



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 19, 2021 12:00-2:00 p.m. ET

This meeting was held virtually

Participants

Shawn Armstrong, University of Alaska

Joe Breen, UTEN/University of Utah

Rich Carlson, DOE/SC

Bobby Cates, NASA/Ames

Basil Decina, NRL

Bill Fink, NASA/GSFC

Mark Johnson, University of North Carolina

Jonah Keough, PNWGP/Pacific Wave

Michael Lambert, PSC/3ROX

Paul Lang, NASA/GSFC

Tom Lehman, Virnao

Paul Love, NCO/NITRD

Joe Mambretti, StarLight/MREN

Linden Mercer, NRL

Inder Monga, ESnet

Ed Moynihan, Indiana University

Aruna Muppalla, NASA/GSFC

Anne Richeson, Lumen

Dan Taylor, Internet2

Kevin Thompson, NSF

George Uhl, NASA/GSFC

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

I. Action Items:

- Continue the discussion on topics for the JET to brief to the LSN - what other topics would of interest. Then narrow down to one or two for discussion with the LSN.
- Internet2 and ESnet updates on their respective new networks.

II. Review of the Minutes of the November 2020 meeting: No corrections were received.

III. An Update on FABRIC and FAB – Tom Lehman and Inder Monga

For reference, the slides for this talk are online at:

<https://www.nitrd.gov/nitrdgroups/index.php?title=JET-Meetings-2021>

Two projects discussed: FABRIC, an NSF mid-scale project, and FABRIC Across Borders (FAB), an NSF IRNC project. FABRIC was awarded more than a year ago with the planning starting more than year before that. FAB was just funded. FABRIC is a four-year project of design and construction. Full operation is expected during year four.

A. FABRIC: Designed to enable new paradigms for distributed applications and Internet protocols

a. A nation-wide programmable network with compute and storage at each node.

i. Designed to run computationally intensive programs and maintain information in the network. Experiments will be able to design and test applications, protocols and services that run at any node in the network.

- ii. Designed from the start to be open to and interconnected with the world.
 - 1. Will retain the sanctity of a testbed.
 - 2. Designed for repeatable results therefore the core nodes and the core edge nodes are connected via dedicated 100G paths.
- iii. Interconnections with national facilities: HPC, cloud and wireless testbeds, domestic and international R&E networks, commercial clouds, commodity internet and local (to a node) resources: computing, storage and BYOE (PCI card, server or switch) via layer 1 or 2. Will use Internet2's AL2S or ESnet's OSCARS where necessary. FABRIC will only advertise public IPv6 address space and receive the full IPv6 route table. IPv4 is used for intra-fabric routing.
- iv. Services will be optimized for science workflows.
- v. 300+ TB non-persistent storage along with programmable GPUs, FPGAs and NICs inside the network. They will have:
 - 1. Network programming at the level of OpenFlow, P4, eBPF, DPDK
 - 2. GPUs to support ML applications
 - 3. Ability to interpose compute, memory and storage into the path of fast packet flows
 - 4. Processing speeds at 25Gbps, 40Gbps, 100Gbps, Nx100Gbps
 - 5. Access to the hardware directly for experiments (programmable network cards, GPUs, FPGA cards) via either a FABRIC bare metal server or a VM/container in a FABRIC server
 - 6. FABRIC dataplane switch/router (Cisco NCS 5500) using MPLS-SR
 - 7. Multiple servers
 - 8. All data plane interconnects at 100G+
 - 9. Sliceable, programmable switching, hierarchical storage and in-network compute
- vi. Quality of service (QoS) with dedicated optical 100G.
 - 1. Priority: Network control, deterministic services (choice of oversubscribed traffic being dropped or place in best effort queue), best effort
- vii. Measurement hardware based on ESnet developed High Touch hardware (see minutes of February and October 2020 JET meetings: <https://www.nitrd.gov/nitrdgroups/index.php?title=JET-Meetings-2020>):
 - 1. GPS clock, programmable port mirroring,
 - 2. Able to measure power consumption, optical layer, CPU, memory, disk, port/interface utilization, etc.
- b. Deployment plans for initial 29 nodes – through June 2023:
 - i. Development Phase (3 nodes): April 1, 2020 – September 30, 2021
 - 1. RENCi, University of Kentucky and LBNL
 - ii. Phase 1 (16 nodes): July 1, 2020 – September 30, 2021 (all gear for Phase 1 is at RENCi)
 - iii. Phase 2 (10 nodes) : April 1, 2022 – June 30, 2023

- c. Node types:
 - i. 9 ESnet Core nodes co-located at ESnet6 Points of Presence.
 - 1. Connected via dedicated 100G DWDM across the new ESnet6 open-line optical system
 - 2. Sited at intersections of multiple high-capacity dedicated optical links.
 - 3. Four of these sites will be upgraded to Terabit SuperCore nodes at some point during Phase 2: Washington, StarLight (Chicago), Dallas and San Diego
 - ii. 20 edge nodes distributed across the R&E community at various regional networks, major cyberinfrastructure facilities, and university hosting sites.
 - 1. Working to get as many connected via 100G layer 1 as possible (CoreEdge nodes). Where layer 1 isn't possible then layer 2 (Edge nodes).
 - 2. Located on campuses, regional networks, and R&E facilities.
 - iii. All three types of nodes have the same basic functionality – programmability, compute and storage.
 - d. FABRIC's network control plane uses Cisco's Network Services Orchestration
 - e. FABRIC's network services: Layer 3 routing, layer 2 connections and layer 3 VPNs
 - f. An experiment (slice) can be completely inside FABRIC or can make use of FABRIC's multiple external connections to reach external experimental and production resources.
 - g. FABRIC Security
 - i. Specific plans and best practices are under development. Section 7 of the FABRIC Network Services and Perring Design document contain more details.
 - ii. Default approach:
 - 1. Experiment slices are isolate by default
 - 2. Specific approvals are needed to enable automated external access connections. The granularity of this access is to be defined
 - h. FABRIC Design Documents
 - i. <https://fabric-testbed.net/resources/design-documents>
 - ii. Most relevant: Topology Design, Network Services and Peering Design, Site Questionnaire (Deployment and Hosting Plan) and Partner Security Questionnaire
 - i. Questions and additional details:
 - i. info@fabric-testbed.net
 - ii. <https://fabric-testbed.net>
- B. FAB (FABRIC Across Borders): Enabling Global Experimentation
- a. Global FABRIC nodes: Japan (University of Tokyo), UK (University of Bristol, EU (University of Amsterdam) and CERN
 - b. Science use cases and partners for FAB
 - i. Astronomy (Vera Rubin Observatory/LSST, Chile)

- ii. Cosmology (CMB-S4)
- iii. Weather (University of Miami & CPTEC, Brazil) - Ben Kirtman, Atmospheric Science & Paolo Nobre
- iv. High-Energy Physics (CERN) - Rob Gardner, FAB Co-PI, Physicist
- v. Urban Sensing/IoT/AI at Edge (University of Bristol) - Dimitra Simeonidou, Professor of Networking
- vi. 5G across borders, P4/SDN - (University of Tokyo) Aki Nakao, Professor of CS; KISTI (Korea Institute of Science and Technology Information)
- vii. Censorship Evasion - Richard Brooks, Professor of ECE
- c. Europe facility partners
 - i. NEA³R (Networks for European, American, African and Arctic Research)
 - ii. ANA (Advanced North Atlantic)
 - iii. ESnet
 - iv. CERN
 - v. GEANT Open Exchange London & Paris
 - vi. NetherLight Open Exchange
 - vii. SURFnet
 - viii. University of Bristol
 - ix. University of Amsterdam
 - x. University of Antwerp
 - xi. SAGE (NSF mid-scale project)
- d. Asia-Pacific facility partners
 - i. TransPac
 - ii. WIDE (to TransPac)
 - iii. University of Tokyo
 - iv. Japanese Science Information Network (SINET)
 - v. Korean Institute of Science and Technology Information (KISTI)
 - vi. Korea Research Environment Open Network (KREONET)
 - vii. StarLight International/National Communications Exchange Facility
 - viii. The additions of Guam and Hawaii are under consideration
- e. South America facility partners
 - i. AmLight International Exchange Point
 - ii. FIU
 - iii. University of Miami
 - iv. Center for Weather Forecast and Climatic Studies (CPTEC)
 - v. AmLight Express and Protect
 - vi. Academic Network at Sao Paolo (ANSP)
 - vii. RedCLARA

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

- A. Prototype/pilot: The various pilots are progressing. Working to get basic measurement from different universities and RONS.
 - a. The C-Light pilot (Clemson University) has recently completed an iteration. Dan Doyle's augmented traceroute shows the measurements crossing five different

- domains: Indiana University, Indiana GigaPoP, Internet2, SOX and ending at C-Light. Next steps will be to cross C-Light and into Clemson.
- b. Work is progressing on a similar pilots with AmLight and OSHEAN.
 - c. Dan Doyle has finished an iteration on the project's Telegraf container so there is now a plugin. Karl Newell just bundled it and it will be tested this week in the University of Utah. The goal is to deploy at NCSA in the next few weeks as NCSA want to use it with their existing Telegraf.
- 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 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
The spreadsheet also has an embedded link to measurement templates for campus, regional and national networks setting out what data is desired. See: <https://drive.google.com/drive/folders/1l-LRyriI6u4AvBeY6NlvyYYaINRpiByA>
 - 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. While NASA polices 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. 3ROX (Michael Lambert): Greenbone is drastically cutting back on the freely available versions of its Security Manager packages.

VI. Network roundtable

OMB signed the new IP-v6 policy on November 19, 2020. It contains several due dates for various steps – 45 days (January 3, 2021), 180 days (May 19, 2021), end of FY2021, etc.

<https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-07.pdf>

- A. International Networks – Indiana University (Ed Moynihan):
 - a. TransPAC:

- i. One of the Guam<>Hong Kong 10G circuits was decommissioned last week. The second 10G Guam<>Hong Kong remains in operation.
 - ii. International Networks – Indiana University is working to be part of a consortium that is planning a 100G Guam<>Singapore circuit by the end of 2021Q1.
 - b. NEA³AR: The new Amsterdam<>New York via Copenhagen 100G circuit is operational. The second new 100G, New York<>London, has experienced some delays, probably COVID related, and is now anticipated mid-February. The old New York<>London 100G remains in use.
- B. NASA GSFC (Bill Fink/Paul Lang/George Uhl): No update today.
- C. NRL (Linden Mercer): Nothing to report today.
- D. Pacific Wave (Jonah Keough):
 - a. Pacific Wave (PW) is part of the consortium working on a new Guam<>Singapore 100G circuit. In Guam it will make use of the relocated GOREX.
 - b. PW is working with UC San Diego to place IGROK servers at PW nodes for monitoring and measurement. Internally PW has been testing these over alien waves.
- E. 3ROX (Michael Lambert): No updates for 3ROX or PSC. For XSEDE the process of connecting the Massachusetts Green High Performance Computing Center (MGHPCC) to XSEDE's L3VPN is just starting. This is in support of its Open Storage network node.
- F. University of Alaska (Shawn Armstrong): No update today.
- G. University of Utah (Joe Breen via chat/email):
 - a. Utah Education and Teaching Network (UETN) is rolling additional 100G circuits around the state of Utah to build out the metro optical rings across the state.
 - b. UETN is building out 400G to the University of Utah (UofU) Science DMZs and to Internet2's POP. Anticipating completion in February.
 - c. University of Utah (UofU) is finishing rollout of the 400G Science DMZ with new clustered 100G DTNs, new 100G perfSONAR boxes, new Protected Environment (PE) DTNs with new perfSONARs and new virtual hosts for PE science. All equipment deployed and functional with final testing and tuning remaining to be complete.
 - d. UofU is connecting a new Cloudlab3 Testbed at 400G to its Science DMZ infrastructure. Anticipated to be complete by February.
 - e. UofU and UETN are working on data center requirements (rack, power, space, access for GPS antenna) and circuits to connect the incoming FABRIC Testbed at layers 1, 2, and 3. Equipment is starting to arrive with deployment in February.

VII. Exchange Points Round Table

- A. StarLight (Joe Mambretti):
 - a. StarLight (SL) is in discussions with GÉANT to connect their P4 testbeds.
 - b. SL continues to work with the international AutoGOLE/SENSE consortium to develop a global multi-point testbed.
 - c. SL is working with CERN and a high energy physics consortium for a point-to-point path CERN<>FNAL<>SL.

- d. SL continues to support the FNAL integration project of the Rucio data manager, FNAL's Big Data Express and DOE's SENSE data orchestrator.
 - e. SL's FABRIC node is due to arrive shortly.
 - f. SL continues to work with Internet2 (I2) on I2's new optical backbone. The hardware is expected by month's end. As part of I2's NGI deployment in Chicago, IL, it is building an 800G metro ring between 600 West Chicago, Equinix and SL.
 - g. Plans are already underway for SC21 in St Louis, MO.
 - h. Plans are being developed for the Global Research Platform September 2021 meeting.
<http://www.theglobalresearchplatform.net/meetings/2nd-global-research-platform-workshop/#more-72>
- B. Ames (Bobby Cates): Current status of USGS' migration from Menlo Park, CA, to Ames:
- a. The new network to support the migration is complete. Due to COVID, unfortunately, USGS cannot move and make use of this parallel network.
 - b. A multi-agency TIC was not developed. USGS will use an existing DOI TIC in the Sacramento, CA, area, extending it to Ames.
- C. PNWGP (Jonah Keough): Nothing to report today.

IX. LSN update – Rich Carlson

The Large Scale Networking IWG (LSN) is the group the JET reports to, part of the Office of Science and Technology Policy. The LSN is composed of the JET and several other teams. For 2021 the LSN is restructuring its meetings slightly to ensure that what each team is doing is visible at the larger, federal level. Part will be a monthly, short report out of the teams' activities. Part will be a clearer mechanism for each team to request more detailed information from the LSN either on topics it's working on or on something one of the other teams is working on.

A third part will be each team organizing one, perhaps two, detailed briefs a year to the LSN on a topic the team is working on. Each team would work with the LSN in the selection of a topic. These detailed briefs would be given by somebody working on the project, rather than the team's chairs or coordinator.

The JET needs to develop a list of possible topics for these briefs. From discussion four were put forward as a start:

- FABRIC & FAB
- The JET's tasking on tools to help with inter-domain issues lead by the team from Internet2's Performance WG
- ESnet 6
- Internet 2's NGI

An example of a good topic to make the federal program managers who comprise the LSN aware of would be the FABRIC project just described. It would be good if they knew of it, what an asset it is and how their researchers might have an opportunity to use it. The goal for these briefs to the LSN is to spread the word beyond the brief's own community. In the case of something like FABIRC, to spread the word beyond the community of NSF researchers.

In response to a question asking if higher levels of briefing should be considered the reply was that briefs to the LSN would be focused within the NITRD community.

Action Item: • Continue the discussion on topics for the JET to brief to the LSN - what other topics would of interest. Then narrow down to one or two for discussion with the LSN.

Meetings of Interest 2020

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

Jan 17-20, 2021	PTC'21 , in person cancelled, moved to a virtual meeting
Feb 1-5	APAN 51 , virtual
Feb 3 (morning)	Asia-Pacific Research Platform Workshop , virtual (collocated with APAN 51)
Feb 8-10	NANOG 81 , virtual
Feb 9-11	The Quilt Winter Member Meeting , virtual
Feb 17-18	HIC , virtual
Mar 2-4	Supercomputing Asia 2021 (SCA21) , virtual
Mar 3 (afternoon)	Asia-Pacific Research Platform Workshop , virtual (collocated with SCA21)
Mar 8-12	IETF 110 , virtual
Apr 11-14	ARIN 47 , in person cancelled, moved to a virtual meeting
Apr 18-21	Internet2 Global Summit , virtual
Jun 14-16	NANOG 82 , Minneapolis, MN
Jun 21-25	TNC21 , hybrid cancelled, moved to a virtual meeting
Jul 24-30	IETF 111 , San Francisco, CA

Next JET meetings

Note: It is anticipated that many of the JET's meetings in CY2021 will be virtual due to COVID-19 guidelines.

Feb 16, 2021	12-2 p.m. ET
Mar 16, 2021	12-2 p.m. ET
Apr 20, 2021	12-2 p.m. ET