IETF NVO3 WG
Virtual Interim Meeting Agenda
2015-01-28 10:00-11:15 EST

Chairs: Benson Schliesser, Matthew Bocci
Secretary: Sam Aldrin

Note Takers: Ignas and Sam

NOTE: Time-slots include both presentation AND discussion.
Please plan accordingly.

1. Meeting Introduction and Agenda Bashing (10 min)

- Future interim meetings were scheduled. Could change a little. Next one is on 2/12.
- Encouraged to read Note well
- General topics to the end of the meeting
- Benson: starting, still waiting for Matthew to join.
- Experimental blue sheet via etherpad
- Link in chat window
- Agenda: as listed, plus a discussion on a control plane
- Recording the meeting now.

- Benson: Announced interim dates before Dallas might change, announcement will be sent in advance.
- For today: agenda
  - CP requirements, would like to have a discussion
- And a general topic at the end on the control plane aspects.

- Silence on bashing the agenda.

2. Update on Control Plane Requirements (10 min)
Chairs

- Requirement documents at the same time as solution documents, contrary to the existing milestones
- Authors do take a note on that, where solutions are buildable as well.

Larry: when we talk about solutions, what about the VDP cjhanges needed? Given that it is not an IETF protocol.
Benson: guess I do not know the answer instantly myself We could want to liason to IEEE. NVO3 may want a draft on how to use the protocol. If there is a need for extensions, there may be more things to be done.
Alia: suggestion. We can liase with IEEE, writing an applicability draft may be the start. But if we asking for changes to VDP then we have to liase formally.

Larry: if I understood correctly - generic requirements would be covered by this?
Alia: you need requirements for NVO3 that covers VDP applicability.
Alia: I would combine applicability and requirements into a single draft covering VDP
Benson: my view on having multiple req docs is when there are many solutions. In case of VDP it is simpler.
Benson: please bring to mail list this topic.
Benson: similar approach may be in the data plane too.
Larry: I am confused now. Will there be a single document only?
Benson: in case of VDP we would put everything into one document.

3. User Plane Signalling (10 min)
https://tools.ietf.org/html/draft-zu-nvo3-user-plane-signalling-01
Zu Qiang zu.qiang@ericsson.com

- This draft is to discuss how to handle tenant traffic. Not sure where it falls, data or control plane.
- Data forwarding is based on inner-outer addr table.
- No additional handling is required.
- L2 services provided by NVE, from tenant perspective NVE is just a bridge.
- Drafts includes STP. Future versions to include other protocols as well.
- NVE need to learn root bridge.

Larry> LACP is aggregating physical links. How you would stretch physical links.
Zu> Link aggr is between tenant systems
Lawrence> What is to be aggregated, as they are physical links?
Zu> Didn't understand the question. Will take it offline.

- There is no need for NVE to support handling of ARP.
- When NVE receives unknown address, NVE may query other NVE, creating security issues.
- In multi homing scenario, multiple NVE's could respond.
- In routing, a virtual routing instance will take care of routing function in a tenant system.
- L3 routing is handled by centralized GW function.

Tom> Why do we need GW here?
Zu> Use case is centralized GW is what I am discussed here. The other use case is distributed GW.
Tom> Are we talking about triangular routing
Zu> It is triangular routing in this case.
Benson> We might use this centralized GW for scale.
Anoop> This is not talking about one or the other, right?
Zu> Only talks about centralized GW function.

- In Distributed GW function, VR is located to Tenant system and attached to distributed GW's.
- Updates routing polices using routing protocols configured.
- Question is how to update remote NVE?
- To update remote NVE, I do have few alternatives.
- 1. Disallow dynamic routing
- 2. Using NVE-NVE interaction
- 3. Use NVA-NVE signaling
- 4. Collocate NVA and GW function.

Benson> As we ran out of time, please take it to the list. If time permits at the end, we could discuss then.

4. CP issues of Layer 2 Gateways (10 min)
https://tools.ietf.org/html/draft-xia-nvo3-l2gw-02
Anoop Ghanwani anoop@alumni.duke.edu

- This is about L2 GW. Presenting for the second time as the draft was significantly updated.
- L2 Overlay networks were deployed in DC''s but traditional L2 Bridging are still being used.
- L2GW is basically a NVE
- There are physical wires connecting network or TS to NVE, causing some issues.
- First issue is MAC learning
- Learn addresses at local L2GW and exchange with remote L2GW
- In multihoming, active standby can be handled by loop detection protocol.
- Active active is harder. More protocol work is required.

Linda> Is it L2 over L2?
Anoop> It is L2 over L3. Underlay is L3.

- In summary, L2GW is a physical NVE.
- Major areas to be addressed as loop detection and active-active connectivity to L2GW.

Benson> Running over time. Please hold comments to the end or take them up on mailing list.

5. CP issues of Tenant System Address Migration (10 min)
https://tools.ietf.org/html/draft-merged-nvo3-ts-address-migration-01
Linda Dunbar linda.dunbar@huawei.com

- We presented this last time. Got adopted 3 years ago.
- Removed solution portion as it was discussed at a different WG
- Added solution with anycast for TS
- NVA manages all unused VLAN-IS's pool
- NVA to manage the first switch attached to TS
- Dynamically interconnect NVE's.
- Various solutions were presented for outbound and inbound traffic.

- Next step, Draft is ready for WG adoption.

Benson> Does this documents specifies how protocol should be implemented or applied?
Linda> Both.

6. NVA Address Mapping Distribution (NAMD) Protocol (10 min)
https://tools.ietf.org/html/draft-dunbar-nvo3-nva-mapping-distribution-00
Linda Dunbar linda.dunbar@huawei.com

- We added 3 author, Tom Herbert, to the draft
- TS could be attached to underlay and overlay network
- First type is, centralized NVAwhere it has all info. Second one is distributed NVA.
- Suggesting to use reachability TLV
- This is to advertise set of addresses in a VN.
- In a push service, uses VN scoped instances of the IS-IS to announce all VN's.
- For incremental push, a new TLV is need to carry timeout and flag for NVA for indication.
- Proposed a Pull query format.

Tom> What is the purpose of sequence number in the query
Linda> If this query is sent by multiple fragments
Tom> Does seq number has state associated with it?

7. Open Discussion (15 min)
Benson> Time to ask question. Silence if fine.
Benson> Next interim is on 2/12 on multicast topic.
Eric> Where is recording available?
Benson> Will let you all know.

See you all in the next interim meeting scheduled for Feb 12th 2015