

NTAC Peering and Routing Working Group meeting notes

7/23/2019 4p EDT

1. Agenda Bash
 - a. There was no bashing of the agenda.
2. Update on peering and TR-CPS/I2
 - a. Jeff Bartig: Not a lot of updates largely due to recent travel. Working on wrapping up 100G peer upgrades. Number of google 100G peering connections.
 - b. Bill Owens: recent email exchange with Google, asked about IRR filtering and they indicated they had no idea what was going on. Jeff Bartig: they just point to FAQ at this point. Guessing that they will delay past September.
 - c. Tony Brock: will Internet2 be doing outreach to get folks registered with IRR? Jeff Bartig: Yes, messages are being drafted to CIOs and other administrative folks this week. Steve Wallace: on Friday's RPKI webinar, will also discuss in general terms. Another webinar on 8/21 from MERIT on how to use IRR and RADB.
 - d. Jeff Bartig: don't know specific google requirements, assuming that best practices will be sufficient. I2 will need to build AS-SET. Still determining process to do this.
 - i. HE is still filtering based on AS-SET. Putting out a partial AS-SET would be an improvement.
 - e. Matt Valenzisi: MCNC is almost ready to publish their AS-SET. What is the approach? Jeff Bartig: publish anytime you are ready. Could reference in peeringDB. Matt: intended to get objects in order via ARIN.
3. Network Weather Update - Trends and Coming Events
 - a. It's summertime; traffic is low. Fall semester will likely be the next shift in backbone traffic.
4. I2 Network update
 - a. New equipment being installed for MAN LAN and WIX. International R&E peering focused.
 - b. Vacating cabinet in Chicago.
 - c. 100G connections to new routers.
5. I2PX Community Stripping/AS-Path Maximum
 - a. Have had issue with iBGP flaps on I2 backbone. Has impacted Seattle, LA, one router in Ashburn. Would occur in differing frequencies. Origin was an I2PX peer, shut down the peer to stabilize. Was able to capture specific routes that were causing the problem.
 - b. These were normal routes but had a lot of BGP communities and pre-pended ASes, which triggered a bug in JUNOS related to a specific feature. If attribute-list is close to 255, a corrupt iBGP update is generated, causing a flap.
 - c. Current recommendation from Juniper is to filter routes to prevent the bug from being triggered.
 - d. Waiting to hear from Juniper as to how the bug will be fixed.

- e. Looking for feedback from community as to how detrimental this workaround would be until the bug is fixed. What should we be doing with what are essentially unnecessary communities?
 - i. Brad Fleming: should be OK to strip off communities that are not transitive. Come up with reasonable AS-path limit. Maybe 30?
 - ii. Tony Brock: is anyone aware of communities that NEED to be transitive?
 - iii. Richard Hicks: could pass on what you're receiving, or you could strip and re-tag that you learned it from a specific ISP. Is there an easy way to have a set of communities that you would allow?
 - iv. Noted that these communities are not inherently bad.
 - f. Jeff Bartig: current plan is facing commercial peers, strip all communities. Other peers, drop routes with prepends >25, we get the idea that they don't want to use that path.
 - i. Positive feedback to this approach.
 - g. Tony Brock: General agreement that I2 should do what needs to be done to keep network stable.
 - h. Bill Owens: noted that this does not affect the R&E table.
 - i. Jeff Bartig: longest path today is 12 hops. So, limit would be double this.
 - j. Conversation will continue on the mailing list.
6. Segment routing working group
- a. Topic came up in MN workshop.
 - b. Suggested that should be a separate working group. General feedback?
 - c. Michael Lambert: what is purpose of working group?
 - d. Dale Finkelson: people wanted to know how it would impact what they were doing and how would it function in our infrastructure.
 - e. Bill Owens: maybe should be more of an interest group or fan club.
 - i. General agreement.
 - ii. Tony Brock: want to make sure that local decisions mesh with community decisions/directions.
 - iii. Michael Lambert: as an example I2 had an SDN wg while the backbone was running SDN but not MPLS.
 - iv. Steve Wallace: Given that this is more like SDN, it might be better to come out of the broader NTAC.
 - f. Tony: let's send back to the NTAC. General agreement.
 - g. Drop a note to Dale Finkelson if interested in participating.
7. IRR filtering
- a. Intended to be topic that Jeff already discussed.
 - b. Steve Wallace: upcoming webinar on RPKI. Some RONS will be updating IRRs on behalf of their members. Would like to know what preference is from RONS; expect survey soon from Linda Roos.
8. Comments/feedback from last week's working group
- a. Matt Valenzisi: glad he went; good info and discussion. Some priorities that need to be sorted out. Need to be sure that we don't put 90% effort into 10% of problems.

- b. Bill Owens: continue to be skeptical of dynamic stitching networks/federations or testbeds. Pushback from some. Campuses and regionals need to speak up with opinions.
 - c. Michael Lambert: hard to push orchestration unless there is downstream demand.
 - d. Jeff Bartig: talked about a lot of technologies, but didn't talk about how they would all fit together and what expectations from community would be. Interpreted Farmer's comments to mean folks needed automated connectivity to Cloud.
 - e. Matt Valenzisi: EVPN can be very helpful as compared to VPLS, as can segment routing. Not so much concerned about interop between networks, but about how to make local networks happier.
 - f. Tony: appreciated that I2 held the workshop and was genuinely seeking feedback and was open to changing course if needed. Glad to see that L1-L3 design was being genuinely re-thought. Looking forward to seeing outcomes.
 - g. Eric Boyd and Joe Breen are convening a workshop on telemetry and automation. Contact either of them or Dale if interested.
 - h. Tony Brock: has an optical vendor been selected?
 - i. Dale Finkelson : Can't say.
9. AOB
- a. None.
10. Adjourn at 5:03pm