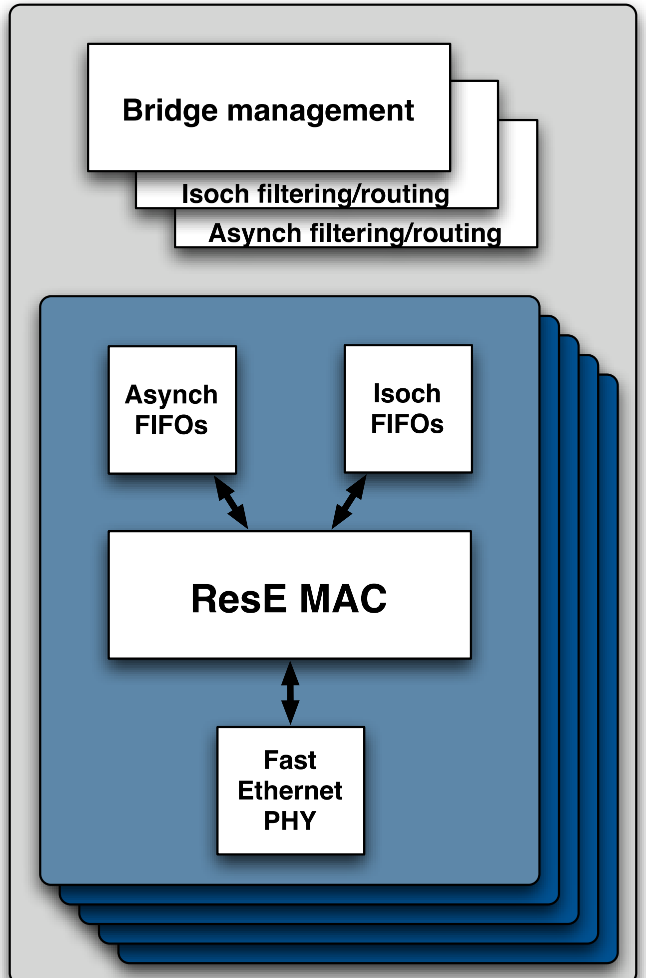

- Enhanced 802.1D bridge
  - basic infrastructure building block

Possible implementations: Combo ResE/1394 bridge

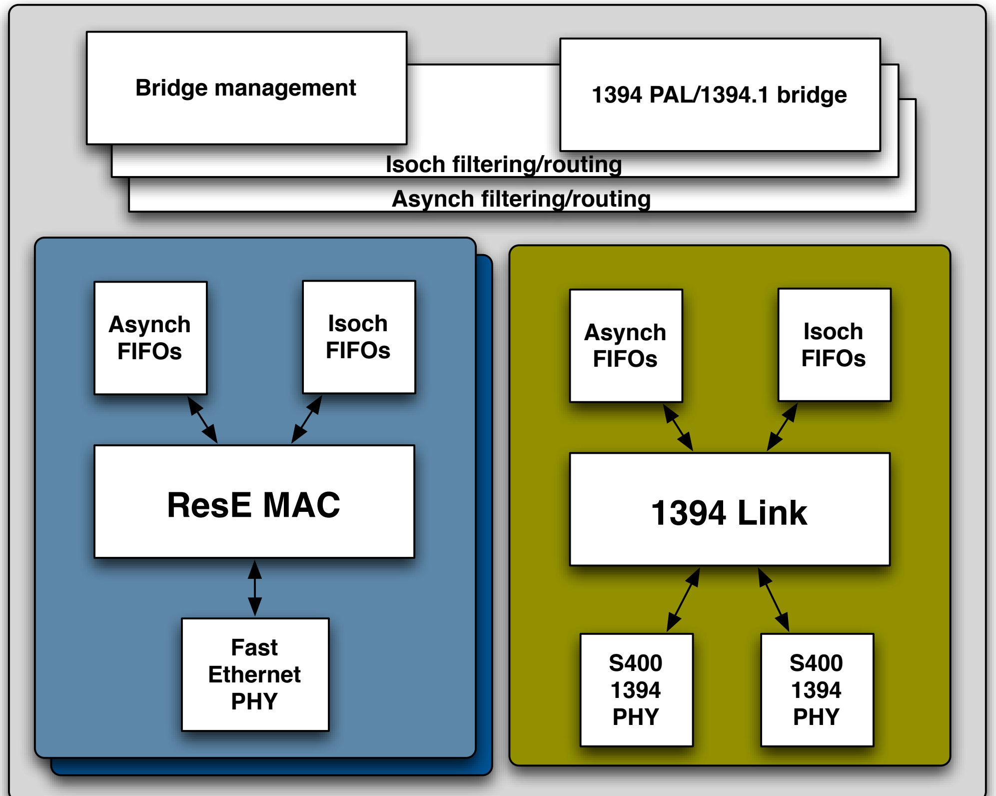

- **Infrastructure for combination 1394/ResE/legacy Ethernet networks**
  - **Transparent bridging of 1394 and ResE**

Possible implementations: Combination ResE NIC/1394.1 bridge

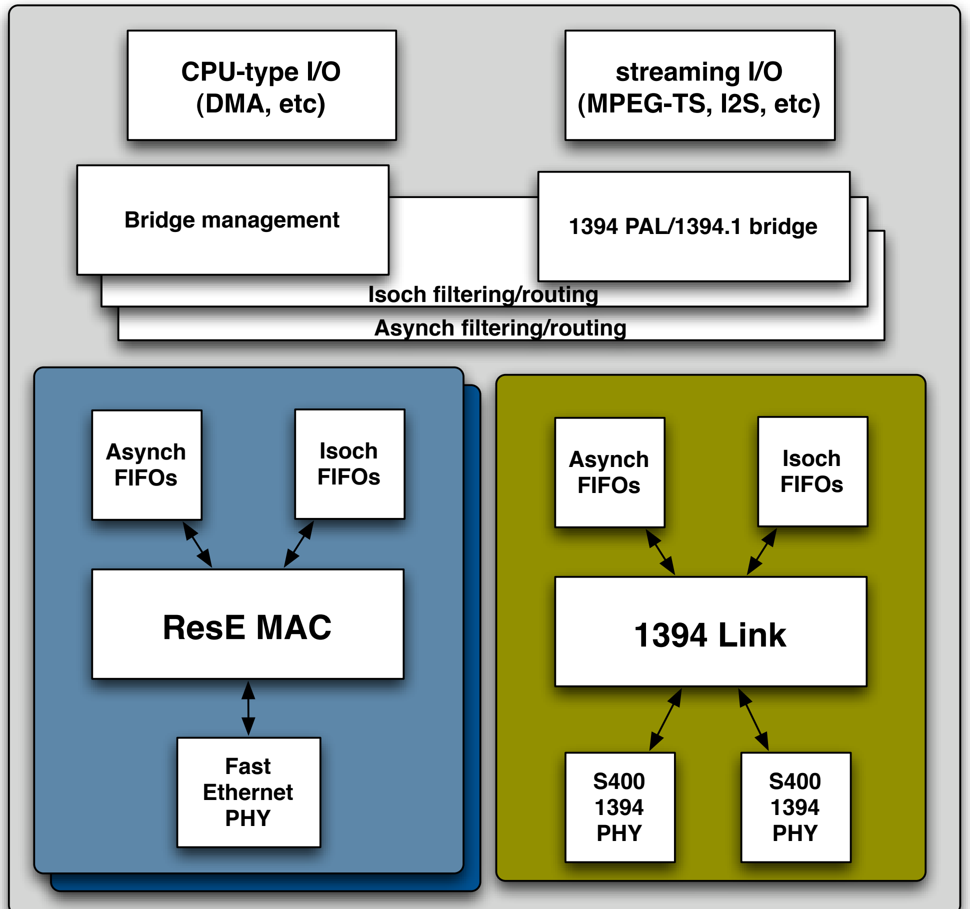

- For high-value products at the center of the A/V cluster
  - DTV
  - Set-top-box
  - A/V receiver

Compatibility

- with existing Ethernet
- with 802.11 and 802.15.3 wireless
- with 1394
- with DLNA
- with UPnP

With existing Ethernet

- All existing Ethernet devices will communicate with ResE devices using "best effort" QoS
- All existing Ethernet devices will communicate exactly as they do now if ResE bridges are in the path between them
- Any existing Ethernet device in the path between ResE devices will act as a block to the communication of isochronous streams
- ResE will be a strict superset of legacy Ethernet

With 802.11 and 802.15.3 wireless

- **All existing 802.11 and 802.15.3 devices using "best effort" QoS (including priorities) will communicate with ResE devices at least as well as current Ethernet devices**
- **There will be a mapping of 802.15.3 "scheduled access" (and the 802.11e equivalent) onto the new ResE bridges.**
- **A ResE backbone will allow all QoS attributes of 802.11e and 802.15.3 to be respected.**

With 1394

- Isochronous 1394 streams will be carried by ResE with no changes
  - It will be possible to build a low-cost hardware-based bridge for isochronous streams between 1394 and ResE
- Asynchronous 1394 transactions will be mapped onto ResE frames using a ResE PAL
  - Protocol Adaptation Layers for 802.15.3 already exists, and one for 802.11e is almost complete.
- Higher layer protocols used by 1394 will all work unchanged if they are "bridge-ready"
  - ... compliant with the requirements of 1394.1

With DLNA

- Existing (and planned) DLNA data/ control works without change
- If DLNA adopts 1394-type streaming, those would work without change
- To take advantage of ResE, DLNA must allow streaming to be done using isochronous capabilities of Layer 2 (ResE/1394/USB/WAN/ 802.15.3/etc)

With UPnP

- Existing UPnP services work without change.
- Planned UPnP QoS 3 "Parameterized QoS" will work well with ResE
  - Please consider contributing

Thank you!

Possible structuring of ResE services

Michael Johas Teener
Plumblinks
mike@plumblinks.com

Agenda

- What services are needed
- All services in MAC?
- All services above MAC?
- Shared?
- Conclusions

Assumptions

- Objectives list from September 2004 interim ResE SG and subsequent informal meetings in San Jose
  - fully backwards compatible with 802.3 and higher layers
  - all existing PHYs supported that are at least 100Mb/s and full duplex
  - add precise synchronization, admission controls, and low latency isochronous services based on 8kHz cycles
  - no topology restrictions beyond what is required for 802.1D spanning tree bridges
- General approach as outlined in "Possible Architecture" part of this tutorial

What services are needed

- Global precise synchronization
  - "house clock"
- Admission controls
  - management of resources
- Low latency isochronous transport
  - schedule packet for transmit during particular isoch period ("cycle")

Synchronization services for client

- Clock synchronization direction control
  - from/to network (clock master or not)
- Clock to network
  - if clock master
- Clock from network
  - higher level scheduling of services
    - need to know current time to know when in the future an event can be scheduled
  - time stamping of streaming data

Synchronization in bridge

- protocol to select master clock in network
  - if no bridge, just uses “best” MAC address
- accept clock from port connected to grand master
- forward clock to other ports with appropriate filtering
  - time-of-day update quick
  - phasing of clock changes slowly

Admission controls for client

- Request bandwidth from link neighbor output
  - bytes/cycle … makes reservation in output queue of link neighbor
- Release bandwidth from link neighbor output
- Accept bandwidth request from link neighbor
  - bytes/cycle … makes reservation in output queue of self, if no resources, tags request
  - pass request on to higher layer
- Accept bandwidth release from link neighbor
  - pass release on to higher layer
- Release local bandwidth reservation

Admission controls in bridge

- allocate channel using GMRP?
- forward bandwidth requests/releases to talker if first request/last release
  - respond directly without forwarding if already routing channel
- forward bandwidth responses to listener

ResE inter-bridge protocol

- in parallel with best-effort bridge protocols

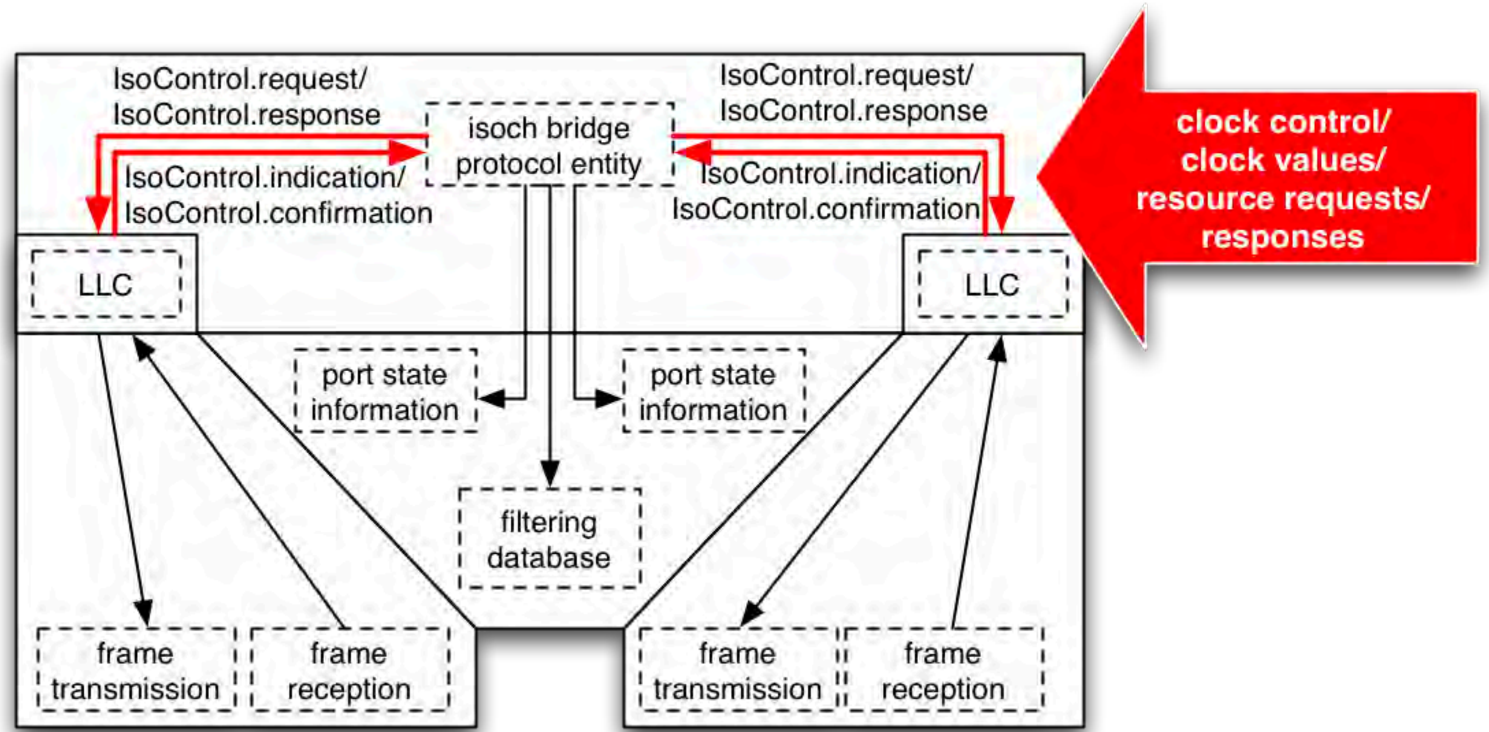

Isochronous transport

- Request transmit of isochronous packet
  – DA, SA, data, cycle “n”
- Receive isochronous packet
  – DA, SA, data, cycle “n”

Isochronous frame relay

- in parallel with best-effort frame relay

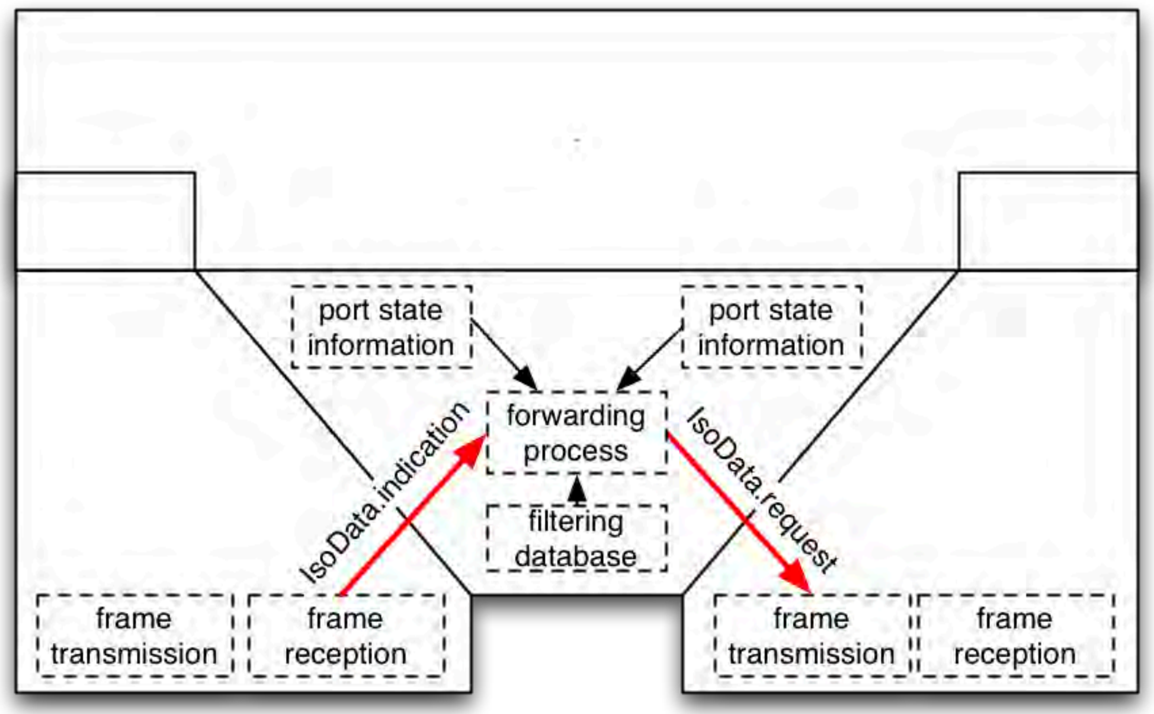

All services in MAC?

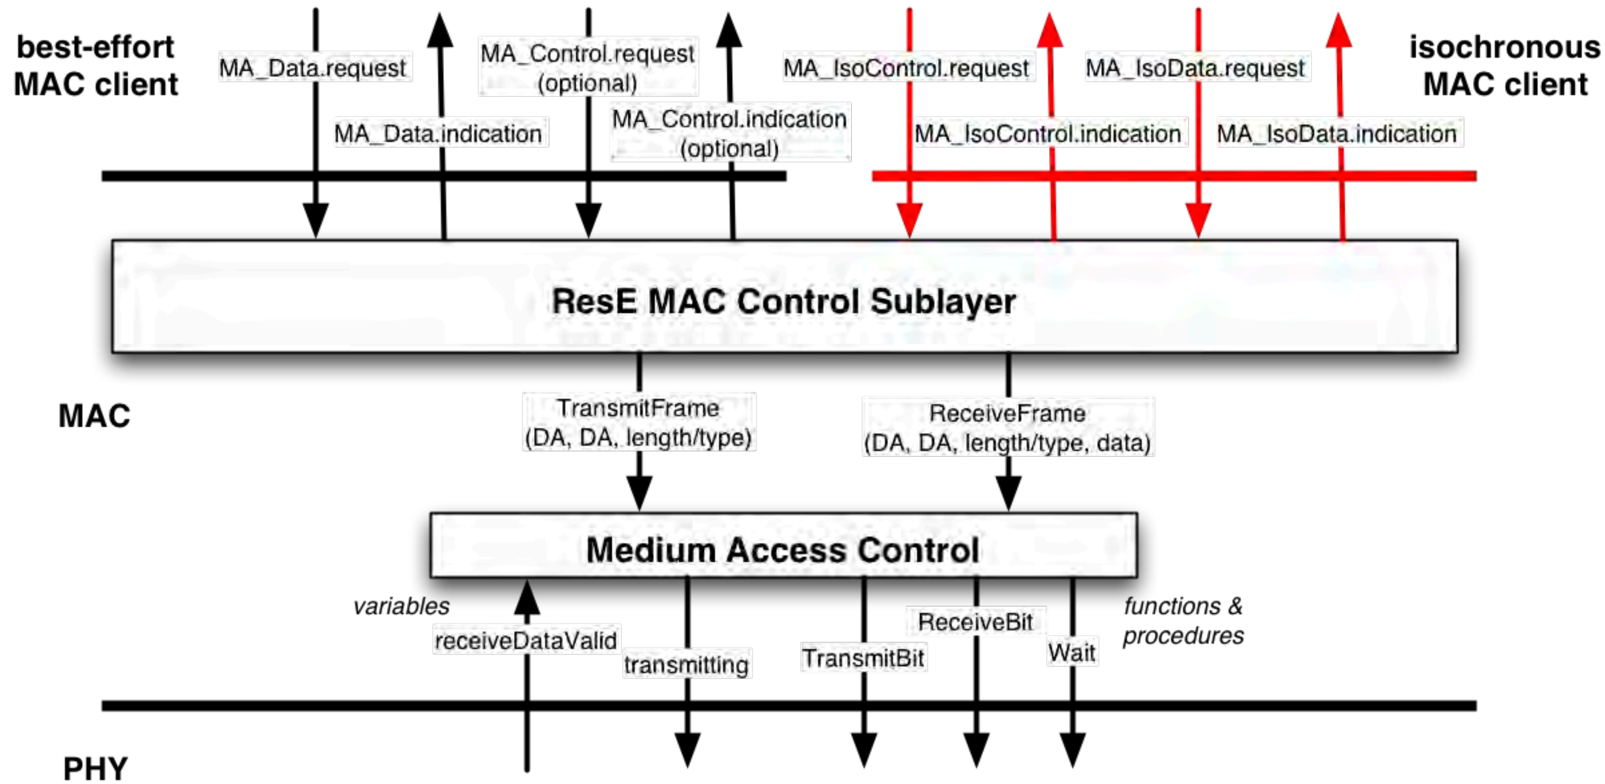

MAC-based services

- Advantages:
  - All best-effort services/protocol stack unchanged
  - New services totally in parallel
  - Similar to 802.11 and 802.15.3 methods
- Problems
  - reinvent registration and control services that may already be defined (e.g., GARP-based services)
  - uhh ... don’t we need bridges?

All services above MAC?

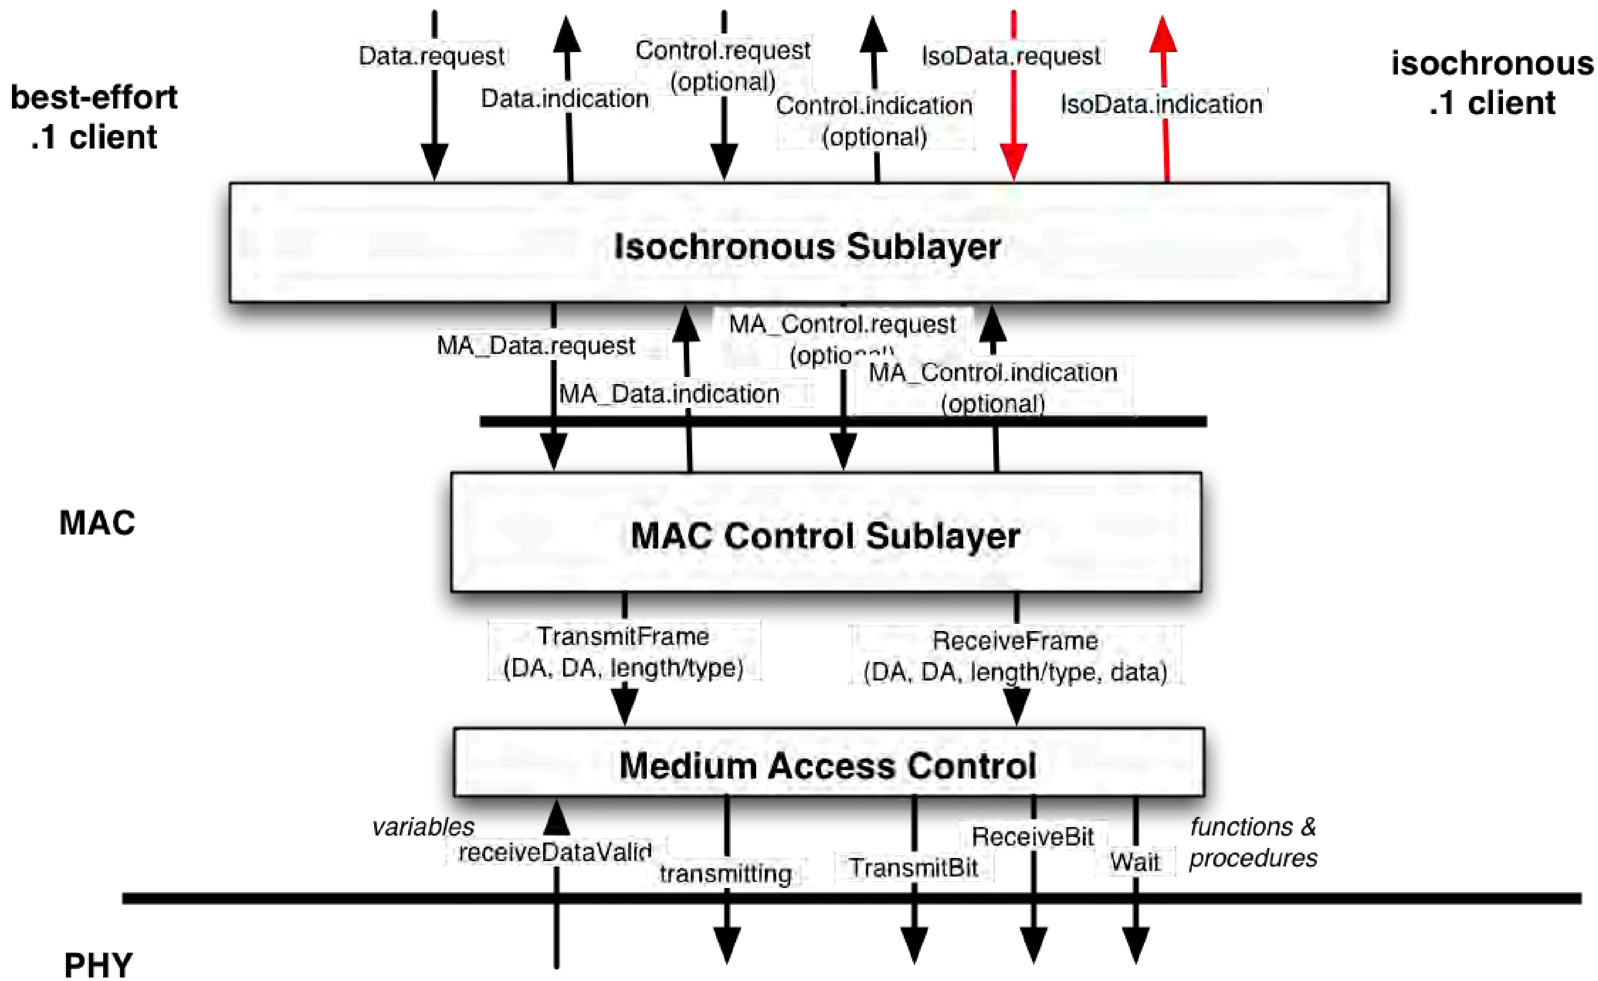

Bridge-based services

- **Advantages**
  - may be easier to specify
  - queues and scheduling concepts already in 802.1D
  - admission control protocols are similar to GARP
- **Disadvantages**
  - non-bridge devices need many of the services as well
  - odd layering for precise timing (how does something above LLC see bit timing?)

How about a combination?

- put admission control services into 802.1D
- put isochronous transport services into 802.3
- share synchronization services

Thank you!

ResE dependencies to existing standards and their complexities

Yong Kim
Broadcom

802.1 Dependencies –
802.1D Bridge

- Residential Ethernet is compatible and deemed to be complaint to 802.1D Bridges for existing services.
  - Spanning Tree compatibility
  - Compatible to the generally accepted 8 bridge hop
- Likely to require MAC services addition to support time sensitive frame forwarding function.

802.1 Dependencies – 802.1Q VLAN

- **802.1Q VLAN Tag/802.3ac Ethernet VLAN Tag**
  - **Compatible with ResE**
- **802.1Q-Rev**
  - **Stacked Q-Tag**
- **802.1ad Provider Bridging**
  - **Stacked Q-Tag and control plane protocol management**
- **May adopt 802.1Q GARP protocol to provide Residential Ethernet BW registration, although the favorite now is to do a subset of RSVP at layer 2.**
- **Residential Ethernet is compatible, and embrace 802.1Q framework without effecting any changes (except.. the GARP thing..)**

Ethernet Frame

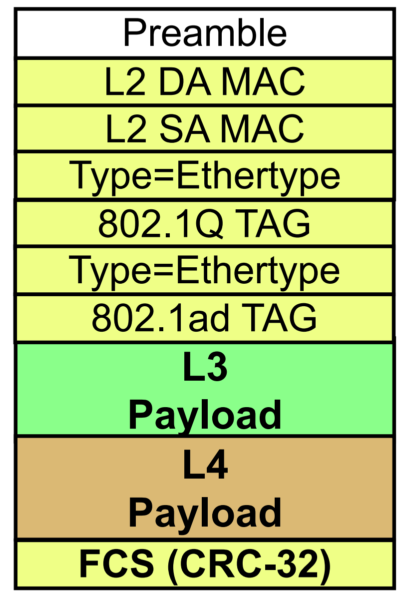

802.1 Dependencies - 802.1ab MAC Connectivity Discovery (LLDP)

- Single Packet Advertisement to the connected neighbor
- "Discovery", not Discovery-&-Configuration.
  - ITU work in process to do Discovery-&-Configuration of IP Phones
- Upper Layer protocol available (uPnP)
- Residential Ethernet MAY use 802.1ab for neighbor discovery (device type) and provide this information to the upper layer mgmt and application – in this case, no changes are expected to 802.1ab.

802.1 Dependencies - 802.1X-Rev Port-Based Network Access Control

- **Port-based authentication**
- **Residential Ethernet is compatible to 802.1X-Rev and does not effect any changes.**

802.1 Dependencies - 802.1AE MACSec & 802.3as Frame Extension

- **802.1AE MAC Security motivated and address pent-up demand for other encapsulation allowances.**
- **Current thought is to leave MAC DA, SA, and 802.2/ Ethertype in clear text, and encrypt the rest**
- **Consider ResE**
  - **MACSec interface to a MAC is clear text, therefore, there is no change (need to have key security mgmt to be done before end-point synchronization).**
  - **MAY want to leave ResE header in clear text and MAY allow for MACSec and/or upper layer DRM.**

MACSec Ethernet Frame

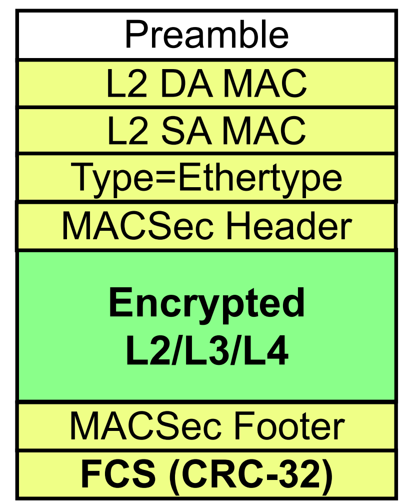

802.1 Dependencies - 802.1ag Connectivity Fault Management

- **This standard specifies protocols, procedures, and managed objects to support transport fault management. These allow discovery and verification of the path, through bridges and LANs, taken for frames addressed to and from specified network users, detection, and isolation of a connectivity fault to a specific bridge or LAN**

- **Residential Ethernet is compatible, and embrace 802.1ag framework without effecting any changes.**

802.3 Dependencies – 802.3af POE, 802.3 POE-Plus SG

- Power over Ethernet, more Power over Ethernet.
- ResE is compatible to 802.3af – only change expected is that if 802.3af implemented, then no optional parameters (of course, power negotiation will still happen).

802.3 Dependencies – 802.3 Auto-Negotiation

- Residential Ethernet expect to add its capability in the selector field
- Will adopt new auto-negotiation page(s) under consideration in 802.3an and 802.3ap.

802.3 Dependencies –
802.3ad Link Aggregation

- **Residential Ethernet is compatible to 802.1ad Link Aggregation and does not effect any changes**

802.3 Dependencies – 802.3ar Congestion Management TF

- **Current direction of 802.3ar's scope is orthogonal to Residential Ethernet.**
  - **Current scope is just add "shaping and rate-limiting" to Ethernet MAC services.**
  - **May help to reduce dropped frames in Residential Ethernet bridges for the best effort services (not really for TCP-IP).**
- **Residential Ethernet is compatible to 802.3ar and does not effect any changes**

Summary

Standard	Description	ResE Dependencies
802.1D	Bridging + RSTP	Time-sensitive Queue/MAC Services
802.1Q, Q-Rev, 802.1ad	VLAN, Q-tag, provider bridging	None
802.1ab	Link Layer Discovery	None – May be used.
802.1X	Port Authentication	None
802.1AE/ 802.3as	MACSec, Frame Extension	None
802.1ag	Conn. Fault Mgmt	None – May be used
802.3af, af+	POE, POE_Plus	Will require no options.
802.3	Auto-negotiation	Add selector encoding
802.3ad	Link Aggregation	None – May be used
802.3ar	Congestion Mgmt	None

Thank you!

`