

IEEE 802 Interim Meeting
 September 2010, York, United Kingdom
802.1 Meeting Minutes, 13-16 September 2010

Attendees

| Last Name | First Name | Middle Name | Affiliation |
|------------------|-------------------|--------------------|---|
| Alon | Zehavit | | Nokia Siemens Networks |
| Ashwood-Smith | Peter | | Huawei Technologies Co. Ltd |
| Boiger | Christian | | Hochschule Deggendorf University |
| Bottorff | Paul | | Hewlett-Packard Development Company, L.P. |
| Bragg | Nigel | | Ciena Corporation |
| Carlson | Craig | | QLogic Corporation |
| Congdon | Paul | | Hewlett-Packard Development Company, L.P. |
| Crupnicoff | Diego | | Mellanox Technologies |
| DeSanti | Claudio | | Cisco Systems |
| Ding | Zhemin | | Ericsson AB |
| Eastlake, 3rd | Donald | | Stellar Switches, Inc. |
| Farkas | Janos | | Ericsson |
| Fedyk | Donald | | ALCATEL-LUCENT |
| Finn | Norman | | Cisco Systems, Inc. |
| Garner | Geoffrey | | Samsung |
| Ghanwani | Anoop | | Brocade |
| Goetz | Franz | | Siemens |
| Gravel | Mark | | Hewlett-Packard Development Company, L.P. |
| Gray | Eric | | Ericsson |
| Gu | Yingjie | | WG802.1 HARMAN INTERNATIONAL INDUSTRIES, INCORPORATED |
| Gunther | Craig | | |
| Haddock | Stephen | | Extreme Networks |
| Hazarika | Asif | | IP Infusion Inc. |
| Jeffree | Anthony | | hp/broadcom |
| Kamath | Daya | | Blade Network Technologies |
| Kashyap | Prakash | | WG802.1 |
| Kiessling | Marcel | | Siemens |
| Kumar | Vinod | | Tejas Networks Ltd |
| Lemon | John | | ADTRAN Inc. |
| Li | Lin | | ZTE Corporation |
| Li | Yizhou | | Huawei Technologies Co. Ltd |
| Lynch | Jeff | | IBM |

| | | | |
|------------|------------|-------|--|
| Mack-Crane | Thomas | | Huawei Technologies Co. Ltd |
| Magee | Anthony | | ADVA Optical Networking Ltd. |
| Messenger | John | | ADVA Optical Networking Ltd. |
| Mizutani | Masahiko | | Hitachi, Ltd. |
| Multanen | Eric | | Intel Corporation |
| | | | HARMAN INTERNATIONAL INDUSTRIES, INCORPORATED |
| Olsen | David | | |
| Pannell | Donald | | Marvell Semiconductor |
| Parsons | Glenn | | Ericsson |
| Pearson | Mark | | Hewlett-Packard Development Company, L.P. |
| Pelissier | Joseph | | Cisco Systems |
| Raeber | Rene | | Cisco Systems |
| Rouyer | Jessy | | ALCATEL-LUCENT |
| Sabony | Mazda | | Atlantik |
| Sajassi | Ali | | Cisco Systems |
| Saltsidis | Panagiotis | | Ericsson AB |
| Seaman | Michael | | Mick Seaman |
| Shao | Hong | | ZTE Corporation |
| Sharma | Rakesh | | IBM |
| Stanton | Kevin | B | Intel Corporation |
| Sultan | Robert | | Huawei Technologies Co. Ltd |
| Teener | Michael | Johas | Broadcom Corporation |
| Thaler | Patricia | | Broadcom |
| Ting | Ao | | ZTE Corporation |
| Unbehagen | Paul | | WG802.1 |
| White | Martin | | Marvell Semiconductor, Inc. |

Interworking Minutes

Monday, September 13

AM

Meeting called to order at 9:15 AM.

John Messenger welcomed us to York and gave some introductory remarks.

Presented patent policy slides. No response to the call for patents.

Stephen Haddock presented the preliminary agenda for the week.

(see <http://www.ieee802.org/1/files/public/docs2010/admin-interworking-agenda-0710-v1.pdf>)

802.1Qbe MIRP: Norm Finn conducted comment resolution of 802.1Qbe Multiple I-SID Registration Protocol sponsor ballot. There were comments received from two balloters, one with an approve vote, and one with a disapprove. All comments are resolved and a recirculation ballot will be conducted. Options for how to integrate this with Q-Rev were discussed. Current plan is to hold the MIRP document after the next recirculation so that we can run one final recirculation after Q-Rev completes. Need to put in a comment that says “Retarget this to Q-Rev when Q-Rev.”

802.1Qbc RCSI: Ben Mack-Crane provided an editor’s report and conducted comment resolution of 802.1Qbc-d1.2 PB Remote Customer Service Interface sponsor ballot. The comment dispositions are available in the bc-drafts folder on the web site.

<http://ieee802.org/1/files/public/docs2010/bc-mackcrane-editor-report-0910.pdf>
<http://ieee802.org/1/files/private/bc-drafts/d1/802-1bc-d1-3-Sponsor-dis.pdf>

802.1Q-REV VLANs: Tony Jeffree conducted comment resolution of 802.1Q-REV-d1.2 Virtual Bridged LAN sponsor ballot.

<http://www.ieee802.org/1/files/private/q-rev-drafts/d1/802-1Q-REV-d1-2-pdis-v1.pdf>

Break for lunch at 11:55am.

PM

Resumed at 1:30pm.

802.1Q-REV VLANs: Tony Jeffree completed comment resolution of 802.1Q-REV-d1.2 Virtual Bridged LAN working group ballot.

802.1AC MAC Service: John Messenger provided an editor’s report and conducted comment resolution of 802.1AC-d1.1 MAC Service Definition ballot.

<http://www.ieee802.org/1/files/public/docs2010/ac-messenger-editor-report-0910-v01.pdf>
<http://ieee802.org/1/files/private/ac-drafts/d1/802-1ac-d1-1-dis-v1.pdf>

Adjourned at 5:30pm.

Tuesday, September 14

Meeting called to order at 9:00am.

Presented patent policy slides. No response to the call for patents.

802.1AC MAC Service: John Messenger continued comment resolution of 802.1AC-d1.1 MAC Service Definition ballot.

Break for lunch at 12:15am.

PM

Resumed at 1:45pm.

802.1AC MAC Service: John Messenger completed comment resolution of 802.1AC-d1.1 MAC Service Definition ballot.

802.1Qbf SegProt: Bob Sultan presented an editor's update and conducted comment resolution of 802.1Qbf-d0.2 PBB-TE Infrastructure Segment Protection task group ballot. Proposed comment dispositions are available in the bf-drafts folder on the web site.

<http://ieee802.org/1/files/public/docs2010/bf-sultan-editors-report-0910-v01.pdf>

<http://ieee802.org/1/files/private/bf-drafts/d0/802-1bf-d0-4-dis-v00.pdf>

Liaison: Stephen Haddock showed two liaison letters that have been received since the San Diego plenary meeting:

A liaison from the MEF regarding the Mobile Backhaul Implementation Agreement Phase 2.

<http://www.ieee802.org/1/files/public/docs2010/liaison-itut-sg11-ls49-0910.doc>

A liaison from ITU Q11 regarding a proposal to request an Ethertype for Q.flowstatesig use.

<http://www.ieee802.org/1/files/public/docs2010/liaison-mef-34119-000mbh-0710.doc>

Adjourn at 3:45pm.

Wednesday, September 15

AM

A joint meeting was held with Data Center Bridging task group. See DCB minutes below for details.

The joint meeting adjourned at 12:05 AM.

PM

Interworking task group meeting resumed at 1:45 PM.

802.1aq SPB:

Mick Seaman discussed the documents he posted on the 802.1aq agreement protocol and the proofs of the loop-free forwarding rules:

<http://www.ieee802.org/1/files/public/docs2010/aq-seaman-agreement-protocol-0910-v2.pdf>

<http://www.ieee802.org/1/files/public/docs2010/aq-seaman-link-state-0910-v4.pdf>

Don Fedyk presented an editor's status report on 802.1aq-d3.0 Shortest Path Bridging. Final comment dispositions on the d3.0 working group ballot are available in the aq-drafts folder on the web site.

<http://ieee802.org/1/files/public/docs2010/aq-fedyk-SPB-Update-summary-0910-v1.pdf>

<http://ieee802.org/1/files/private/aq-drafts/d3/802-1aq-d3-0-dis-v2.pdf>

New Work:

Presentations on Resilient Network Interconnect from Norm Finn and Zehavit Alon:

<http://www.ieee802.org/1/files/public/docs2010/new-nfinn-light-nni-0710-v04.pdf>

<http://www.ieee802.org/1/files/public/docs2010/new-alon-NNI-Resiliency-LACP-approach-09-10-v03.pdf>

Adjourn at 5:35pm.

Thursday, September 16**AM**

Meeting called to order at 9:00 AM.

Presented patent policy slides. No response to the call for patents.

New Work:

Presentations on Resilient Network Interconnect from Stephen Haddock, Prakash Kashyap, and Janos Farkas:

<http://www.ieee802.org/1/files/public/docs2010/new-haddock-resilient-network-interconnect-LAG-0910-v3.pdf>

<http://www.ieee802.org/1/files/public/docs2010/new-pkashyap-resilient-network-interconnect-LAG-0910-v1.pdf>

<http://www.ieee802.org/1/files/public/docs2010/new-farkas-network-interconnect-functionalities-0910-v01.pdf>

Break for lunch at 12:20am.

PM

Resumed at 1:30pm

New Work:

802.1 received a proposed resolution to our AAP comment on the last call for G.8021v3. Glenn Parsons presented the proposed resolution to call it to the attention of participants. Ben Mack-Crane led a walk-through of a related presentation posted by Maarten Vissers. A 802.1 liaison responding to the proposed comment resolution, and possibly to Maarten's presentation, will be drafted at the November plenary. 802.1 participants are encouraged to provide individual feedback to Maarten.

ITU resolution:

<http://www.ieee802.org/1/files/private/liaisons//liaison-ITUT-G8021v3-LC-comment-resolutions-0910.zip>

Maarten Vissers:

<http://www.ieee802.org/1/files/public/docs2010/new-vissers-elan-etree-multi-vid-architecture-0810-v04.ppt>

Presentations on Equal Cost Multi-Path enhancements for SPB from Peter Ashwood-Smith, Ali Sajassi, and David Allan:

<http://iee802.org/1/files/public/docs2010/aq-ashwood-spb-futures-0910-v1.ppt>

<http://iee802.org/1/files/public/docs2010/new-sajassi-link-utilization-and-convergence-0916-v0.pdf>

<http://www.iee802.org/1/files/public/docs2010/new-allan-load-spreading-for-SPB-0910.pdf>

Adjourn at 3:45pm.

AVB Meeting Minutes

None Provided. For the information currently available, the following agenda was provided in E-Mail to the mailing list:

Monday, September 13

AM

Patent slides shown.

Call for essential patents. There were none.

802.1AS D7.2 comment resolution

PM

802.1AS D7.2 comment resolution continued

Tuesday, September 14

AM

Geoff Garner presented: Multiple Replication Simulation results for 802.1AS Synchronization Transport with Clock Wander Generation and Updated Residence and Pdelay Turnaround Time

- <http://www.iee802.org/1/files/public/docs2010/as-garner-simulation-results-mult-replic-0910.pdf>

PM

Don Pannell presented his discussion of AVB Latency Math

- <http://www.iee802.org/1/files/public/docs2010/ba-pannell-latency-math-0910-v4.pdf>

Michael Johas Teener presented compliance questions for AVB

- <http://www.iee802.org/1/files/public/docs2010/ba-mjt-compliance-questions-0910.pdf>

Wednesday, September 15

AM

New work items were presented and collected by Don Pannell

PM

Comment resolution for

- 802.1BA 2.0
- 802.1AS

End of meeting

DCB Meeting Minutes

Chair's proposed DCB Agenda for the York Interim:

Start at 1:30 PM on Monday, 13 September and run through Thursday, 16 September.

The following time allocation is tentative:

- Monday PM - 802.3bd and 802.1Qaz comment resolution
 - It would be helpful if the disapprove voters on 802.1Qaz could join us for this - if not, please let me know when you can be available and we will work with you then.
- Tuesday through Thursday - EVB projects (the following is tentative depending on Tony's availability)
 - Starting with editor's summaries of the draft status and open issues on 802.1Qbg and 802.1Qbh on Tuesday morning, then we will begin 802.1Qbg review followed by 802.1Qbh.
- Wednesday morning joint DCB and Interworking.

Monday, 13 September

AM

Read the patent policy and made the call for patents – no responses

Created proposed agenda for the week – Meeting will more officially start Monday PM

- <http://www.ieee802.org/1/files/public/docs2010/admin-dcb-agenda-0910-v0.pdf>

Bob Sultan took the extra free time to discuss a proposal for flood reduction

- <http://www.ieee802.org/1/files/public/docs2010/new-sultan-flood-reduction-with-districts-0910-v00.pdf>
- A scheme to reduce flooding of unknown unicast packets by dividing the data center network into 'districts' defined on a per-VLAN basis
- To guarantee you know all addresses within your 'district' you need some kind of address registration and deregistration protocol with appropriate keep-alives
- His concept is that you would create a 'static' type entry with a port-map of NULL. Then you could 'learn' the port assignment at some point later. Perhaps by using a special multicast address that remained within the district
- Managing addresses registration/un-registration would be using a special reserved address
 - how reliable would this need to be?
 - There's an open question as to how reliable things need to be
- What this proposal does is allow you to create smaller address tables within a district, such that you don't have to know all the addresses of people you are communicating with (even outside your district)

- there is no savings for broadcast/multicast
- there is no savings on the district boundary machines as they must learn everything anyway.

PM

Read (again) the patent policy and made the call for patents – no responses

Ballot Resolution for 802.1Qaz

- Paul Bottorff's comments on drop DE bits and 802.1Qaz's lack of addressing what to do with it
 - currently, 802.1Qaz is completely blind to DE bits
 - perhaps Paul's comment is asking for a feature addition to the spec, but it may be that one could view this as a significant missing feature and thus a flaw
 - in Paul's view, to change 802.1Qaz to support this would require a change to the TLV definitions. Some note that the problem is perhaps more of a general 802.1Q problem that scheduling use of DE bits are not specified and thus may need to be a separate project
 - the comment was also made against text that is not part of the re-circulation, so if necessary the comment can be invalidated this time around
 - everyone agreed there is a good point here and a potential larger issue that should be raised (perhaps against Q-REV), but the comment has been withdrawn.
- Other comments from Paul Bottorff against the use of LLDP and described state machines in DCBX are bugs that will be fixed.
- Ben Mack-Crane noted a comment about the portioned allocation of excess bandwidth that is not testable and should not be written as a normative statement. The resolution is to make it non-normative, put it as a note and avoid should/shalls/mays and simply make a statement about an implementation.

802.1Qbh Comment Resolution – Joe Pelissier

- Many people have commented that generating the MCID from a port map is hard to do and would rather fetch the MCID from the MAC/VID look-up. If you want to do ACLs and any other function like port mirroring on the packet and you are getting the MCID from just the MAC/VID look-up, then you would need MCID->MCID translation tables on egress or somewhere after the look-up to get to the correct final MCID.

Tuesday, 14 September

AM

Pat made the call for patents – no responses.

LOAs on 802.1Qbb have been posted to the IEEE 802 website

802.3bd was discussed. One technical comment regarding a missing pad field to make the packet a proper minimum sized packet. Also there were some other editorial comments. Otherwise the document is done.

Tony presented a status update for the 802.1Qbg draft. Getting the state machines agreed to and updated is one of the biggest issues to resolve for the draft. Another big issue is the use of the Ethertype for S-channel tagged frames. CDCP state machines that exist will be used because we

have decided to not to support tie breakers. EDCP is really not needed per-se and will be better described as simply the EVB TLV to be exchanged in LLDP. See:

- <http://www.ieee802.org/1/files/private/bg-drafts/d1/802-1qbg-d1-1.pdf>

ECP state machines were simplified and cleaned-up in this draft. We clearly need to document what type of service the protocol is providing to the higher layer protocols. It will provide an acknowledged service that avoids buffer over runs, only delivers a message once in order, but is not a guaranteed service and assumes the network is ordered. Since we assume that frame delivery is ordered we need to make sure there isn't a network in-between – not a problem. A counter of dropped transmissions will be added.

Rakesh ran through his state machines for VDP

Paul B ran through his state machines for VDP. A bug was noted in all machines regarding a lost Disassociate. The keep-alives will bring back the VSI instance.

PM

Joe reviewed a discussion on a previous ballot comment about how to determine how to control the relationship between PEs and CBs. If you plug two extended bridges together, they might want to reform into one bridge based on bridge priority, etc. However, if you use the mechanism that was decided the day before, you can prevent this by not allowing the extended bridge ports from participating in the CB association process. Joe's proposal is to only allow a single port, by default, as the link to the CB and if you want a more advanced configuration (e.g. like an HA scheme) you must do some manual configuration.

Pat described how link aggregations can be accomplished over S-Channels for some channels and not others and why this might be useful. The key differentiation between VEPAs and PEs in these cases is the fact that PEs don't have a hypervisor present in all cases and thus must have some kind of configuration or policy on what to do in the case of cable reconfigurations

Daya presented her proposal to add Channel Type IDs.

- <http://www.ieee802.org/1/files/public/docs2010/bg-kamath-channel-typeid-0810-v2.pdf>
- This is about having a default profile for the channels. One approach is to provide a channel-ID that is persistent and another idea is to use a default VSI profile for the channel
- It seemed like most people are in favor of simply running VDP and provide a default profile for the channel. The proposal is to add a channel-ID to EDCP that would cause policy to be applied and this effectively results in two ways to have a policy applied

Yizhou presented her proposal for VLAN-ID assignment via VDP.

- <http://www.ieee802.org/1/files/public/docs2010/bg-yizhou-VLAN-info-in-VDP-0910-v1.pdf>
- What she would like to do is allow the bridge to assign the VLAN-ID during the VDP exchange. The VLAN-ID is known by the station via two mechanisms:
 - one is via the VSI type information retrieved during the VM provisioning
 - another is known by the hypervisor as a port group name
 - the latter case could create a configuration problem
- A note that passing information in this way changes the way VDP validates the response to a particular command because you can no longer do a string compare of the TLV
 - you either need transaction IDs or a more complex comparison
- The current text in VDP actually already allows for the VLAN-ID to be set to 0 by the station, but:
 - we lack the text that tells you what to do with the response and

- there is a current issue with validating the response.

Wednesday, 15 September

AM – Joint Interworking/DCB

Joint meeting with Data Center Bridging task group called to order at 9:00am.

Pat Thaler called for patents. No response to the call for patents.

802.1Qbh BPE:

- Joe Pelissier and Paul Bottorff gave presentations on the architectural models of the Port Extender and a Controlling Bridge that supports Port Extension.
 - <http://www.ieee802.org/1/files/public/docs2010/bh-pelissier-genericbridges-0910.pdf>
 - <http://www.ieee802.org/1/files/public/docs2010/bh-bottorff-arch-0810-v1.pdf>
- Joe presented Port Extension support for Generic Bridge Types
 - He is proposing that we use the connection identifier as a way to identify the multi-cast frames that are the same so they can be appropriately filtered down to the primary port.
 - The difference between this proposal and Paul Congdon's previous proposal is that in Paul's proposal, the M-Component was doing everything, including determining the MCID for the frame. What Joe has changed is kept the MCID identification in the C-Component, but used it as a way to do ingress filtering in the M-Component so only a single copy is sent.
 - For the ingress filtering inside the M-Component to work properly, at most one of the ports in the multicast group will be allowed to receive the frame and all others will filter the frame. The port that 'allows' the frame to pass may have other MCIDs associated with it, or it may not.
 - It was noted that each C-component port will need reflective relay enabled in order to allow multicast frames received from a PE to head back to the other ports of the PE
 - Panos is concerned with the passing of the connection identifier between components. Steve agrees we need to be careful on how we are using the connection identifier and making sure we are consistent with our intended use
 - To support generic bridge types there are a few changes required to Qbh. The main change is that any component that wanted to support PE, that component would need to support the ability to generate an MCID from a frame and send it as a connection identifier across the internal LAN
 - To support building an extended provider bridge, we really have two options, use a different Ethertype for the S-Channel, or only use an M-tag in all cases. Joe prefers to use a separate M-tag. This keeps the M-Component from having to understand two different tag formats and only have a single relay that is single tag aware
 - It was pointed out that the size of the MCID needs to be greater than 12-bits in order to support all VLANs, but the source port identifier does not need to be 12-bits and could be more or less as desired. The limit is the number of external bridge ports supported on a single cascade port of the controlling bridge
 - Joe pointed out that multi-channel can still be stacked into the architecture. When the multi-channel is implemented by a port mapping S-VLAN component, it doesn't need to generate an MCID because the default one will work just fine. If you want to generic S-VLAN component that will be doing multicasting like a standard S-Bridge and also want to be an extended bridge, then that S-VLAN component will have to generate a MCID just like the C-VLAN component

- The question of whether it needs to be an interworking project or not was discussed and perhaps DCB and Interworking need to either rejoin or make more time to get together. The editors and chairs should discuss.
- Most people are happy with the concept of a single relay that this proposal provides, but seem to be a bit uncomfortable about the implications of this 'new' type of component.
- Paul Bottorff's presentation on his architectural diagrams for Port Extension
 - Pat pointed out that this presentation should have gone first to provide the background material for Joe's presentation. All this is new material for both DCB and Interworking
 - Paul's diagrams show a way to address option1 of Joe's proposal. If we used a new Ethertype for both the M-tag and the multi-channel, then you could adopt his approach.

802.3bd: Pat Thaler reviewed selected comments from the 802.3bd sponsor ballot.

Joint meeting adjourned at 12:05 PM.

PM

Rakesh presented his latest state machines for VDP

- <http://www.ieee802.org/1/files/public/docs2010/bg-sharma-evb-VDP-Description-0810v05d.pdf>
- fixed the problem with clearing localChange on keep-alive by using an 'if' clause in the associated and pre-associated states
- We converged in general on the design and will attempt to resume in the morning to review the consensus.

Paul Bottorff provided his comments based on the latest draft of state machines that will be uploaded

Donald Eastlake III briefly described an IETF TRILL draft on DCB support

- <http://tools.ietf.org/id/draft-eastlake-trill-rbridge-dcb-00.txt>
- The biggest issue is with Congestion Notification and it is similar to issues that provider bridges also had.

Thursday, 16 September

AM

Pat made the call for patents – no responses.

Paul and Rakesh presented the new state machines and everyone was happy:

- <http://www.ieee802.org/1/files/public/docs2010/bg-bottorff-vdpmch-0910-v1.vsd>

Rakesh presented a new discussion on Link Aggregations and EVB.

- <http://www.ieee802.org/1/files/public/docs2010/bg-sharma-evb-lag-v01.ppt>
- The question of where does LLDP and ECP run comes up again. Certainly we will have an instance on the physical link, but if an instance runs on the aggregation, we have the problem of disambiguating the instance again
- Also the question exists on performing a link aggregation of channels as well. We also looked at lots of scenarios of lagging channels and links using different addresses for

- LACP (nearest bridge and nearest customer bridge). It gets very complex quickly, but in many scenarios work architecturally.
- We need a way to know which instance of the LLDP and ECP are the aggregate and which are the physical link.
 - A proposal for disambiguating LLDP instances over LAGs has been presented before and can be seen at:
<http://www.ieee802.org/1/files/public/docs2010/new-congdon-linkag-LLDP-0110.pdf>
 - The EVB use of LLDP (CDCP and the EVB TLV) will always run at the aggregate level and will be using nearest customer bridge. The other LLDP agents running on the physical links will be using nearest bridge address anyway.

To close on the channel ID proposal, we will create a new address format for the VSI TLV that is VLAN only and we will specify that 0xfff is a wildcard that means all VLANs, effectively, all traffic will match the VSI type – thus creating a default VSI for the channel.