



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

July 20, 2021 12:00-2:00 p.m. ET

This meeting was held virtually

Participants

Shawn Armstrong, University of Alaska
Nick Buraglio, ESnet
Rich Carlson, DOE/SC
Bobby Cates, NASA/Ames
Basil Decina, NRL
Bill Fink, NASA/GSFC
Jonah Keough, PNWGP/Pacific Wave
John Heidemann, USC/ISI
Kevin Kranacs, NASA/GSFC - EOS
Michael Lambert, PSC/3ROX
Paul Lang, NASA/GSFC
Ann Von Lehmen, NSF

Paul Love, NCO/NITRD
Joe Mambretti, StarLight/MREN
Linden Mercer, NRL
Edward Moynihan, Indiana University
Aruna Muppalla, NASA/GSFC
Amy Philipson, PNWGP
Glenn Ricart, US Ignite
Anne Richeson, Lumen
Kevin Thompson, NSF
George Uhl, NASA/GSFC
Chris Wilkinson, Internet2

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

I. Action Items:

- Internet2 and ESnet updates on their respective new networks.

II. Review of the Minutes of the June 2021 meeting: Corrections were received and were incorporated in the final version which has been posted on the JET's web page.

III. Observing the IPv4 Internet: from Outages to Work-from-Home - John Heidemann

Note: The slides for this talk are posted on the JET's web page at:

<https://www.nitrd.gov/nitrdgroups/images/2021/JET-John-Heidemann-07202021.pdf>

- A. John and his team have been looking at the IPv4 Internet since 2006. The current goal is to identifying outages, dependencies and the topology of its networks. This understanding in turn leading to a more robust and resilient internet. Their tools can also be used to see outages in near real time and infer the cause. That is, what does the Internet tell about what's happening in the world (for example, the migration to work from home or coups).
- B. A map of the paths of US internet cables shows that outside of the northeast there are relative few paths producing pinch points. In the northeast, while there are more cables there are a lot more places to interconnect so again there are pinch points. When there is an issue, such as a fire in a tunnel between New Jersey and New York City, the tools

developed can reveal these pinch points and cascading dependencies. Once identified steps can be taken to enhance diversity.

- C. Scanning the IPv4 address space:
- a. Currently the IPv4 space (2^{32} , approximately 4B addresses) are scanned every 11 minutes using pings with as light a touch as possible.
 - b. The space is broken into 8 bit blocks (256 addresses, referred to as a “/24”) with available/not available being tracked block by block.
 - i. For clearer visualization the linear address space is show as a square using a Hilbert curve
 - ii. At 600 dpi, with each address a single pixel, the entire space takes 2.8x2.8 meters to show
 - c. One tool looks at a single block over time with each address represented by a row. Over time addresses within the block are pinged and shown as reachable (green) or not (black). A ping may not respond for a variety of reasons – a computer crashed or is turned off, a firewall may intervene, etc. As long as at least one address within a block responds then the block is considered reachable
 - d. Blocks which have never responded are flag as non-responsive and not considered as markers of an outage.
 - e. In pinging the goal is to be as least intrusive as possible – to probe politely:
 - i. If the first address in the block responds then go to the next block. Over time the first address probed in each block is changed
 - ii. If it doesn’t respond try again at a different address within the block
 - iii. If several addresses within the block don’t respond only then is the block flagged as down
 - iv. ISI has been tracking outages since 2013 and has worked to be as non-intrusive as possible with mechanisms in place to let networks op-out, describe what they are doing, etc.
 - f. The projects three key properties in Internet edge outage detection by adaptive probing:
 - i. Principled: probe only when needed (informed by Bayesian inference)
 - ii. Precise: Only some short outages (less that probing interval of 11 minutes) are missed
 - iii. Parsimonious: The project sends between 13 and 19 probes/hr. to a /24 block - only +0.7% background radiation.
 - g. A review of the projects data by an outside team found that some very sparse blocks were falsely flagged as down. These blocks are now fully probed (spread over several rounds of probes) before being marked as down.
- D. After the fact and near real time reporting allows users to see what has happened – say, after a hurricane has passed through. The rate of recover can be examined and used as a basis for improvements for the next event. With the near real time view a user can also watch a country go offline when there has been a political upheaval or the internet is otherwise shit down – say to prevent cheating during a nationwide school exam.
- a. Near real time outages can be observed at:
<https://outage.ant.isi.edu/>

- E. Hidden dependencies
 - a. The project has collected over 45TB of data since outages starting being tracked in November 2013. Approximately 20k observations of 5M blocks.
 - b. To make sense of it two techniques are used:
 - i. Interactive visualizations
 - ii. Automated clustering.
 - c. Visualization - Goal: Make patterns stand out.
 - i. Non-geographic seems the most useful to spot outages. These also help to reveal dependences between networks.
 - ii. By sorting with multi-timescale similarities, rather than IP block address, outage dependencies become much clearer
 - iii. This sort also shows networks that go offline every day at the same time – diurnal networks
 - iv. The sort algorithm has a fairly efficient run time.
 - d. Automated clustering.
 - i. Automate the discovery of network dependencies
 - ii. Insight: failure at the same time, multiple times => dependency
 - iii. Cluster on similarity of fail/recovery events
 - iv. Using this clustering algorithm networks that are dependent on another provider will be clear as their down & uptimes will match, even though their address blocks have no relationship to each other. This is something that pings alone do not reveal and helps to clarify the Internet's topology.
- F. Diurnal networks
 - a. 150-280k address blocks (out of approximately 5M) show patterns that change daily. Usually reflecting people coming into the office during the day or turning on their computer at home in the evening. (*n.b.*: Not as clear in the US as most home networks use an always on modem per an FCC rule in 2001. Western Europe also has less change.)
 - b. Once these networks were identified by examining changes in their patterns of use it was clear when a major change in work habits took place. For example when a country or city went into lockdown due to COVID or when some other event happened which affected much of the population.
- G. Next steps:
 - a. Looking for what else the data can be used to show.
 - b. Working on how to do the same measurement in the IPv6 space (current techniques aren't applicable due to the size of the IPv6 address space).
 - c. How can the project help others use the data?
 - i. Worked with the FCC
 - ii. An API is in the works allowing the near real time export of the data. Who would find that useful? It could help first responders know where they will have network access and where not
 - iii. Are there other applications of the data?

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

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

- A. Prototype/pilot: The various pilots continue to progress. Work continues to get basic measurement data from different universities and RONS.
 - a. Continuing to work with different groups to pull in their data. Additional working sessions need to be scheduled.
 - b. Work continues on different ways to visualize the data.
- B. 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_PNVpeT42nAxa3bW4QostMxcCHTXkWSPbZOpIFwE/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_YoMpa3wjpgwcmIzgy0QsTi1bHb4Qk1cV6qfAM/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 No updates were received.

VI. Network roundtable

- A. International Networks – Indiana University (Ed Moynihan):
 - a. NEA³R:
 - i. Both trans-Atlantic 100G circuits are up and running.
 - ii. International Networks – Indiana University (IN/IU) continues to work with its partners on getting the monitoring and measurement in place and reflected in NetSage.
 - b. TransPAC:
 - i. Circuits are stable.
 - ii. Continue to work within the APOnet consortium on the Guam<>Singapore 100G circuit. Estimated turn-up next month - August.
 - c. IN/IU continues to work with its partners to understand big data flows, who is using them and optimize the performance. IN/IU continues with other trans-oceanic R&E networks to normalize routing to make good use of the new infrastructure. A working group has been brought up inside the GNA with this as a focus.
- B. NASA EOS (Kevin Kranacs): All stable.
- C. NASA GSFC (Bill Fink): Making preparations with Joe Mambretti and Linden Mercer for SC21.
- D. NRL (Linden Mercer): NRL is working with GSFC and StarLight on plans for SC21. It's discussing local connectivity with MAX to support.
- E. PSC/3ROC/XSEDE (Michael Lambert): No updates this month.
- F. University of Alaska (Shawn Armstrong): Nothing to report this month.
- G. US Ignite (Glenn Ricart):
 - a. US Ignite (USI) is working with seven communities to connect the unconnected in their locales. All new to the national networking community. They are particularly interested in for connections health care, governmental services and educational resources. These are primarily coming in through existing ISPs but some are using universities or other research institutions.
 - b. A good number of new institutions should be coming online with new funding:
 - i. NTIA has \$1B to support networking for tribal colleges and tribal areas
 - ii. There is also \$288M in broadband infrastructure
 - iii. Minority serving institutions have also received some finding.
 - iv. Proposals are due 17 August.

VII. Exchange Points Round Table

A. StarLight (Joe Mambretti):

- a. Planning for SC21.
 - i. StarLight (SL) has a panoply of demonstrations planned with GSFC, NRL and others.
 - ii. SCinet is working with SL to support it in general and with a focus on Network Research Exhibition (NRE) demonstrations.
 - iii. To support the testbed needed for the SC21 demos, SL is working with SCinet for 800G between McLean, VA, and SL; and 1.2 Tb between SL and SC. These will include both 400G and 100G circuits.
 - iv. SL is also working on 4x100G to the west coast to support NRE sites. They will probably land at SL but still to be resolved.
 - v. SL is working with the Open Science Grid on a 500G demonstration that could serve High Energy Physics.
 - vi. SL continues to work with CERN on its ML developments in the NOTED project. NOTED attempts to predict large flows and provision circuits in anticipation. The model developed has about a 95% accuracy in its flow prediction. SL has started working with TRIUMF in Canada to use the same techniques. This will be shown in an NRE demo.
 - b. SL and GÉANT are starting to run experiments with their integrated p4 testbeds.
 - c. SL is developing plans to integrate the Chameleon testbed with FABIC.
 - d. The Data Mover Challenge for SupercomputingAsia that SL has assisted in designing has been provisioned.
 - e. SL is doing initial investigations on the technologies needed to provide production services on circuits of 400G and above capacity.
 - f. SL is preparing for the Global Research Platform workshop in September. It is organizing several sessions including Eli Dart describing his vision for DMZ 2.0
- ### B. Ames (Bobby Cates):
- Ames has started getting ready for the new DREN being built by Verizon.

Meetings of Interest 2020

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

Jul 24-30	IETF 111 , virtual meeting
Aug 2-6	APAN52 , virtual meeting
Sep 20-24	The 2nd Global Research Platform (2GRP) Workshop , virtual meeting
Sep 28-30	The Quilt Fall Members Meeting , virtual meeting
Oct 19-20 & 28	ARIN 48 , virtual meeting
Nov 1-3	NANOG 83 , Minneapolis, MN, hybrid meeting
Nov 4	ARIN 48 , Minneapolis, MN, hybrid meeting
Nov 6-12	IETF 112 , in person cancelled, moved to a virtual meeting
Nov 14-19	SC21 , St. Louis, MO. Anticipated to be a hybrid meeting
Jan 16-19, 2022	PTC'22 , Honolulu, HI

Feb, date TBA [APAN53](#)
Feb 14-16 [NANO 84](#), Austin, TX

Next JET meetings

Note: It is anticipated that JET meetings will remain virtual for the foreseeable future. The possible exception will be the November meeting if SC21 remains a hybrid conference.

Aug 17, 2021 12-2 p.m. ET

Sep 21, 2021 12-2 p.m. ET

Oct 19, 2021 12-2 p.m. ET

Nov 16, 2021 8:30-10 a.m. CT,

n.b. This meeting is concurrent with SC21. It will be held in room 123 of the America's Center Convention Complex, St. Louis, MO. Remote access is expected to be available.