TRILL WG

Date: Friday, March 7, 2014, 11:50 – 13:20
Location: Blenheim Room

Jabber scribe: AD Ted Lemon
Scribe: Sue Hares
Co-chair: Jon Hudson
Co-Chair: Donald Eastlake

Agenda:

Administrativia Chairs [4 minutes]
Scribes selection
Agenda Bashing:
Document Status & milestone report Chairs [5 minutes]

TRILL Implementation Report Susan Hares [5 minutes]
TRILL OAM Status Update, Tissa Senevirathne [10 minutes]
co:authors: Tal Mizrahi, Deepak Kumar
draft-ietf-trill-oam-fm-02
draft-ietf-trill-loss-delay-02
draft-ietf-trill-oam-mib-00

TRILL Active-Active: YiZhou Li, Weiguo Hao, Mingui Zhang [24 minutes]
draft-yizhou-trill-active-active-connection-prob-02
draft-hao-trill-analysis-active-active-01
draft-zhang-trill-zz-multi-attach-00

Trill over IP Margaret Wasserman [10 minutes]
draft-mrw-trill-over-ip-04

Directory Assisted Edge Linda Dunbar, Donald Eastlake [8 minutes]
draft-ietf-trill-directory-assist-mechanisms-00
draft-ietf-trill-ia-appsubtlv-00
draft-dunbar-trill-directory-assisted-encap-04
(draft-ietf-trill-channel-tunnel-00)

Wrap-Up, Chairs [4 minutes]

[1] Adminisitrivia (scribes etc), Agenda Bashing, Chairs [4 minutes]

- The reading of the Note Well. IPR should be disclosed
- No changes were made to the agenda

[2] Document Status, Chairs [Donald Eastlake]
[slides-89-trill-0.pptx] [10:00 – 18:00 on audio recording]

- TRILL Standard Track RFCs
  - RFC 6325, ñRBridges: Base Protocol Specificationò
  - RFC 6326, ñTRILL Use of IS-ISò ï published as updates
  - RFC 6327, ñRBridges: Adjacencyò ï published as update
  - RFC 6361, ñPPP TRILL Protocol Control Protocolò
  - RFC 6439, ñRBridges: Appointed Forwarders
  - RFC 6850, ñDefinitions of Managed Objects for RBridges
- TRILL informational RFC
  - RFC 5556, ñTRILL: Problem and Applicability Statementò
  - RFC 6847, ñFCoE over TRILLò
  - RFC 6905, ñRequirements for OAM in TRILLò
  - RFC 7067, ñDirectory Assistance Problem and
- In RFC Editorôs Queue
  - draft-ietf-isis-rfc6326bis-03.txt
  - draft-ietf-trill-clear-correct-06.txt
  - draft-ietf-trill-fine-labeling-07.txt
  - draft-ietf-trill-o-pw-06.txt
  - draft-ietf-trill-oam-framework-04.txt
  - draft-ietf-trill-rbridge-bfd-07.txt
  - draft-ietf-trill-rbridge-channel-08.txt
  - draft-ietf-trill-rbridge-extension-05.txt
  - draft-ietf-trill-rfc6327bis-02.txt
- Drafts in TRILL WG Last Call
  - draft-ietf-trill-esadi-05.txt
- Other Drafts in TRILL WG
  - draft-ietf-trill-channel-tunnel-00.txt
  - draft-ietf-trill-cmt-02.txt
  - draft-ietf-trill-ia-appsubtlv-00.txt
  - draft-ietf-trill-loss-delay-02.txt
  - draft-ietf-trill-oam-fm-02.txt
  - draft-ietf-trill-oam-mib-00.txt

    - draft-ietf-trill-rbridge-vlan-mapping-10.txt
    - draft-ietf-trill-resilient-trees-00.txt
- Other TRILL Related Drafts in Other WGs
    - draft-ietf-l2vpn-trill-evpn-01
- Draft in Call for TRILL WG Adoption
    - draft-mrw-trill-over-ip-04
- **Discussion:**
    - **[Donald show a graph of draft normative dependencies.]**
    - Comment: This dependency graph looks like a state machine.
    - Question: How do you read the dependency diagram?
    - Donald: The drafts are mostly dependent on the TRILL use of IS-IS (rfc6326bis draft at the bottom). This will unblock a great number of the TRILL draft that wil probably have sequential RFC numbers. They may come out next month.
    - Donald: Please read the drafts that are upcoming.
- Milestones:
    - We have only 2 milestones overdue
        - Jan 2014 Initial WG draft on TRILL over IP ï draft is progress
        - Feb 2014 Submit OAM Fault Management to IESG as Proposed Standard ï draft in progress
    - **Upcoming milestones**
        - Mar 2014 Submit ARP/ND Optimization to IESG as Proposed Standard
        - Mar 2014 Initial WG draft on TRILL Implementation
        - April 2014 Submit Active-Active to IESG for Proposed Standard
        - April 2014 Submit OAM Performance Management to IESG as Proposed Standard
        - Jun 2014 Submit TRILL-over-IP to IESG for Proposed Standard
        - Jul 2014 Submit Multilevel to IESG for Proposed Standard
        - Aug 2014 Submit MultiTopology to IESG as Proposed Standard
        - Nov 2014 Initial WG draft on TRILL IS-IS Security
        - Mar 2015 Submit RBridge support of DCB (PFC, ETS, CN) to IESG for publication as Proposed Standard
        - Dec 2015 Re-charter or shut down the WG
    - Please take a look at what is on the milestones. We can do other things not listed as one of the milestones as long as they are in our charter.

[3]TRILL implementation Report, Susan Hares [19:00 ï 24:00]

- Demonstration of online TRILL Drafts
    - Live demo for the TRILL work
    - https://www.surveymonkey.com/s/HJFLX3N
    - Sue: Volunteers ï I would like people

- Discussion:
    - Sue: I would like people to help me review the drafts and go through and find the options.
    - [**Erik Nordmark (?)]:** This is going through the RFCs indicating I did implement this work or I did not implement.
    - Sue: This online form is a demonstration for getting the information. What I need right now is people who will read the drafts and send me a copy of the survey.
    - Jon Hudson: Look at this as an opportunity. Going through this in detail with someone that has as much opportunity as Sue will be helpful. You will understand how this works and the details within the draft. If you do not have the experience, take this as an opportunity.

[4] TRILL OAM Status Update: Tissa Senevirathne

[slides-89-trill-1.pptx] co-authors: Tal Mizrahi, Deepak Kumar [audio: 27:00 - ~34:00]

- Status:
    - Formal presentation done to the IEEE 802.1 Interim committee meeting requesting formal allocation of block of CFM opcodes code points and block of CFM TLV types for the purpose of IETF OAM - Formal ballot is in progress
    - Published RFC 6905 ï Requirements for OAM
    - In RFC Editors queue: rdraft-ietf-trill-oam-framework-04
    - Updates published to:
        - draft-ietf-trill-oam-fm-02
        - draft-ietf-trill-loss-delay-02
    - Still working on:
        - draft-ietf-trill-oam-mib-00

- Status of CFM code points [Donald Eastlake]
  Slides: [draft-89-trill-7.ppt]
    - Quick update on the code points based on IEEE 802.1 liaison received yesterday 3/6/14
        - The IEEE 802.1 working group informing the IETF, based on the request in the IETF liaison (9/24/13), had voted to allocate these code points to the IETF.
        - draft-eastlake-iana-cfm-consideration will be an AD sponsored draft that will create an IANA registry for these code points. This will draft will allow code points to be assigned based on IETF standards action.
        - Please let me know you know if you have a question.
        - this was the gating factor for getting the following drafts WG LC
            - draft-ietf-trill-oam-fm-02
            - draft-ietf-trill-loss-delay-02

        - Both drafts will be WG LC from 3/7 to 3/24.
        - Any comment on the list.
- Discussion:
    - [Unknown] What does the draft with these registries do?
    - Donald: The draft (draft-eastlake-iana-cfm-consideration) simply creates the registries and puts the values in the registries. It states the assignment of values is based on a standard action. The Fault management draft (draft-ietf-trill-oam-02) makes requests off that registry.
    - Tissa (?): When will these drafts go to WG LC?
    - Donald: The two OAM drafts (draft-ietf-trill-oam-fm-02, draft-ietf-trill-loss-delay-02) will go through a 2 week WG LC, a bit of processing (at IETF AD), and the 2 week IETF LC. The AD sponsored draft(draft-eastlake-iana-cfm-consideration) will go through a IETF LC, but it is a 4 week LC because it is not WG sponsored.
    - Tissa: How soon will this AD Draft (draft-eastlake-iana-cfm-consideration) go to the IETF LC?
    - Donald: It should be within a week. The processing time will be about the same for IETF AD process for WG LC and IETF AD processing for AD sponsored draft.

- **Discussion on Tissa's slides**
    - Weiguo Hao [Huawei]: Should the OAM be combined with the netconf-yang model?
    - Tissa (Cisco) : These drafts should be separate.
    - Weiguo Hao [Huawei]: Thank you

[6] TRILL Active-Active: Yizhou Li, Weiguo Hao, Mingui Zhang

- **draft-yizhou-trill-active-active-connection-prob-02 [slides-89-trill-2.pptx] [36:00-38:16]**
    - **Presentation**
        - update 01 to 02: mostly editorial
        - update 00 to 01:
            - Some common understanding for MC-LAG as it is vendor specific
            - Remove RPF failure from problems sub-section. Re-phrase it as goal
            - Add high level goals, and Co-authors
    - **Discussion:**
        - Donald: Are there any comments?
        - Jon: Please review this draft and comment. We need to make sure to do this correctly.

        - Donald: This has been adopted on the list and hearing no objection here, it has been adopted by the WG as
            - ñdraft-ietf-trill-active-active-connection-prob-00.txt

- **draft-hao-trill-analysis-active-active-01 –(Weiguo Hauo)** [**audio: 39:00 – 1:05:00]**
    - Presentation : slides-89-trill-3.ppt
    - Discussion
        - [**Problem 1: slide 3]**
        - **Erik Nordmark**: If there is another MC LAG involving the same set of three RBRIDGES, in one of these solutions (consistent or independent) will you allocate a different name for each different lag connected to both?
        - **Weiguo:** Yes
        - **Erik Nordmark**: Is this true for both the solutions or just the consistent solution? Even in a consistent allocation, if it is CE-2 that is sending something on MC-LAG-2 that packet come to RB3, it needs to get forwarded out to RB1, assuming that RB will need a nickname?
        - **Jon:** You need a unique identifier in either solution.
        - **Erik:** Is this true for both the solutions?
        - **Weiguo:** yes, for consistent allocation the nickname is allocated per global allocation.
        - **Erik**: Per LAG, not per connection of RBridges associated with a MLAG. Is this true for the consistent as well as the independent?
        - **Jon:** It should be true for both (independent and consistent)?
        - **Erik:** I do not understand the nickname consumption difference?
        - **Weiguo:** Perhaps we can discuss this on the mailing list.
        - [Slides 4-
        - **Frank Xia (Huawei):** This is a comment. If you choose any one solution, does one solution have a problem co-existing with the other solution? I think you study this part of the problem.
        - **Weiguo:** This is a good comment. Each Solution is independent. It can be stitched together with each other.
        - **Jon:** In theory.
        - **Weiguo:** Any solution can be stitched together.
        - **Tissa:** I do not understand your last slide. Frank brought up an important point. At the end of the day, for TRILL to be successful, one active-active solution needs to be chosen. We have been talking a lot about it, and we have not accomplished anything. My proposal is to pick one, and go with it. The lack of a single solution is preventing deployment in the real world. Trying to pick all solutions will end up with a solution where interoperability will be broken. Please pick one and go with it.

- Yi**zhou Li (Huawei):** There is one document that is a base document for this investigation. It is a base document that should start with the Psuedo-name (draft-hu-trill-pseudonode-nickname-06). I think this with the pseudonode nick that will provide all of the solutions.
- **Jon:** We should definitely discuss this on the list.
- **Donald:** Do we want to adopt the solutions draft at this time?
- **Tissa:** I think we should have additional changes before adoption. There are flaws that need to be addressed on the list. Please have one more review on the list.
- **Donald:** You want to see it improved.
- **Tissa:** Please see the centralized replication part which the list has indicated.
- **Weiguo:** The centralized he short comings have been updated. I will be release this later.

- **draft-zhang-trill-zz-multi-attach-00**
  [draft-slides-trill-4.ppt [Audio: 1:05 – 1:17]
  - **Discussion:**
    - Tissa: I do not understand MAC information with a Data label ï with the claim that it does not need hardware. Can you please explain this point?
      Mingui: The draft is now the starting point with suggestions asking for discussion and help, rather than a detail solution.

    - Tissa: If you are not addressing this in the draft, I do not understand how this works. Can you please help me understand this? Mingui: Basically, the data label has to be configured per MC-LAG in an Active/Active Edge group, so that it can be used to distinguish MC-LAGs.

    - Tissa: What is the definition of your data label?
    - Mingui: The data label is the Fined Grained Label/VLAN.

    - [Tissa] In that case I do not understand how you are claiming the solution is hardware independent

    - [John] Yes sounds like pre-mature and need more investigation. Heôs looking for help on the solution, and he is looking for input.

    - Tissa: This is why he is providing no discussion on the hardware. Minguiôs cliam to have no hardware dependency ï is not precisely correct. It is premature to make that claim until you have all the details completed.

- Erik Nordmark: I would like to understand your assumptions. At first I thought you had different VLANs which are carried across the TRILL Fabric [in control packets]. For example, packets on VLAN 10 will be transmitted to a VLAN Label 20 will be separated from VLAN 20. These VLANs will never meet in TRILL. You do not have any looping. VLAN 10 will always be separated from VLAN 20 by RB3. This is a completely local matter. If I already have VLANs, I can use an active-active methodology, and it will split the loads.
- Jon: This works even in a case where the interface has 6 VLANs on the interface.
- Erik: You do not need to change anything in TRILL to do active-active in this fashion. Both Rbridges are denoted as designated forwarder, but you have to have consistent configuration to make this work. Are you intending VLANs to handle this balancing across the network, or are you intending VLANs to handle things at the Edge. These are two completely different problems. This needs to be defined carefully.
- Weiguo: How does the root RBridge learn that a C-MAC is associated with multiple Ingress Rbridges ? Is it through the control plane or through the data plane?
- **Mingui:** It depends. If the remote RBridge chooses to store multiple locations at the ESADI-LSDB while install only one location into the data plane.Then it doesn't require new silicon. If the remote Rbridge chooses to
- store the multiple locations at the data plane, it needs new hardware.
- Weiguo: A classical LAN switch hardware can only learn one C-MAC [?] with one ingress Nickname with one ingress port.
- Donald: This requires silicon changes. Due to time, we must move on to the next presentation.

- **TRILL Over IP, Margaret Wasserman [10 minutes] [draft-ietf-trill-over-ip.txt] [slides-89-trill-8.ppt] [Audio:1:18 – 1:30]**
  - draft-mrw-trill-over-ip-04.txt
  - presentation
    - Document defines a UDP/IP encapsulation for TRILL
    - Connects two remote TRILL sites into a single TRILL campus over any IP network
    - Two Scenarios:
      - Remote office scenario
      - IP Backbone scenario
    - **Document status:**
      - 04 version published in January 2014,

        - Changes were additions of UDP header, congestion consideration section, updated references and editorial changes.
    - **TRILL has no useful source port to hold so this has issues in the header**.
        - The 5 tuple used for load sharing (ip source address, ip destination address, ip source udp port, ip destination upd port, and the IP protocol) ï this could be the same for all traffic on the link.
        - We set the source port set to an entropy label that is based on the 5 tupel or the like to allow ECMP to work. This allows the flows to hash and load share.
    - **Open issues:**
        - Only one issue: Should we do a UDP/IP encapsulation or a custom header?
            - The alternative is to have TRILL run directly on IP. Direct header might require TRILL specific header for entropy, and a checksum is the only field that are necessary.
            - We provide entropy in to checksum?
            - Is checksum of the TRILL packet needed or desirable?
        - We need to complete the section on handling multicast, and middle box commands.

- **Discussion**:
    - TRILL: The gain is the small header size with UDP.
    - Margaret: The UDP header is 8 bytes. Our only saving is 6 bytes of saving.
    - Tissa: It will be much simpler. UDP is deployed and worked out. It is a better solution.
    - Margaret: We have not written a middle box solution section. Sometime middle boxes block everything that is not UDP or TCP.
    - Tissa: This another reason to go with UDP.
    - Margaret: Does anyone else have a concern?
    - Donald: You have additional slides.
    - Margaret: One more. However, how many people have read this document? If you could ask for an opinion?
    - Donald: Do you think UDP encapsulation? There are about 10 people who have read the draft. How many people think UDP encapsulation is better?
        - There is zero for special. It is strongly in favor of UDP encapsulation (>10). Consensus for UDP that will be confirmed on the mailing list.

    [second discussion]

- Stephen osage (UPNC): Would control plane be another issue? How do get IP address of entry point of the UDP tunnels?
- Margaret: We have defined this as static. We have not considered dynamic consideration.
- Jon: I am seeing more interest among our customers in Just-in-time tunnels.
- Margaret: You can start with this static approach and then layer this on top of it a control protocol.
- Stephen: This is for machine migration across 3 campuses a use case you are interested in? Are you able to scale using UDP/IP encapsulation?
- Margaret: Are you looking for data centers?
- Stephen: Is the destination you encapsulate to something that could change?
- Margaret: In data centers where machines are moved? I have not considered it. I do not object to having a discussion.
- Jon: You can use fine-grain labeling to create replicated space. The tunnels should not impact the stack of the VM. You might get a few dropped packets.
- Donald: The IP tunnel is a link. TRILL does not have to be aware of what the link type is (IP or something else).
- Margaret: Is your question would I accept a control protocol written that brings up this links dynamically? If so, the answer is yes. If your question is, does this basic protocol need a dynamic protocol ï then the answer is no.
- Stephano (lucent): In France, we are seeing a number of Data Center providers who have an open source implementation of TRILL at different sites, and interconnected via IP. They are consider LISP as a control plane protocol to handle move machines dynamically. It is this use case that causes a question.
- Weigo: I have comments about multicast issue. The multicast issue is a most difficult and complicated problem. For the TRILL multicast issue, there is more effort that is needed to handle this part in the draft.
    - Please consider how the IP network multicast solution needs to work with TRILL multicast solution? You need more details here.
- Margaret: There are two cases: a) IP network supports multicast, and b) IP network doesnôt. The second case requires serial unicast.

Chair: Unfortuntely, we are out of time to talk about the following at this meeting:

1. The Directory Assist Status ï
2. Directory assisted TRILL Encapsulation

Chairs: Thanks for being here. See you on the mailing list and in Toronto.