Residential Ethernet Objectives, Requirements and Possible Solutions

Michael Johas Teener
mikejt@broadcom.com

Agenda

- Background
- Derivation of requirements
- Current approaches
  - and problems
- Proposed solutions
  - emphasis on bridge-based mechanism(s)

Residential Ethernet Background

Before we start ...

- Please defer all discussion about *where* in 802 the whole of “Residential Ethernet” needs to be specified
  - Some ideas on layering are at the end
- Treat “Residential Ethernet” as an 802-based system that will typically be implemented on top of “Ethernet”

Thanks!

Scope (from July '05 CFI)

- Residential Ethernet provides time-sensitive delivery between plug-and-play stations over reliable point-to-point full-duplex cable media. Time-sensitive data transmissions use admission control negotiations to guarantee bandwidth allocations with predictable latency and low-jitter delivery. Device-clock synchronization is also supported. Ensuring real-time services through routers, data security, wireless media, and developing new PMDs are beyond the scope of this project.

Purpose

- To enable a common network for existing home Ethernet equipment and locally networked consumer devices with time-sensitive audio, visual and interactive applications and musical equipment. This integration will enable new applications, reduce overall installation cost/complexity and leverage the installed base of Ethernet networking products, while preserving Ethernet networking services. An appropriately enhanced Ethernet is the best candidate for a universal home network platform.

ResE objectives for 802

- 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 (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)

- 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 and follow-ons)
  - … minimize configurations that “don’t work”
- backbone for IEEE 802.11 and IEEE 802.15.3 such that all QoS parameters are respected

ResE objectives for "other groups"

- Isochronous bridging to IEEE 1394
  - take advantage of wide experience and standardization
  - support IEC 61883 payloads (MPEG/DV/digital audio)
  - "DCAM" machine vision cameras (uncompressed)
- Perhaps even bridging to USB, MOCA, and higher level protocols such as RTP
  - additional “isochronous” payload types

... and one more thing ...

- All objectives must be met at “virtually no cost”
- For CE companies, this means that if it can be done using 100-base TX PHYs, then that’s what they want

And there are strong arguments that it can be done

Detailed requirements

- Compatible with existing* and planned 802.1 and 802.3 standards

  *at least those that support 100Mbit/sec full duplex and better
- Timing synchronization
  - minimize jitter-induced distortion
  - minimize wander-induced loss of data or excessive buffering
  - guarantee inter-stream synchronization of related data
- Guaranteed low latency for AV streams
  - typical network latency <2ms
  - scalable
    - deterministic based on network topology, not on implementation
- Guaranteed bandwidth for AV streams
  - once connection is made, QoS is consistent and unchanging

Summary of Jitter/Wander/Synchronization Requirements
(from G. Garner, to be presented 5/17/05 at Austin ResE SG Interim)

Requirement	Uncom-pressed SDTV	Uncom-pressed HDTV	MPEG-2, with network transport	MPEG-2, no network transport	Digital audio, consumer interface	Digital audio, professional interface
Wide-band jitter (UIpp)	0.2	1.0	50 μs peak-to-peak phase variation requirement (no measurement filter specified)	1000 ns peak-to-peak phase variation requirement (no measurement filter specified)	0.25	0.25
Wide-band jitter meas filt (Hz)	10	10			200	8000
High-band jitter (UIpp)	0.2	0.2			0.2	No requirement
High-band jitter meas filt (kHz)	1	100			400 (approx)	No requirement
Frequency offset (ppm)	±2.79365 (NTSC) ±0.225549 (PAL)	±10	±30	±30	±50 (Level 1) ±1000 (Level 2)	±1 (Grade 1) ±10 (Grade 2)
Frequency drift rate (ppm/s)	0.027937 (NTSC) 0.0225549 (PAL)	No requirement	0.000278	0.000278	No requirement	No requirement

Network Interface MTIE Masks

(from G. Garner, to be presented 5/17/05 at Austin ResE SG Interim)

MTIE - Maximum Time Interval Error - peak-to-peak phase variation for a specified observation interval, expressed as a function of the observation interval

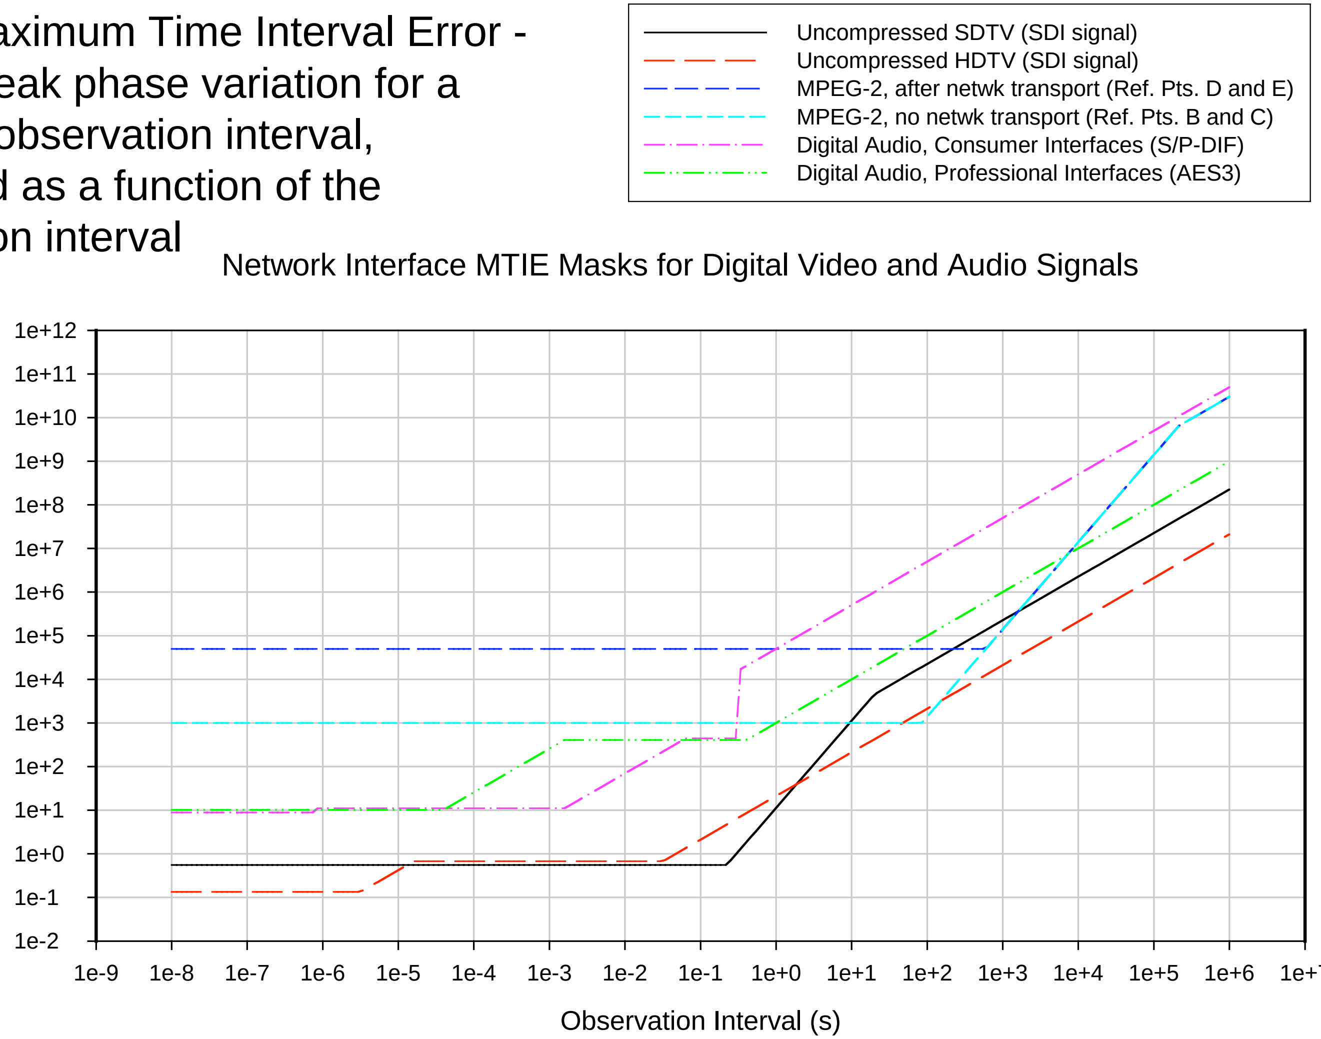

Inter-stream synchronization requirements

(from G. Garner, to be presented 5/17/05 at Austin ResE SG Interim)

- Tightly coupled audio (e.g., audio streams delivered to multiple speakers) - 10 μs
- Lip-synch - 80 ms
- Video animation with accompanying audio - 80 ms
- Additional examples and requirements given in [Ralf Steinmetz, Human Perception of Jitter and Media Synchronization, IEEE JSAC, Vol. 14, No. 1, January, 1996, pp. 61 – 72]
- Inter-stream synchronization difficult without network-wide time base

Latency requirements

- Typical human-interface control systems require <100ms response time
  - best if <50ms
  - guess at network budget of 25ms for round trip, or 12ms one way
  - the application(s) will use up the rest
- Worst case control system for home is playing/recording musical instrument(s)
  - not a small market!

Playing music requires timely feedback

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

Home recording

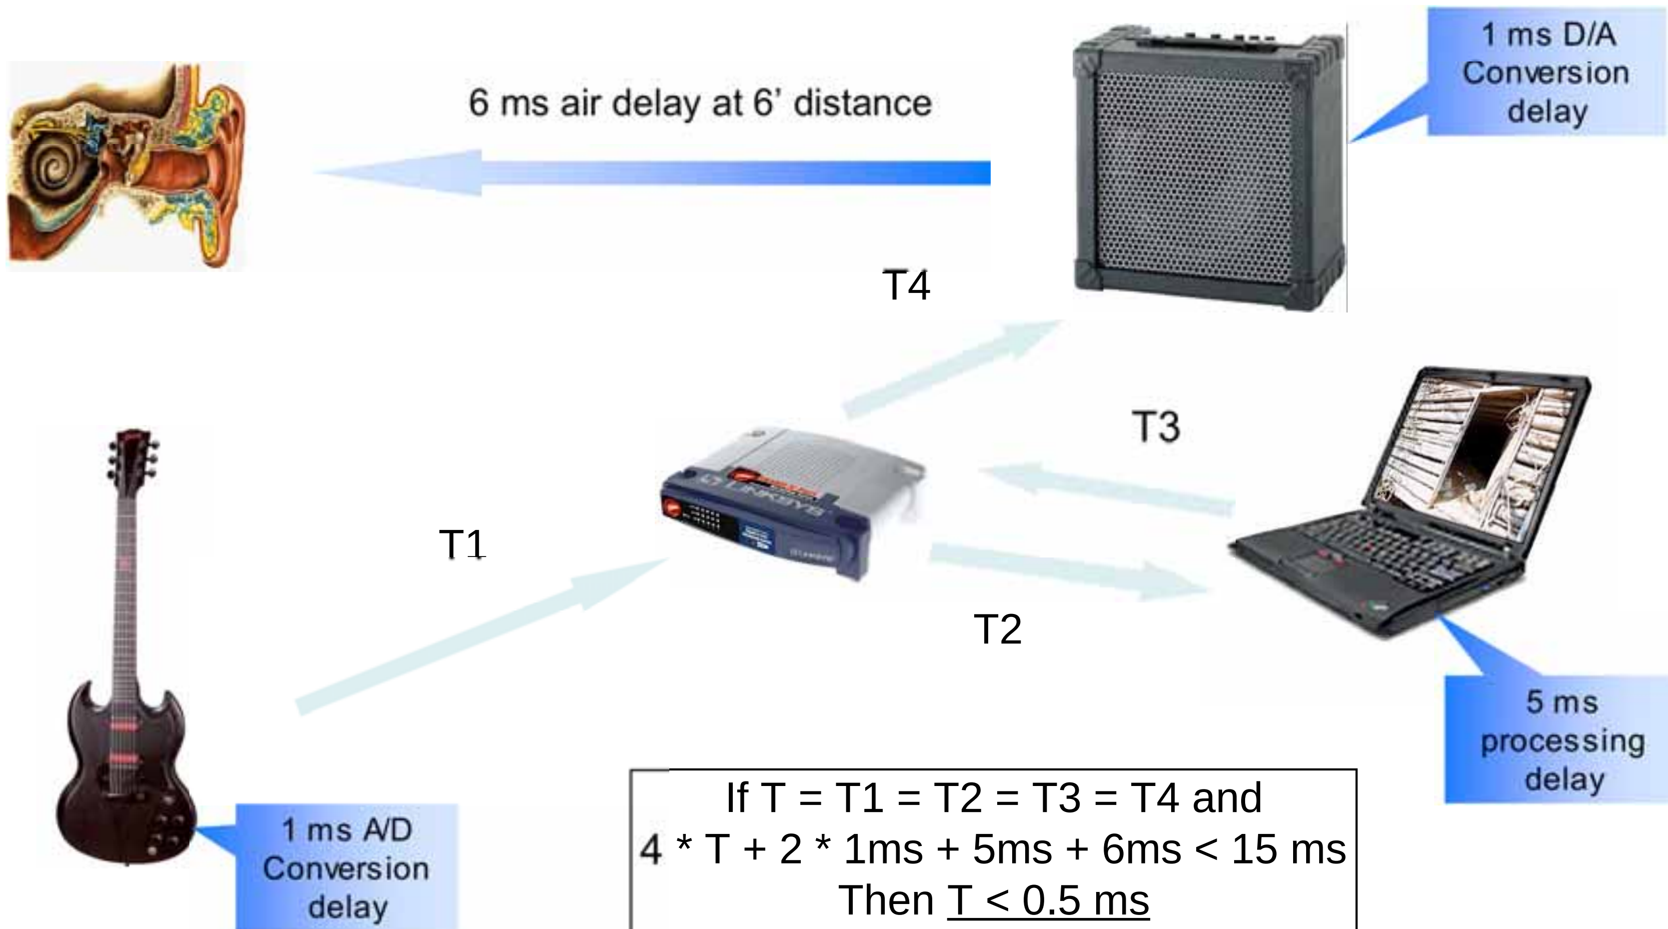

If T = T1 = T2 = T3 = T4 and
4 * T + 2 * 1ms + 5ms + 6ms < 15 ms
Then <u>T < 0.5 ms</u>
(Tn is total network delay included queuing inside networked device)

Garage jam session

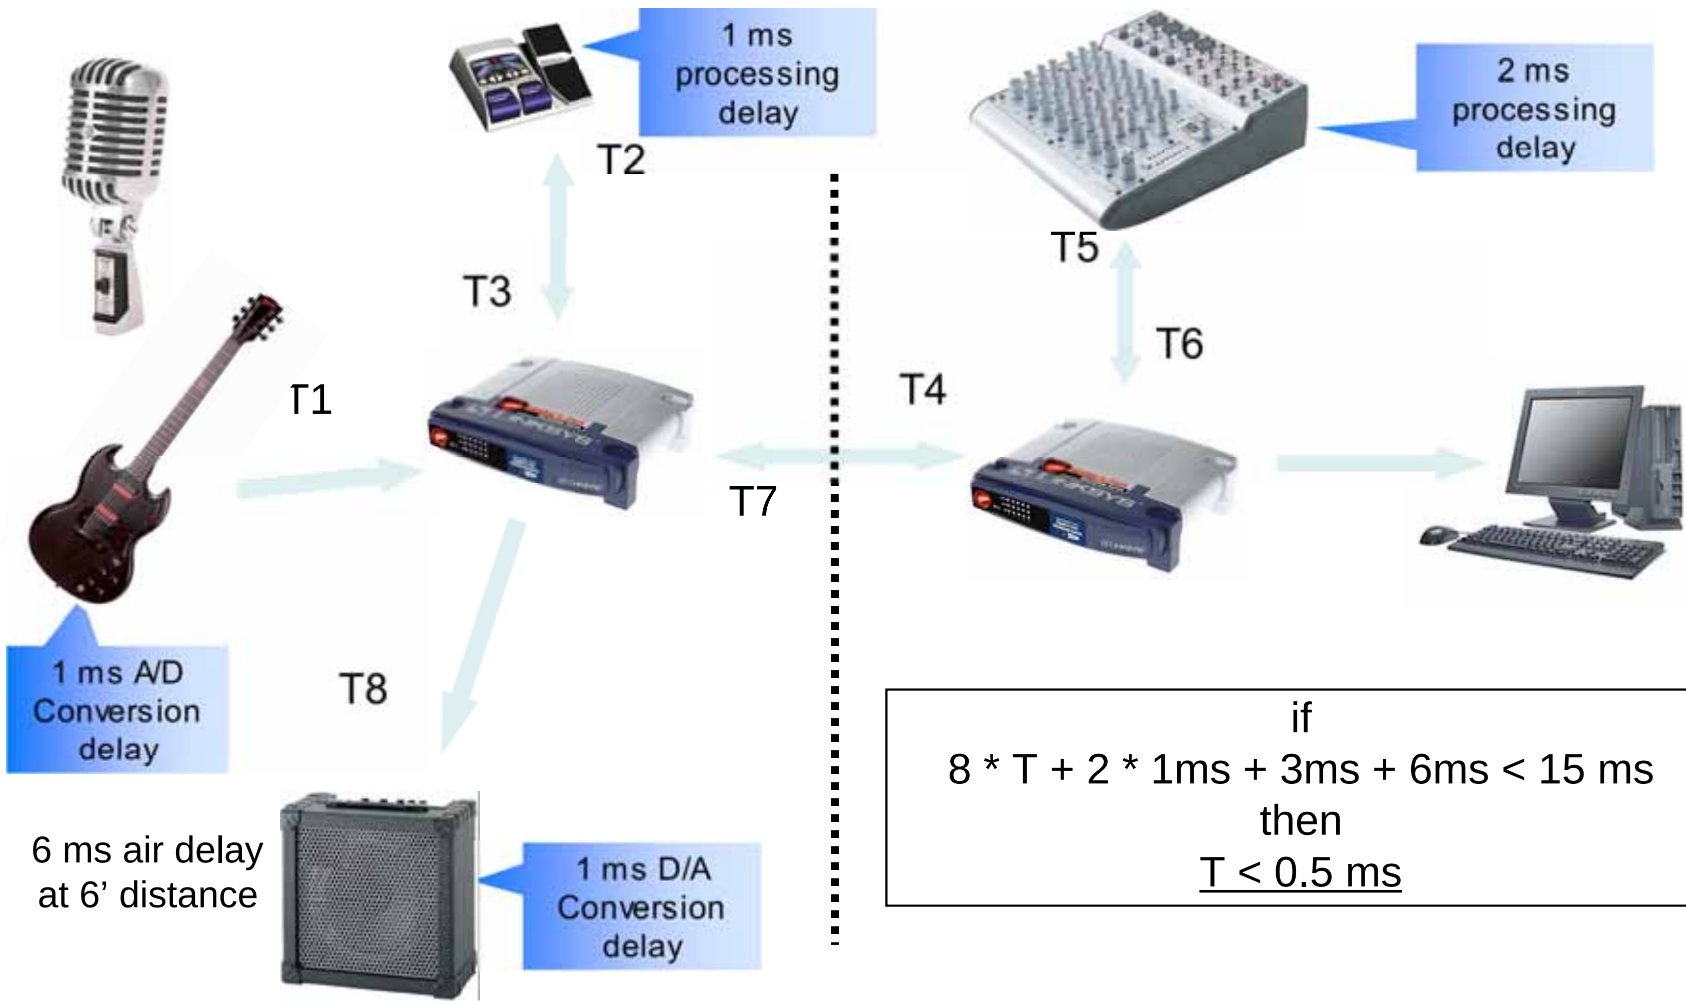

if
8 * T + 2 * 1ms + 3ms + 6ms < 15 ms
then
<u>T < 0.5 ms</u>

Guaranteed bandwidth for connections

- Typical statement made by CE representatives:
  - "If the network is reaching its limits, denial of service is preferable vs. allowing a new application to disrupt the audio or video quality of another application already running on the network." Jim Battaglia, Pioneer
- In other words, this is a desire for a connection-oriented QoS
  - "intserv" preferred over "diffserv"
- For applications where statistical QoS is acceptable, then the CE industry will use wireless

Manageability preserved

- Some percentage of bandwidth reserved for best effort traffic
  - proposing 25 %
- using priorities for higher level management as is currently proposed

Existing solutions and problems

”Higher layer" timing synchronization

- Timing synchronization can be done at higher layers
  - NTP and SNTP from IETF
    - slow to converge, insufficiently accurate
  - IEEE 1588
    - measures timing at start of packet launch on media
- 1588 requires that bridges participate
  - using priorities would not work
    - Dirk S. Mohl, “IEEE 1588 - Precise Time Synchronization as the Basis for Real Time Applications in Automation”
  - all would require a 1588 timing element and coordinator
  - drive up the cost of “CE” oriented bridges
- Much easier to do as a service of the MAC (or just above the MAC if the MAC does not introduce any delays)

Guaranteed QoS: Overprovisioning

- Proposal: provide enough network bandwidth to guarantee that AV streams get adequate bandwidth and bridge queues never overflow nor add excessive delay
- Problem: no way to enforce
  - How do you stop a consumer from plugging in a GigE-based PC and using it?
  - How do you guarantee that the consumer won't start plugging in legacy FE bridges that "mostly" work?
  - How do you guarantee that some future applications don't outstrip your overprovisioning assumptions?
    - e.g., uncompressed or lightly compressed video, trick play

Guaranteed QoS: add priorities

- Proposal: have AV streams run at a higher priority, combine that with overprovisioning
  - better than straight overprovisioning
- Problems:
  - still require that consumer has a managed network. E.g., unqualified configurations will sometimes work, sometimes not
  - still no guarantee for future applications
  - no guarantee that PC applications won't start using priorities
    - no rules there!

Proposed enhancements

- ResE SG has three proposals on the table
  - Allow a “hub-like-thing” to be used in GigE to switch streams in real time (Spyder)
  - Use a variation of the EPON timeslot management
  - Add real-time or scheduled priority gating to bridges
- The first two do not address how the enhanced QoS gets through bridges
  - And that *is* a requirement for a total solution
  - PERSONAL preference is to solve a problem just once
  - So that’s what I’m talking about here

Goals for bridge-based solution

- meet requirements
- minimal additions
  - use 802.1 and 802.3 mechanisms and layering
- easy to implement
- easy to specify
- scalable
  - useful for all MAC/PHY combos that have similar performance (>= 100Mbit/sec, full duplex)
  - useful in all kinds of cascaded bridge environments

Bridge-based synchronization?

- Run 1588-like protocol at the MAC layer or just above
  - precise measurement of launch of packet
  - implement on output port *after* any queuing
  - could be MAC addition, but requires change to MAC services or another MAC control layer
- Run 1588 master selection/grand master selection at LLC level as a bridge protocol

802.1-based admission control

- Control protocol using some much simplified RSVP - “SRP?”
- Specific labeling of "AV" streams so admission control can be enforced
  - Use “well known” multicast addresses
    - AV streams are frequently multicast anyway
  - Interesting project for protocol to assign addresses without using a centralized server
- Enforcement of limits on bandwidth usage at output ports
  - limits set by SRP *and* by absolute limit of 75% of port bandwidth capabilities

Pacing

- AV streams are paced using a common rate at all output ports
  - propose 8kHz pacing as used by IEC 61883 higher layers used by 1394
- Pacing is enforced at bridge output ports as well
- Spreads out traffic to reduce “bunching” at queues
- No increase in maximum latency
- Allows fixed maximum buffer size at each output port
  - No matter how many bridges in a path, the buffer will *never* overflow

Conclusion and discussion

- Timing synchronization, guaranteed low latency and guaranteed bandwidth are needed for CE-based streaming data
- Existing methods of carrying streaming data on bridged Ethernet networks either require management or fail to have guaranteed QoS
- Modest additions to bridges and DTE will meet all requirements
  – a proposed architecture has been examined by implementers and blessed as “virtually no cost”
  – the Devil’s in the details, however, so ...
- Please contribute!

Thank you!