Residential Ethernet Study Group (ResE)
September 2005 joint ResE/802.1 interim meeting
Submitted by Alexei Beliaev (ieee802@brainpot.com), Gibson Labs

September 29, 2005
Meeting starts at 9:20 am.
About 25 people are present.

Mr. Michael J. Teener (ResE group chair) starts the meeting with the announcement of the meeting agenda and goals. The agenda includes several presentations and discussion of future work of the study group.

Goals
- Study group status review
- Discussion of moving to 802.1
- Receive technical presentations
- Review objectives
- Plan future work
- Build consensus!

Mr. Teener presents information about the study group reflector and the web. He mentions that the Residential Ethernet Interest Group is also available on Yahoo website. The Yahoo group is useful for storing public materials for ongoing discussions. Mr. Teener presents the meeting ground rules and gives overview of the IEEE standard structure, standard process, by-laws and rules. Then he reads IEEE patent policy and reminds participants of the meeting about inappropriate topics during discussions.

The purpose of the Study Group is to generate PAR, 5Criteria and gain their approval from the working group. Non-function of the Study Group would be choosing a solution.

The next presentation is provided by Mr. Mick Seaman. The presentation title is "Residential 802.1 bridging".

This is a transition meeting. We may have to deal with a mix of working groups styles. The summary:
- Purpose of PARS
- Some .1 style observations
- Objectives for this meeting
- Some related .1 goals and developments
- .1 compatible technical architecture
- .1 compatible document architecture

The purpose of PARs is to authorize the project (i.e. writing or amendment of one document), set the scope as a defense against delays due to scope creep, endless arguments about what was meant or agreed, mindless horse trading. The PAR is not required for exploratory work.

Some .1 style observations

- consensus based and driven
    - recourse to formal voting to decide issues is rare
    - never in task groups
    - years of avoiding short-term manipulation
- The production of the standard is the process
    - The test of consensus is draft balloting
    - The final test of consensus is sponsor ballot (?)
    - Very rarely vote about anything else
- Strong preference for very highly focused PARs
    - Fundamental project management
    - Avoiding fears, public positioning, delays
    - What you see is what you get
    - Continuous step by step development

Objective for this meeting is to develop proposal for .1 considerations

- To support Residential Bridging applications
- What existing 802.1 documents need to be modified
- What new 802.1 standards are required
- What additional supporting standards are required to or assumed

Some .1 goals

- Broad applicability across applications (clearly identify common base solution)
- Leave no undercutting economic alternatives (what alternatives exist)
- Don't mortgage the future (very wary of architectural oddities that constraint future developments)
- Don't destroy the present and its growth (continuous compatible development)

Some related .1 developments

- Rapid reconfiguration (in .1D/.1Q) – minimize complex binding of resources to paths
- P802.1ad & 802.1Q-REV (drop precedence and flow metering – "abv: for admission control")
- P802.1AQ Shortest path bridging (remove single spanning tree requirement, without complex management, may supplant GVRP/MVRP in some areas)
- P802.1ag Connectivity Fault Management (simple in-band tools for checking connectivity)
- P802.1ak Multiple Registration Protocol ("abv: GARP replacement")
- Other potential "class of service" proposals ("abv: congestion management")

.1 compatible technical architecture

Strong emphasis on hard shell/soft core network

- very simple class-based forwarding functions
- Admission control / rate control / policing at edge
- No flow control state within network

QoS as a set of successive improvements and approximations

- 100% loading never achieved
- Performance bounds can be realized

Technical architecture elements

Bridge

- Performance ("abv: follow spirit")
- Flow metering and drop precedence (admission control and traffic profile enforcement)
- Class based queuing
- Queue service algorithms

End station

- class based admission control
- Frame class/priority marking
- Admission control / traffic profile enforcement (not heavy tailed or even Poisson)
- 'Management' / 'user' reporting

Exists in .1: Use of traffic classes, regeneration of priorities for admission control/traffic profiling

.1 compatible document architecture

- Maximize leverage of existing applicable items (Piece parts and placeholders already in .1Q)
- Set out the whole application and assumptions
  - Possibly a Recommended Practice
  - Possibly a .1Q informative Annex
  - Clarify the need for non-transport, non-802 elements
  - A possible record of technical architecture
- End-station behaviors
  - Possibly in .1Q, possibly separate standard
  - Opportunity to leverage recent .3 rate control work

Document architecture elements

- Bridge performance (.1D 16.1, 16.2 – much to be done)
- Flow metering (and policing)
  - .1Q 8.6.5 and potential Annex addition
  - Include definition of what bandwidth = X means

- Use of classes and class-based queues
    - .1Q Annex G additions and 'application profiling'

- Queue service (transmission selection) algorithms
    - .1Q 8.6.8 and potential Annex addition

- Residential real-time 'domain' identification
    - Document where?

- Admission control protocol
    - Large separable item once metering/policing units decided

Why focus on documentation
- Until written down, what is done is not clear
- Until written and integrated, impacts are not clear
- Until clear, false fears and hopes dominate
- Most of every project that adds or modifies occurs after it is thought technically 'complete'
- Opportunities for staged completion
- Begin with the end in mind

Mr. Seaman provides username and password for 802.1 on-line documents access.

Discussions

Mr. Teener offers to use this afternoon (09/29/05) for posted presentations and tomorrow afternoon (09/30/05) - for planning of November meeting.
Mr. Grow points out that the room for November 2005 RE meeting must be reserved through Steve Carlson. He also assumes that 802.3 will not renew RE SG during November meeting.
Mr. Seaman informs the meeting that .1 usually would not rush the PAR. But … class specification can be scoped quickly.

Reflector discussions will continue using the .3 reflector.

[Break until 1:00 pm]

Mr. Geoffrey M. Garner of Samsung Electronics gives presentation "Meeting RESE requirements: a simulation study".
In his presentation, Mr. Garner summarizes ResE audio/video application performance requirements, shows simulation results and conclusions.

Delay requirements of applications are based on documents from ITU-T and interactive audio/musical applications that were described in presentations from previous ResE meetings (end-to-end delay requirement for ResE: <2ms for up to 7 hops).

Jitter/wander requirements of applications are based on requirements of uncompressed SDTV and HDTV, MPEG-2, and professional and consumer digital audio interfaces.

Mr. Garner provides network interface MTIE masks of applications showing that uncompressed video has the most stringent requirement. The second stringent is digital audio.
Mr. Garner describes the common application MTIE mask as the lower envelope of the application MTIE masks. Then he provides simulation results for several cases of transporting of multiple time-sensitive traffic streams over current Ethernet.
The conclusion includes requirements for separate timing synchronization, subscription protocol for network resources, admission control, and methods to describe and normalize the time-sensitive traffic (PAL).

Mr. Felix Feifei Feng of Samsung Electronics gives presentation "Subscription protocol and admission control in Residential Ethernet"

Deterministic low latency and low jitter for data delivery can be provided only if the availability of network resources is guaranteed and intermediate bridges are appropriately configured along the entire transmission path. This requires a subscription protocol for explicit negotiation for network resources.
Subscription protocol (SRP) is employed to manage isochronous streams by updating isochronous filtering database of each bridge along the stream path.
SRP signaling is based on GARP. 'Listener' uses GARP to show its intention of joining isochronous stream. 'Talker' sends reservation signaling towards its 'listeners'.
Mr. Feng shows a picture describing GARP-based SRP architecture and provides examples of signaling.
Then Mr. Feng illustrates admission control requirements for a case with non-preemptive strict priority schedulers.

Mr. Seaman suggests specifying admission protocol enforcement ('enforcers') at network boundaries.

Mr. Geoffrey M. Garner of Samsung Electronics presents "Analysis of clock synchronization approaches for Residential Ethernet".

This is a slightly modified version of presentation for the next IEEE 1588 meeting on October 10-12, 2005.
Mr. Garner describes the application reference models based on example for MPEG-2 transport and MTIE masks for end-to-end requirements of SDTV, HDTV, MPEG-2, and digital audio application interfaces.
Then he lists different synchronization approaches for frequency and phase adjustments.

The presentation includes simulation results for some of the listed synchronization approaches for cases with and without additional clock noise and phase granularity.

After presentation discussions touch the implementation cost for the synchronization approaches and relationship between 1588 and ResE parameters.

Meeting adjourned at 5:20 pm.

September 30, 2005
Meeting starts at 9:20 am.
About 30 people are present.

Mr. David V. James gives a presentation on "Clock synchronization"
The presentation illustrates the concept of the house reference clock and shows the application cases where precise timing synchronization is required.
Mr. James speaks about cascade synchronization, and then he describes some possible grand-master selection protocols.
The synchronization protocol is based on periodical (~10 ms period) messages with pipelined computation for allowing inexpensive software implementations of the protocol.
Following slides describe some design details.

The following discussion is centered on ability to adjust the period of synchronization messages in the range from 1 to 100 ms, and on importance of pipelined computation.

The following presentation by Mr. David V. James is on "Bursting and bunching considerations".
The presentation assumes a topology with seven hops that may be non-realistic but is useful for illustrating the worst-case scenario that unlikely happens in simulations.
Then Mr. James describes pacing and shows how it works as a tool against bursting effects. A 'plan-ahead smart' scheduler operation is illustrated.

The discussions after presentation include the question of the amount of built-in memory that every switch would require and an approach to supporting different latency requirements using different traffic classes is discussed.

Admission control consists of three following components:

1. Ingress metering (possibly in bridge)
2 Transmit shaping (possibly in the end-station)
3 Subscription signaling

In the bridge, incoming traffic shaping is per port. Buffering inside bridge is required for reshaping the stream and having bursts undone. Another approach is to reshape the traffic based on class. Unfortunately, this doesn't solve bunching within a class. Credit-based shaper (set priorities based on the target transmission time) would be equivalent to putting a maximum rate scheduler on the output port.

Defining two classes of streaming traffic may be better than trying to pack everything into one streaming class. A class based shaper must police the best effort traffic. Metering requirement for data that comes from outside are required.

Mr. Yong Kim of Broadcom makes a presentation “802.1 & ResE”.

The presentation agenda includes the following topics.

- Examine the top level ResE objectives
- Examine existing and on-going 802.1 standard work
- Examine “pent-up” need to do something within 802.1
- Identify how ResE objectives may fit into 802.1 framework
- Explore other 802.1 standards that may apply to ResE work

ResE requirements affect forwarding, admission control, timing synchronization, and recommended practice. Current 802.1 standard works affect MAC services, a bridge edge, and a bridge core.

ResE requirements

- Forwarding
    - Queue/flow definition
    - Scheduler (beyond strict priorities) enhancements
    - Congestion handling

- Admission Control
    - Establish controlled environment (“ResE cloud”)
    - Control protocols (manage 802.1 controls, metering, scheduling, etc.)

- Time Sync
    - Network timing awareness (bounded delay, end-point synchronization)

- Profile, Rec. Practice
    - VLANs, class usage
    - Multi-path/STP/link aggregation
    - Rate control in MAC (congestion management)
    - Management of scheduling queues
    - Class (priority level) regeneration

Then Mr. Kim described existing 802.1 features and the recommended 802.1 features. The recommended features include

- Egress shaping (MAC services)
- Egress shaping and scheduler definition for min and max bandwidth guarantees for egress classes
- Ingress meter configuration specification
- Scheduler enhancements (rate controlled priority queuing, scheduler algorithms)
- Discovery and GARP enhancements

- Time synchronization support
- CE end-point profile and bridge profile

We can add minimum bridge performance guarantee to recommended ResE Features. It is desirable to provide some terms and definitions.

Anticipated changes to 802.1q (non-controversial)

- Add an additional item to 8.1.6 that explains that queue scheduling configuration supports traffic expediting
- Rewrite 8.6.8 to describe new parameters and scheduling algorithm
- Admission control (ingress metering)
- GARP extensions to add ResE features (reservation)
- Administrative edits for consistency

Possible additions to 802.1 (controversial)

- Time Sync (MAC services, forwarding)
- Discovery (LLDP annex or something else)
- DRM considerations (802.1X, MAC security (AE/AF), or outside)
- ResE profile annex (with attention to usability for other purpose)

The group needs to decide the right approach for time synchronization standard (maybe 1588 would add it). Encryption can affect timing protocol. Sending timing information "after" is solving the problem. We need to sufficiently (broad) capture possible timing solutions, broad range of applications.

Mr. Kim provides a list of other 802 standards that would be applicable to ResE. Then he lists major future work and discussion items. It is mentioned that shortest path bridging and link OAM are also important.

Mr. Geoffrey M. Garner of Samsung Electronics makes a presentation "Delay and delay variations; Simulation results for multi-hop conventional Ethernet cases with bursting/bunching".

The presentation shows in details and analyses the results of simulation of several cases of competing constant bit rate (CBR) streams going through up to seven hops of Ethernet network with priorities (traffic classes).

Mr. Garner presents a formula (rule of thumb) for estimating the worst case end-to-end delay and delay variation. He shows that the simulation results comply with the formula.

It was mentioned during discussions that in the case of real applications, multiple streams are normally synchronized with the same clock.

Mr. Felix Feifei Feng of Samsung Electronics presents “Simulation results on 802.1p and pacing approaches”.

The presentation shows results of several simulation scenarios of conflicting time-sensitive and best-effort data streams going through Ethernet network with multiple bridges. The presentation compares results of cases with conventional Ethernet bridges versus cases with bridges that implement egress traffic shaping.

Mr. Feng makes a conclusion that in the absence of the absolute worst case, switching with shaping (pacing) introduces larger end-to-end delays and smaller delay variations. Jitter is also not accumulated along the multi-hop path for switches with pacing.

The following discussions address the issues of building up the required .1 attendance and getting the voting rights to ResE group participants.
Then the group discusses coordination of preparation for the November plenary meeting and how to avoid possible scheduling conflicts.

The meeting adjourns at 4:00 pm.