



### 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  
March 17, 2020 12:00-2:00 p.m. ET

#### Participants

Shawn Armstrong, University of Alaska  
Joe Breen, UTEN/Univ of Utah  
Nick Buraglio, ESnet  
Rich Carlson, DOE/SC  
Bobby Cates, NASA/Ames  
James Deaton, GPN  
Dave Diller, MAX  
Bill Fink, NASA/GSFC  
Dale Finkelson, Internet2  
Andrew Gallo, CAAREN/GWU  
Michael Gill, NIH  
Ann Keane, NOAA  
Kevin Kranacs, NASA/GSFC – EOS  
Padma Krishnaswamy, FCC  
Michael Lambert, PSC/3ROX  
Paul Lang, NASA/GSFC

Joyce Lee, NCO  
Paul Love, NCO/NITRD  
Joe Lyles, ORNL  
Joe Mambretti, StarLight/MREN  
Linden Mercer, NRL  
Inder Monga, ESnet  
Alex Moura, RNP  
Aruna Muppalla, NASA/GSFC  
Mark Mutz, NOAA  
Glenn Ricart, US Ignite  
Anne Richeson, CenturyLink  
Dan Taylor, Internet2  
George Uhl, NASA/GSFC  
Ralph Wachter, NSF  
Chris Wilkinson, Internet2

**Proceeding:** This meeting was chaired by Rich Carlson (DOE/SC).

#### **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 February meeting: No corrections were offered during the meeting. One was received after the meeting.

#### **III. Brief on the FABRIC project – Inder Monga**

For reference, the slides from this presentation are online at:

<https://www.nitrd.gov/nitrdgroups/images/4/41/JET-Inder-Monga-03172020.pdf>

- A. FABRIC is a testbed that brings the network together with storage and compute. It is designed to fill a gap in the CISE research infrastructure: future edge cloud and future core network.
- B. Why FABRIC
  - a. The mantra of the last 20 years – “Internet is showing its age”
    - i. Applications designed around discrete points in the solution space.
    - ii. Inability to program the core of the network.

- c. What changed?
  - i. Cheap compute and storage that can be put directly in the network.
  - ii. Multiple established methods of network programmability (OpenFlow, P4, eBPF, DPDK, BGP flowspec).
  - iii. Advances in both machine learning (ML) and artificial intelligence (AI) that can be combined with the in-network compute and storage.
  - iv. Emergence of 5G, IoT, and various flavors of cloud technologies.
- d. Opportunity for the community to push the boundaries of distributed, stateful, “everywhere” programmable infrastructure at scale. Questions were:
  - i. More control plane or data plane state, or some combination? Multiple architectures coexist in this space?
  - ii. Network as a big-data instrument? Autonomous network control?
  - iii. New protocols and applications that program the network?
- C. Security must be an integral component from the beginning.
- D. The FABRIC team is made up of many; Ilya Baldin (RENCI) is the lead PI and Anita Nikolich (IIT) is the co-director.
- E. The goal of Fabric is to provide a research infrastructure with a broad reach:
  - a. Support new network architectures and science applications.
  - b. Advance cybersecurity.
  - c. Provide an environment where existing and new wireless testbeds, HPC and IoT can be interconnected.
  - d. ML and AI integrated. Facilitated by in-network GPUs.
  - e. Support the training of the next generation of computer science researchers.
- F. FABRIC’s core is built on top of ESnet6.
  - a. With ESnet6 built OLS it was easy to give FABRIC an 100G wave across the ESnet footprint. The wave is terminated only at FABRIC nodes thereby completely insulating ESnet production traffic.
  - b. To push the technology to the next generation it was decided to additionally build a 1Tbps path between SDSC, TACC, NCSA and PSC. The timing of this is still to be determined.
- G. The FABRIC partner Internet2 will provide a path to the public clouds for any application. FABRIC also connects to Chameleon and CloudLab.
- H. At the edge FABRIC works with many – supercomputer centers, RONS, and individual institutions. It is anticipated that additional core nodes will be built off the ESnet footprint by RONS.
- I. The FABRIC team is developing the needed software to request, reserve and schedule resources. Part of the software being developed will permit resources to include a guaranteed bandwidth. Resources can be reserved for relative short periods of time or for long running experiments.
- J. FABRIC nodes include both management and user provision-able resources. It can be thought of as a “disaggregated router”:
  - a. All ports interconnected by a 100G+ switch programmable through testbed control software.

- b. Network cards with high speed interfaces: 25G, 40G and 100G. 200G+ when available:
    - i. Programmable interface cards (hardware OVS offload + DPDK).
    - ii. Reconfigurable interface cards (FPGA and P4/network processors).
  - c. High-performance servers equipped with GPUs, FPGA compute accelerators, NICs, SmartNICs, and NVMe drives. Storage is short term & shared high volume. While not meant to be persistent there will be flexibility as experimenters need.
  - d. It is anticipated that users will be able to place their own switches in nodes.
  - e. Node resources are sliceable for concurrent use by multiple experimenters.
  - f. Each node also has management and data plane switches along with a server for its management.
- K. Network Services
- a. Link both different elements of a requested topology together within FABRIC and with resources outside FABRIC.
  - b. Possible services for an experiment to use:
    - i. Layer 2 on-demand with bandwidth provisioning or best-effort (implies the user will build their own layer 3).
    - ii. Layer 3 (IPv6 and IPv4) best-effort.
    - iii. Layer 3 peering between the experiment's topology and an existing production network such as a campus.
    - iv. Layer 2 peering between the experiment's topology and a cloud provider (Google, AWS or Azure, via Internet2's CloudConnect).
    - v. VPN from a FABRIC node to an experimenter's desktop or a campus resource.
- L. Measurement within FABRIC
- a. The measurement framework is designed to be adaptable, programmable, scalable, extensible, and shareable:
    - i. Used to collect, store, and publish measurement data from users and for the management of FABRIC.
    - ii. Supports a common/shared message bus infrastructure based on publish/subscribe technology.
    - iii. Supports efficient filtering and searching. There may also be limited processing of measurement data.
    - iv. Interfaces with multiple UIs and alert systems.
  - b. Fine-grained precise measurements
    - i. Leverages highly accurate PTP signal from a node-local GPS receiver
    - ii. Supports precise timestamping of packets using NIC cards (a.k.a., PacketGPS)
    - iii. This is technology sharing with the SmartNICs being developed for ESnet6.
  - c. There will be supports for on demand high-speed packet capture with limited processing.

- M. FABRIC will use early, “FABRIC ready” experiments from security, IoT, ML in the network, named data networking, and advanced transport protocols to help formulate the design, validate the implementation, and in general debug.
- N. The project kicked off in October 2019. Current planning calls for initial deployment of a small number of core and edge nodes beginning in year 2 of the project, deployment of the balance of the nodes will start in year 3 with full operations by the end of year 4. Between July 1, 2020 and September 30, 2021 19 nodes will be deployed with the initial sites being Chicago/StarLight, RENCi, University of Kentucky, and LBNL.
- O. FABRIC isn’t:
  - a. An isolated testbed. It will peer at layers 2 and 3 with a variety of other networks
  - b. A place for long term production workloads.
  - c. A place for real protected data (PII, HIPA, etc.).
  - d. A new fast pip between connected facilities. It is a place to experiment with new approaches.
- P. FABRIC is:
  - a. An everywhere programmable testbed combining core and edge components with links to outside facilities.
  - b. A multi-user facility supporting concurrent experiments over varying scales using a federated authorization system with allocation controls.
  - c. A place to experiment on new network architectures, protocols and distributed applications.
  - d. Extensible with new facilities connecting over time.
- Q. FABRIC wants to build a community of experimenters, facility partners, regional and national networks, government agencies interested in network research, and partners in industry. It is planning on holding one or two community events and workshops a year.
  - a. The next workshop will be held virtually April 15. See: <https://fabric-testbed.net/events/fabric-community-workshop-2020>

**IV. Operational network security roundtable** (only those who had comments are noted)

- A. Internet2 (Chris Wilkinson): Internet2 is doing enhanced monitoring of its flow data to help with the COVID-19 response.

**V. Networks Round Table**

- A. CAAREN (Andrew Gallo): No update.
- B. ESnet (Nick Buraglio):
  - a. ESnet has worked with its sites making sure there is adequate capacity for remote working as staff transition to working from home.
  - b. ESnet6 activities are continuing:
    - i. The packet RFP is nearly finished. Anticipated to be awarded in April or May.
    - ii. The facilities and deployment of the management infrastructure underway and proceeding well with very few surprises. The management net is IPv6 only.
    - iii. Deployment of the new OLS is also proceeding.

- C. Internet2 (Chris Wilkinson):
  - a. Internet2 (I2) has spent a lot of effort preparing for and dealing with the impact of COVID-19 on the network. It is seeing shifts in traffic flows and expects these to grow as students on spring break return to campus. I2 is seeing spikes with Zoom and residential providers and is moving many of these peerings to 100G.
  - b. Next Generation:
    - i. I2 is working RFP on technology validations with the four finalist vendors for the packet. It is anticipating the award to be made in May or June.
    - ii. For its optical deployment I2 is doing fiber acceptance on most of its footprint. They are aiming for the first transition to be in early April centered on Kansas City, MO. Transitions will expand from that area as more fiber is accepted. As many of the transitions will be hot cuts they are proceeding with caution especial on the initial ones. The target for completion is the end of the year.
  - c. Internet2 has a new software manager: Mike Simpson
- D. NASA EOS (Kevin Kranacs): No update.
- E. NOAA (Mark Mutz):
  - a. The final piece of N-Wave's 100G upgrades are in Colorado on FRGP's Denver<>Boulder path.
  - b. Several small projects in the rest of the network are going forward. N-Wave is working with other portions of the Department of Commerce to provide them with connectivity.
- F. NRL (Linden Mercer): NRL is starting its planning for SC20 next November in Atlanta, GA.
- G. RNP (Alex Moura): RNP is working on a number of core network upgrades.
- H. 3ROX (Michael Lambert):
  - a. 3ROX: Tracking traffic flows as people transition to working from home.
  - b. PSC: No changes.
  - c. XSEDE: Working to get 100G perfSONAR boxes tested and deployed. Best that has been achieved with iPerf3 is 27-29Gbps. iPerf4 can fill the pipe.
- I. US Ignite (Glenn Ricart): Additional Digital Town Squares (DTS) joining. New locations are in Springfield, OR; Phoenix, AZ; and Las Vegas, NV. Some of these are becoming more regional in nature, covering a larger area than previous DTS. US Ignite is working to help DTS to coalesce their technologies from the varieties now used - private networks, LoRanWAN, etc. - that need to go through gateways to interconnect. The internet is being used as an intimate network. Some applications, e.g. traffic lights, are not run over the public internet due to security concerns. Security needs to be enhanced such that these critical applications can be run over a public network. DTS are looking at how they can have edge compute to support critical applications and the IoT.
- J. UTEN (Joe Breen): UTEN is expanding throughout the state. Much is currently at 100G with the capability of going to 400G as needed. It is preparing for the FABRIC roll out.
- K. GPN (James Deaton): GPN is getting ready for FABRIC. It is also watching flows as traffic changes due to jump in working remotely. It is also working hard on measurement sharing to bring more transparency to R&E networks.
- L. University of Alaska (Shawn Armstrong): Nothing new to report.

## VI. Exchange Points Round Table

- A. MAN LAN and WIX (Dale Finkelson): Software testing is completed with the two sites fully functional. The ANA cable was successfully repaired – all ANA circuits are up.
- B. MAX (Dave Diller): MAX is getting ready for FABRIC. No issues to date due to its members users moving to working remotely.
- C. StarLight (Joe Mambretti):
  - a. StarLight (SL) has started its planning for SC with MREN and other partners.
  - b. SL worked with an international collaboration to provide an international testbed for the Data Mover Challenge at SupercomputingAsia in which SL was also a competitor. Though the conference was cancelled, SL was awarded Best Integration of Data Intensive Science.
  - c. SL is preparing for FABRIC.
  - d. SL is working with ESnet to implement the new racks for ESnet6.
- D. Ames (Bobby Cates):
  - a. Campus is closed until April 7 with only mission critical access is permitted.
  - b. Continuing to work on TIC 3.0 and its acceptable cloud topologies. This is ongoing with recommendations being accepted at least through June.
  - c. USGS' migration continues, though paused for the moment. One more large group from Menlo Park, CA, later this year to transition with a few individuals after that.

## VII. Discussion of the JET's tasking on measurement

- A. Current status of efforts lead by Joe Breen, Karl Newell and James Deaton
  - a. Of the four options for telemetry of network data effort has focused on the site operating the SNMP collector and pushing to the central location.
  - b. Work on the draft MOU continues with a preliminary version anticipated in a few weeks. When available this will permit network leaders to communicate with their broader constituents.

## Meetings of Interest 2020

Mar 21-27	<a href="#">IETF 107</a> . In-person cancelled, moved to a virtual meeting
Apr 26-29	<a href="#">ARIN 45</a> , In person cancelled, moved to a virtual meeting
Jun 1-3	<a href="#">NANOG 79</a> , in person cancelled, moved to a virtual meeting
Jul 25-31	<a href="#">IETF 108</a> , Madrid, Spain
Aug 5-7	<a href="#">Network Automation &amp; Cloud NetSec Workshops</a> , Bloomington, IN
3-7 Aug	<a href="#">APAN50</a> , in person cancelled, moved to a virtual meeting
Oct 5-8	<a href="#">TechEX20</a> , Atlanta, GA
Oct 19-21	<a href="#">NANOG 80</a> , Seattle, WA

## Next JET meetings (all meetings are virtual)

Apr 21	12-2 p.m. ET
May 19	12-2 p.m. ET
Jun 16	12-2 p.m. ET