



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
February 18, 2020 12:00-2:00 p.m. ET

Participants

Hans Addleman, Indiana University
Nick Buraglio, ESnet
Bobby Cates, NASA/Ames
Wallace Chase, REANNZ
James Deaton, GPN
Basil Decina, NRL
Bill Fink, NASA/GSFC
Dale Finkelson, Internet2
Renata Frez, AmLight
Michael Gill, NIH
JJ Jamison, Juniper
Mark Johnson, UNC
Ann Keane, NOAA
Kevin Kranacs, NASA/GSFC – EOS

Yatish Kumar, ESnet
Michael Lambert, PSC/3ROX
Paul Love, NCO/NITRD
Joe Mambretti, StarLight/MREN
Linden Mercer, NRL
Aruna Muppalla, NASA/GSFC
Karl Newell, Internet2
Frank Seesink, UNC
Kevin Thompson, NSF
George Uhl, NASA/GSFC
Italo Valcy, FIU
Adil Zahir, FIU
Matt Zekauskas, Internet2

Proceeding: This meeting was chaired by Kevin Thompson (NSF).

I. Action Items:

- ESnet update on its operational network security use of Rapid7.
- Internet2 and ESnet to updates on their respective optical rollouts.

II. Review of the Minutes of the January meeting: No corrections were offered during the meeting or via email.

III. ESnet High-Touch Services – Yatish Kumar

- A. ESnet6 High-Touch precision network telemetry with 10ns resolution
 - a. A 100Mhz clock has a period of 10ns. A 100Gbit stream implies a 125 byte packet every 10ns. 10ns is also approximately 10ft (assuming a vacuum). All can yield detailed insights into how a network is behaving and how flows are performing. In turn this detailed flow information can be used for traffic engineering, capacity planning and anomaly detection (e.g., AI and ML apps).
 - b. The goal is to be able to monitor every packet that goes through ESnet6.
- B. SmartNICs
 - a. These are the basis for the 10ns telemetry using FPGA/ASIC.

- b. For approximately \$2k in hardware a SmartNIC can do full packet capture at 100Gbps or monitor up to about 1Tbps – the bulk of the traffic on ESnet6. A 4x100 card cost about \$15k.
- c. Because of the design of ESnet6 – OLS, line- & client-side interfaces, and multiple 100Gs – a card with multiple SmartNICs on it can be plugged directly into shared spectrum. Alternatively, it can be plugged into a router if useful. Only a top-of-rack switch is needed. Researchers can program these cards in C and assembly language – they don't need to be hardware designers.
- C. FABRIC: A DWDM overlay on ESnet6 with a lot of compute/storage/DSP in the network over dedicated DWDM waves. FABRIC SmartNIC cards will be compatible with ESnet's FPGA development.
- D. One use case of the wide range of uses of SmartNICs: connecting a 4D Scanning Transmission Electron Microscope uses the same underlying hardware as ESnet's precision telemetry.
- E. Summary:
 - a. Precision telemetry with 10ns timestamps is a game changer in networking insights
 - b. Precision telemetry can be implemented now on any 100G server using built in NIC support for hardware timestamps and Linux kernel support – DPDK telemetry
 - c. FPGA based SmartNICs are affordable and powerful tools for telemetry generation and analysis
 - d. SmartNICs are an emerging tool for all scientific instruments. Including real time DSP/ML and other use cases. Providing scale and customized hardware at a cost-effective price point.

ESnet is working on all of these. It welcomes collaboration with any open group or agency on different aspects of the technology.

IV. Sharing telemetry – James Deaton and Karl Newell

For reference, the two slide sets for this presentation can be found at:

https://docs.google.com/presentation/d/1do7PEkF_2Dh-vhuj6L0b3K1KZ9Xfm116axmcNUEqTto/edit#slide=id.p

https://docs.google.com/presentation/d/1s3_IR6o1CBurqxeVSIhGdDSMXdCfNJCObndsgDgoYQM/edit?folder=177sN27BKA9pSjl1r7VJY0tpPcK2W77EU#slide=id.p

- A. Internet2 started the project of sharing telemetry about 18 months ago with the goal of getting away from the need to solve interdomain performance issue via the phone. At the 2018 Technology Exchange Indiana University's Global NOC (GNOC) showed an augmented trace route tool using the already collected SNMP data from the two nationwide and the approximately dozen networks which used GNOC tools and services. In a single picture an engineer can quickly see if any network in the path is having issues.
- B. Proposed models for shared Telemetry (not exhaustive)
 - a. Streaming telemetry
 - i. Local parser/aggregator pushes telemetry to shared message bus
 - ii. Application is community maintained, runs locally

- iii. Example - Logstash Docker container
 - b. Local SNMP collector
 - i. Local collector pushes telemetry to shared message bus
 - ii. Application is community maintained, runs locally
 - c. Central SNMP collector
 - i. Internet2 run and maintained SNMP collector - pushes telemetry to shared message bus
 - ii. Member will need to provide SNMP access to this collector
 - d. API for submitting data.
- C. What about the privacy of the shared data? Much of the data is already shared by RONS with their members. The larger community, outside R&E, is still a work in progress. To help with the privacy concerns this isn't flow data, just link addresses. The group developing the MOU to expand beyond the data the GNOC is already collecting is working on what can be shared. For example, is it a concern that somebody in the community can see that a particular port on another network is moving 3Gbps this week vs. 5 last week? Use of local collectors will help as they give the local network control over how data is collected and what is shared.
- D. As the project develops several items need to be resolved:
 - a. Who wants to participate?
 - b. What model?
 - c. What telemetry to share (ex: not everything, routers/switches in the ScienceDMZ path, DTN)?
 - d. Decide on the format and metadata
 - e. Scope and design: collectors, aggregators, parsers, message bus.
- E. As the technical decisions are made it should make it easier for the needed policy issues to be resolved.
- F. As a reference point, Internet2's project started off by looking at the path between ScienceDMZ's at the University of Utah and Pennsylvania State University, neither of which used the GNOC's tools or services. The needed data to feed into the augmented traceroute was collected with Globus CLI (and many other tools). Issues were also found (e.g., Globus CLI and pScheduler don't work well together)>

V. Operational network security roundtable (No comments from those in attendance)

VI. Networks Round Table

- A. Indiana University Intentional Networks (Hans Alderman): TransPAC and NEAAR are both working well.
- B. Internet2: (Dale Finkleson and Karl Newell): Optical deployment is proceeding.
- C. NASA EOS (Kevin Kranacs): No changes.
- D. NIH (Mike Gill): No changes.
- E. NRL (Linden mercer): Nothing new.
- F. REANNZ (Wallace Chase):

- a. University of Hawaii and REANNZ have been working to get the Hawaiki cable broken out in Hawaii. This will let REANNZ interact with PIREN to reach Guam. This is anticipated to take 6-8 months.
 - b. A lot of work is going on right now with NSF's Office of Polar Operations to support a relocation of a satellite downlink site in New Zealand for the existing Antarctic service. REANNZ is also working with Polar Operations for a possible new satellite service with a downlink site in Guam.
- G. 3ROX (Michael Lambert): No changes for PSC, 3ROX or XSEDE this month.

VII. Exchange Points Round Table

- A. Ames (Bobby Cates): Assisting with the TIC 3.0 changes. Many of these are aimed at reaching cloud service providers.
- B. StarLight (Joe Mambretti):
 - a. StarLight (SL) has been preparing for SupercomputingAsia (SCA) with a coalition to develop a testbed for the Data Mover Challenge. Results have been good. Unfortunately, due to the COVID-19 virus, the conference has been cancelled.
 - b. SL participated in a DOE Office of Science Quantum Internet Blueprint Workshop. One of the plans is to have a quantum network tying together all the DOE labs. This is complementary to SL's current effort with ANL and FNAL to build a quantum testbed.
- C. MAN LAN and WIX (Dale Finkelson):
 - a. Continuing to work on the transition from the old hardware to the new. Anticipate finishing getting the new software in place in the next ten days to two weeks. Will then begin adding new capabilities. The new hardware is working well.
 - b. No new connections.
 - c. One of the ANA circuits will be down for four to five days starting tomorrow as some undersea repairs are carried out.

VIII. Discussion of the JET's tasking on measurement – Rich Carlson, Kevin Thompson, all

Background: One of the JET's tasks this year is to coordinate the development of tools to monitor cross-domain workflows and facilitate the sharing of measurement data between networks. The data to be anonymized as needed.

Discussion:

- A. Taking what Karl and James were talking about and extending would be a huge help. Making at least some of those tools more extensible will enable people to try out in a broader context. REANNZ agreed that this would be a help for them and will participate. The more of this data available, that is shared, the easier it will be to solve and to end performance problems.
- B. It was remarked that SNMP data, unlike Netflow, doesn't seem to be a security problem. Another concern is on the performance of the hardware. With newer devices performance is no longer an issue for SNMP collection.
- C. Next steps for James and Karl: They are working on a draft MOU. This will permit network leaders to communicate with their broader constituents. The current pilot

needs a better way to collect the data. That is being tackled by the Global NOC. One or two RONS are ready to have their data pulled directly. James, Karl or Joe Breen will try to give an update at the JET's meeting during Global Summit.

IX. Discussion of the JET's semi-annual DC area face-2-face meetings – Paul Love, all

After discussion it was agreed that these meetings have severed their purpose and can be discontinued. The JET will continue to have face-2-face meetings each year – for this year during Internet2's Global Summit and SC. (Note: *Global Summit has since been cancelled.*)

Meetings of Interest 2020

Note: Meetings cancelled since the February JET have been removed from this list.

| | |
|-----------|---|
| Mar 21-27 | IETF 107 . In-person cancelled. A virtual meeting is being developed. |
| Apr 26-29 | ARIN 45 , Louisville, KY |
| Jun 1-3 | NANOG 79 , Boston, MA |
| Jun 8-12 | TNC20 , Brighton, UK |
| Jul 25-31 | IETF 108 , Madrid, Spain |

Next JET meetings

Note: It is anticipated that JET meetings through August will be virtual due to COVID-19 guidelines and the JET's usual summer schedule.

| | |
|--------|--------------|
| Mar 17 | 12-2 p.m. ET |
| Apr 21 | 12-2 p.m. ET |
| May 21 | 12-2 p.m. ET |