



Joint Engineering Team (JET) Meeting Minutes

National Coordination Office for Networking and Information Technology R&D (NCO/NITRD)
490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024
January 18, 2022, 12:00-2:00 p.m. ET
This meeting was held virtually

Participants

Joe Breen, UTEN/University of Utah	Linden Mercer, NRL
Nick Buraglio, ESnet	Doug Montgomery, NIST
Rich Carlson, DOE/SC	Edward Moynihan, Indiana University
Bill Fink, NASA/GSFC	Aruna Muppalla, NASA/GSFC
Andrew Gallo, CAAREN/GWU	Mark Mutz, NOAA
Ann Keane, NOAA	Glenn Ricart, US Ignite
Michael Lambert, PSC/3ROX	Frank Seesink, University of North Carolina
Caren Litvanyi, Indiana University	Kevin Thompson, NSF
Paul Love, NCO/NITRD	Steve Wallace, Internet2
Joe Mambretti, StarLight/MREN	

Proceeding: This meeting was chaired by Kevin Thompson (NSF) and Rich Carlson (DOE/SC).

I. Action Items:

- Internet2 and ESnet updates on their respective new networks.

II. Review of the Minutes of the November 2021 meeting: Three corrections were received.

III. SRv6 Fundamentals and update – Nick Buraglio

- A. SRv6 is a technology with similar structure and attributes to SR-MPLS:
 - a. In SR-MPLS Labels are distributed in order to perform tasks such as traffic engineering, fast reroute, etc.
 - b. SR-MPLS requires a consistent (i.e.: “global”) label space to be agreed upon and configured across a given network. The labels are distributed via an internal protocol such as OSPF or IS-IS.
 - c. SRv6 further simplifies this by removing labels and leveraging full, 128 bit IPv6 addresses as the reference pointers for these tasks. Distributed in the same way as SR-MPLS.
 - d. SRv6 uses a block of IPv6 space within a network’s assigned IPv6 space as a “locator block” much like the global label space used in SR-MPLS.

- B. Previous implementations didn't go very far
 - a. They required very large header to house the 128bit segment identifiers which in turn made the headers too large for routing to be supported in hardware. Therefore routing fell back to software support.
 - b. Several competing, vendor specific implementations with poor vendor support due to the header requirements where supporting the 128 SIDs would have required hardware updates.
- C. SRv6 today:
 - a. SRv6 leverages the Segment Routing Header (SRH) to encode the packet processing program in the IPv6 packet headers as a Segment List, together with optional metadata.
 - b. The packet processing instructions may express:
 - i. TE
 - ii. Topology
 - iii. VPN
 - iv. Fast Reroute.
 - c. General limit is 15 SIDs, which still requires $N*16$ bytes to be carried inside the SRH.
 - d. But rather than being an IPv6 address, a current implementation of SRv6 now uses microSIDs (uSID). This implementation currently has a lot of traction
 - e. Still 16 bytes long, but uSIDs provides a model for reducing the overhead by introducing a 16 byte extension containing a "micro-instruction program" with several "steps" encoded in each single 16 bytes to direct the processing of the packet.
 - f. Advantages of uSIDs:
 - i. Significantly simplified routing stack.
 - ii. No longer disseminating labels.
 - iii. Like SR-MPLS, complex tasks can be managed via a controller.
 - iv. Leverages existing protocol stack.
 - v. uSIDs are gaining increasing support from large vendors and host stacks.
 - g. Security implications and disadvantages:
 - i. Uses Routing Header 4 which many platforms do not yet have the capability to filter on.
 - ii. No current BCP for filtering at network edges or for securing SRv6. One is needed and several people are talking about writing one.
 - iii. Not many production implementations.
 - iv. Support missing from some vendors.
 - h. References:
 - i. <https://www.segment-routing.net/scientific-papers>
 - ii. <https://dl.ifip.org/db/conf/cnsm/cnsm2020/1570663490.pdf>
 - i. Where is ESnet on using SRv6?
 - i. Exercised in the lab. Currently not planning to move ESnet6 from SR-MPLS to SRv6. Using SR-MPLS is a solid plan. Perhaps in the future ESnet will revisit SRv6.

IV. Discussion of the JET's tasking on tools to help with inter-domain issues – Joe Breen

This is a community project to collect shared data from all who are willing to share.

The related, live map is at: <https://www.globalresearchmap.org/>

- A. Dan Doyle from the GlobalNOC with his team and the team from LEARN have finished getting all of LEARN's circuits live in the project.
- B. Next RON to be focused on is MERIT. MERIT has an internal project but when that is completed the goal is add its circuits into the project.
- C. Dialog is continuing with RNP to add its circuits shown in South America.
- D. Background on efforts lead by Eric Boyd, Joe Breen, James Deaton, Dan Doyle, and Karl Newell:
 - a. The project gets basic SNMP metrics from groups around the country that are willing to share for trouble shooting and research. Metrics include link utilization, discards and errors. These are collected hop by hop as the path crosses multiple domains.
 - b. Several prototypes are going along with the drafting of a basic letter of intent for those wishing to participate.
 - c. Tools: Telegraf container as an option for local collection. Nearly ready for production use.
 - d. Tracking sheet of networks willing to share data. Please update your network's entry. See:
https://docs.google.com/spreadsheets/d/1pMW_PNVpeT42nAxa3bW4QostMxcHTXkWSPbZOpIFwE/edit#gid=0
Templates for campus, regional and national networks setting out what data is desired can be found at:
Campus template: (for a Science DMZ or research segment)
https://docs.google.com/spreadsheets/d/1v7iFw8_YoMpa3wjgwcmlZgy0QsTi1bHb4Qk1cV6qfAM/edit#gid=1161461998
Regional template:
<https://docs.google.com/spreadsheets/d/1ElqYiLTLn-Q07doDzHb5vtUCUosFLNbNSgiumm145d4/edit#gid=0>
National backbone template
https://docs.google.com/spreadsheets/d/14CQi67LjJ_hlnrpl8WpTbHmQSW112zzvKPBp6fx8Gw/edit#gid=0
 - e. The Internet2 Performance Working Group Community Measurement, Metrics, and Telemetry project holds meetings on the second Tuesday for those participating or interested. If you are interested, please contact Joe:
[Joe Breen <Joe.Breen@utah.edu>](mailto:Joe.Breen@utah.edu)
 - f. General information about this project can be found at:
<https://spaces.at.internet2.edu/display/PerformanceWG/Internet2+Community+Measurement%2C+Metrics+and+Telemetry+Project>

- g. While NASA policies preclude EOS from sharing this data, EOS has an internal perfSONAR (pS) mesh. They are happy to open their firewalls to permit pS testing by prior arrangement. Contact George at:

["Uhl, George D." <george.d.uhl@nasa.gov>](mailto:george.d.uhl@nasa.gov)

V. Operational network security roundtable

A. Internet2 (Steve Wallace): Two items to cover related to RPKI.

a. Visa incident

- i. The National Registry of India created ROAs for some of Visa's IP resources in India unbeknownst to Visa. The originating ASNs for these new ROAs was an unassigned ASN. The National Registry has an API capability with APNIC for resources that are under the National Registry's umbrella. Service providers immediately evaluated these ROAs and accepted them as valid. (Visa has not done its own ROAs.)
- ii. The Indian ISPs were not doing route source validation so they weren't dropping these routes. But the major transit providers were dropping them. On top of that, some of the resources were hosted in cloud providers. Altogether a pretty large service interruption for Visa.
- iii. In resolving this event two things became clear:
 - 1. The ISPs were ill equipped to troubleshoot. They would look, see the routes weren't in their table, weren't being announced, but had no indication of why. Their procedures did not include determining where a route was dropped and so a route being dropped due to being invalid was invisible to them.
 - 2. The tools and training at Tier 1 ISPs aren't sufficient for this kind of outage.

b. ROAs and DDoS mitigation

- i. Background: A Research and Education institution created ROAs for its IP resources. It came under a DDoS attack and engaged its DDoS scrubber which was cloud based.
- ii. The ROAs created the institution did not allow for the announcement of more specific /24's which was what a cloud based scrubbing provider need to do.
- iii. The major networks saw the more specific, saw they were invalid and treated them as if it was attempted hijacking and therefore dropped them.
- iv. The result: The scrubbing provider was trying to do their job but it wasn't effective.
- v. When planning to working with DDoS scrubbers you need to ensure that your ROAs allow for:
 - 1. The scrubber to use longer prefixes.
 - 2. The scrubber to perhaps issue routes from its own assets and therefore a different ASN.

- vi. In working with a DDoS scrubbing provider, or if when bringing one onboard, it's important to work with them to check any of existing ROAs for these issues. They should also be encouraged to periodically check to verify new and modified ROAs are also compatible.
- c. MANRS should consider a new category of participant for DDoS scrubbing providers as these sort of issues will most likely come up many times.
- d. Questions:
 - i. Is the presentation on Visa's issues public?
The author has permission to do it but it hasn't yet been given. It will be at a future Internet2 NTAC meeting.
n.b. It is currently planned for the NTAC's March meeting.
- e. In looking in Internet2's routing table the number of prefixes that have ROAs is approximately 4%. This is way below the average for the Internet. The barrier may well be that well over half of prefixes announced are not covered by either an ARIN LSA or RSA.

VI. Network roundtable

- A. CAAREN (Andrew Gallo): Everything is nice and quiet.
- B. ESnet (Nick Buraglio):
 - a. The ESnet5 to ESnet 6 migration continues. Very intense to finish the migration, decommission POPs, etc.
 - b. ESnet is working on its first IPv6-onlt pilot program that isn't the ESnet6 management network. The scope is still being decided – anticipated in the next six months.
- C. International Networks – Indiana University (Ed Moynihan):
 - a. ANA: All of the links are now officially up after some issues with various ones over the last half year. International Networks – Indiana University (IN/IU) is working with the other partners to load balance.
 - b. NEA³R added an Arctic focus when it was last renewed. IN/IU is partnering with NORDUNet and discussing with it the best ways to support research in the Arctic. There's a new supercomputer in Finland, a Nordic data center, etc. There are several network projects being driven by Europe looking for ways to link Asia and Europe.
 - i. One, Far North Fiber, has good momentum behind it. The European Union has some funding opportunities for NRENs to be involved in it. The current design has it coming from the Nordic/European area through the Northwest Passage into Seattle, WA, and then on to Japan. A good complement to TransPAC and NEA³R.
 - ii. While Arctic Connect (European/Nordic to Asia via the Northeast Passage) is dead as the European/Nordic partners withdrew, Russia is now laying its own fiber from Murmansk to Vladivostok.
 - c. In the Pacific/Asian are:
 - i. TransPAC is stable.

- ii. Speaking on behalf of Internet2 (I2) regarding the Guam<>Singapore circuit collaboration, the circuit was up, then there was the fiber cut and now it's back up. There's some traffic on it but I2 still needs to complete its tests and accept the circuit. These last two steps should be done in the next few days.
 - iii. APOnet will be coming together once Guam<>Singapore is accepted. IU/IN is putting together a dashboard for it on NetSage. The APOnet agreement is very similar to the ANA agreement. With the agreement and the visibility that the NetSage dashboard should provide better science engagement.
- D. NOAA (Mark Mutz):
 - a. NOAA's N-Wave is still working on its 400G backbone upgrade. In the planning stage along with some lab work on the Juniper PTXs.
 - b. N-Wave is discussing with the provider of its DC area DWDM ring about possible expansions.
 - c. Work is proceeding with NOAA's Cyber Security on TICAP 3.0 and related infrastructure changes.
- E. NRL (Linden Mercer): NRL has reviewed the performance it was able to achieve at SC21. Perhaps finer detail at a meeting this spring but In general NRL is very happy to have been able to fill a 400G pipe with just a couple of streams over the path Virginia to Chicago, IL, to St. Louis, MO. The 200G NIC were not as good as the 100G NICs of last year. NRL is planning on discussions with Nvidia and Mellanox on steps to improve.
- F. PSC/3ROC/XSEDE (Michael Lambert): Nothing new this month.
- G. US Ignite (Glenn Ricart):
 - a. US Ignite is continuing to work with seven communities on connecting the unconnected as a precursor to some of the \$42-\$63 billion that will become available to help connect the under- and un-connected.
 - b. An interesting question: How do you connect to the Internet today? To what? Not the use of CBRS or stringing fiber to interconnect sites. But having down that, then to what?
 - i. Some of these communities are teaming with universities. That's fine as universities know how to connect to the Internet.
 - ii. Some are working with WISPs or local/small ISPs. Some of these provides have challenges in the finer technical details so it would be a great help if JET members reached out these providers to offer a helping hand – capacity planning, security issues, etc. US Ignite can help these initial seven. But when the \$40+ billion comes out, with a \$100 million to a state, help will be needed.

VII. Exchange Points Round Table

- A. StarLight (Joe Mambretti):
 - a. StarLight (SL) is working on projects in Asia, Canada and Europe.
 - b. There is some residual work on the Data Mover Challenge for SupercomputingAsia in a month and a half.

- c. A good deal of effort remains in the development of various test beds (SL supports about two dozen at any point in time).
- d. SL is working with various collaborators, particularly NRL and NASA's GSFC, on building a testbed for 400G, 800G and terabyte end to end flows.
- e. SL continues to work with AutoGOLE, the NOTED testbed from CERN and a NSF funded project on Named Data Networking with Northeastern University.
- f. SL is investigating various ways to integrate Chameleon and FABRIC. Additionally FABRIC with some of the other test beds and Chameleon with the Platform for Advanced Wireless Research.
- g. SL has started its initial planning for SC22 in Dallas next Nov.
- h. Meetings collocated with SupercomputingAsia:
 - i. A mini APAN meeting. This will be on March 2, 1:30-3:30p.m. Singapore time.
 - ii. A mini Global Research Platform workshop. This will be on March 3, 1:30-5:30p.m. Singapore time.

Meetings of Interest 2022

Note: Meetings cancelled since the November JET have been removed from this list. Those moved to a virtual format have been updated.

Jan 16-19, 2022	PTC'22 , Honolulu, HI
Jan, 26-27	Hawaiian Intranet Consortium , virtual
Feb 14-16	NANOG 84 , Austin, TX, hybrid
Mar 1-3	SupercomputingAsia 2022 , Singapore, hybrid
Mar 7-11	APAN53 , Bangladesh, virtual
Mar 8-10	The Quilt Winter Meeting , virtual
Mar 19-25	IETF 113 , Vienna, Austria, hybrid
Apr 24-27	ARIN 49 , Nashville, TN
Jun 6-8	NANOG 85 , Montréal, QC, Canada
Jun 13-17	TNC22 , Trieste, Italy
Jul 23-29	IETF 114 , Philadelphia, PA
Aug (date TBA)	APAN54 , China

Next JET meetings

Note: It is anticipated that JET meetings will remain virtual for the foreseeable future

Feb 15, 2022,	12-2 p.m. ET
Mar 15, 2022,	12-2 p.m. ET
Apr 19, 2022,	12-2 p.m. ET