



### 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  
April 16, 2019, 12:00-2:00p.m. ET

#### Participants (\* in person)

Hans Addleman, Indiana University	Padma Krishnaswamy, FCC
Rich Carlson, DOE/SC*	Michael Lambert, PSC/3ROX
Bobby Cates, NASA - Ames	Paul Love, NCO*
Basil Decina, NRL*	Chris Lowe, USDA/ARS
Bill Fink, NASA/GSFC	Rich McGeer, US Ignite/University of Utah
Mike Gill, NIH*	Linden Mercer, NRL*
Patty Giuntoli, ESnet	Glenn Ricart, US Ignite/University of Utah
Greg Grimes, Mississippi State University	George Uhl, NASA/GSFC
JJ Jamison, Juniper*	Chris Wilkinson, Internet2
Kevin Kranacs, NASA/EOS	

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

I. **Action Items** (carry forward): Plans for ESnet6's Optical Core were updated (now to be done as a webinar, date to be announced).

II. **Review of the Minutes** of the March 2019 meeting: No corrections to the minutes for the March 2019 were mentioned. (Note: Some were received after the meeting and have been incorporated into the final minutes.)

III. **Operational network security roundtable** (only those who had comments were noted)

- A. 3ROX: Michael Lambert: 3ROX has recently become MANRS certified
- B. US Ignite: Glenn Ricart: Stay tuned for the Edge Network talk.

IV. **Networks Round Table**

- A. ESnet: Patty Giuntoli: ESnet is working very hard for the final design review of ESnet6 scheduled for May 15-17, 2019. It received proposals for the Open Line System in mid-March. Proposal review is in progress.
- B. Internet2: Chris Wilkinson: Internet2 is making expansions for its operational capacity at its peering locations: Chicago, Ashburn, and the Bay Area. This includes going to 100G with major sources such as Google.
- C. International networking at Indiana University: Hans Addleman
  - a. NEEAR: status quo.

- b. TransPAC: peering with LHCONe is up for the pair of 10G circuits between Guam and Hong Kong. TransPAC hopes to bring up peering with CERnet as well.
- D. SCinet: Hans Addleman
- a. SC's wide area network will be very exciting this year. Old vendor/partners are coming back and are being joined by providers of metro and DCI products.
  - b. If you expect to need bandwidth to SC20 in Denver in November, please complete the bandwidth request sent to the JET's list. It can be found at <https://goo.gl/forms/kAu0mKjnfkaw62ap1>
  - c. If you have any questions, please contact the SCinet WAN Team chair at [addlema@iu.edu](mailto:addlema@iu.edu)
- E. 3ROX: Michael Lambert
- a. 3ROX: No changes.
  - b. XSEDE:
    - i. The migration to the layer 3 VPN over Internet2 has been completed.
    - ii. The old VLAN mess has been turned down.
    - iii. No performance degradation, no complaints from users about the migration.
- F. US Ignite: Glenn Ricart: Stay tuned for the EdgeNet talk.
- G. Mississippi State University/MISSION: Greg Grimes
- a. MISSION has completed the transition from a 10G ring provided by AT&T to a 100G ring provided by CSPIRE. Any bumps in the transition have been worked through.
  - b. MSU is connected at 100G.
  - c. Of the eight public universities in Mississippi four are research universities (MSU, University of Mississippi, Jackson State University and Southern Mississippi University) and have been connected to the original MISSION. The other four (Delta State University, Alcorn State, Mississippi Valley State and Mississippi University for Women) had only commodity level connectivity but are now connected to MISSION at 10G and becoming members of Internet2 - and they are now paying less for their 10G connections than they were paying for 1G of commodity service.

## V. Exchange Points Round Table

- A. WIX and MAN LAN: Chris Wilkinson
- a. Internet2 has turned up a few more 100Gs on those systems.
  - b. Internet2 is looking at the best way to update both before the Next Gen packet service is available. Not only is the gear end of life, one is now out of slots.
- B. StarLight: Joe Mambretti (via email): StarLight is working with a consortium to organize the initial Global Research Platform Workshop. It is scheduled for September 17-18 in the CalIT2 building at UCSD coincident with the GLIF meeting.
- C. NASA Ames: Bobby Cates
- a. USGS' transition from its old home in Menlo Park to Ames is ongoing. The transition team is on-site. The scientists are scheduled to arrive later this year. A

- 10G connection to Menlo Park is in place. Other circuits are a 10G to CENIC's Sunnyvale POP along with their commodity internet service.
- b. Working to support DREN's Technical Interchange Meeting which will be held at Ames in September.

## VI. ESnet6 Project Update: Patty Giuntoli

- A. ESnet is approximately three years into the design, approvals, procurement and construction of ESnet6 with a major milestone, its final design review is scheduled for May 15-17, 2019. Construction is projected to start mid-2020.
- B. ESnet expects at least 18.8 Tbps of bandwidth on the optical core with an upper expectation of 24.6 Tbps.
- C. The RFPs for the optical transmission portion of the project are under review.
- D. ESnet6 design objectives:
  - a. No single point of failure within the ESnet network backbone.
  - b. Agility to add and/or move bandwidth capacity.
  - c. Fine-grained traffic engineering.
  - d. Comprehensive automated management.
  - e. Highly programmable and flexible services.
- E. ESnet6 optical core design characteristics:
  - a. High bandwidth and scalable capacity:
    - i. Variable modulation rates (e.g., 100-400G waves) and Flex-grid for spectral efficiency.
    - ii. Design for full fill (in C-band).
    - iii. 100GE handoffs day 1, 400GE supported opportunistically (requires "FlexO-ish" solution).
  - b. Hardware and Provisioning Resiliency:
    - i. Topology of interconnected rings.
    - ii. High availability and redundant hardware, with distribution of connections across redundant components.
    - iii. SRLG-like provisioning of waves across disparate hardware.
  - c. Flexible:
    - i. Designed for colorless, directionless, and (limited) contention-less.
    - ii. Open Line System with third party transponder support.
  - d. Secure: Isolated (walled-garden) management plane with protected access for 3rd party L1 NOC.
  - e. Programmable: Automated provisioning, will require integrated 3rd party transponder and OLS management.
- F. ESnet6 packet core design characteristics:
  - a. High bandwidth with fine grain traffic engineering:
    - i. MPLS switching within the Packet Core, eliminating overhead of routing.
    - ii. Distinct set of SR-TE LSPs associated to each service instance.
    - iii. QoS for user traffic to support guaranteed bandwidth services.
  - b. Packet layer resiliency:
    - i. TI-LFA FRR will facilitate fast recovery for local failures.

- ii. Unrecoverable SR-TE LSP failures will result in fall back to ISIS-SR LSPs.
    - iii. S-BFD will allow the LSP head-end router to quickly detect non-local failures and react accordingly.
  - c. Flexibility: Use of PCE to manage SR-TE LSPs provides for potential to support non-standard multi-constrained path finding (e.g., predictive usage).
  - d. Secure: Separate forwarding tables and QoS queues for management and control traffic to prevent interaction with user traffic.
  - e. Programmable: Automated provisioning and management of SR-TE LSPs per service instance.
- G. Planning for backbone capacity:
  - a. ESnet looks at their historical traffic growth from several different points of view to gauge Day 1 and future needs on each segment of their ring-of-rings design.
  - b. In planning the capacity ESnet is ensuring that there will be enough capacity at any intersection to cover the loss of any segment.
  - c. ESnet consulted with its major users on their anticipated growth.
  - d. ESnet now has their traffic growth evaluation process automated and plans to revisit it approximately every six months.
- H. ESnet6's types of forwarding services:
  - a. "Low Touch": Secure, programmable with flexible service instantiations.
  - b. "High Touch": Highly programmable data plane with flexible data path configurations. ESnet views this service "as a strategic capability that will initiate new conversations, facilitate research collaborations, and enable novel services."
  - c. "No Touch": A possible future service with opaque forwarding over a Point-to-Point wave service.
- I. Next steps and questions: ESnet is doing the upgrades to ESnet6 incrementally. As the new Open Line System is installed the legacy ESnet5 packet services will be swung over to it – "ESnet5.5".
  - a. Will IPv4 be supported natively or tunneled over an IPv6 native network? Or dual stack?  
 ESnet intends to use IPv6 for management and internal networks wherever possible – using IPv4 only where a device only supports IPv4. For customer and peer facing network ESnet plans for parity of IPv4 and IPv6 at least in the early phase of ESnet6. At some point ESnet may roll out features to IPv6 that may not be immediately supported in IPv4. They have no explicit, intentional notion of de-supporting IPv4 for at least the first few years of ESnet6.

## VII. Brief on EdgeNet – Rick McGeer (US Ignite)

- A. EdgeNet: A lightweight, scalable edge cloud.
- B. Goal of EdgeNet: A presence in every city in the world with a population over 100,000.
- C. First edge cloud PlanetLab (PL):
  - a. Peaked at 1353 nodes at 717 sites.
  - b. Over 150 publications in top 5 venues in 5 years.
  - c. PL very easy for experimenters to use. They get a container and can easily build a network of containers.

- d. Died due to costs of custom software plus the local costs – dollars and staff time.
- D. Takeaways from PlanetLab for EdgeNet:
  - a. CS Funding agencies do 3-5 years and no maintenance.
  - b. Hardware is heinous - incredibly hard to support remotely.
  - c. Use other peoples' software – don't do custom software unless it really is adding value.
  - d. "Be a nice guest" – be as cheap and easy as you can be for the local sites.
  - e. Live off the land. Send as little as possible by using local resources wherever possible. A researcher should be able to download into the local cloud and be able to join the network.
- E. EdgeNet is PlanetLab but with local virtual machines plus Docker & Kubernetes.

### VIII. Update on Internet2's Next Generation Network: Chris Wilkinson (Internet2)

- A. NextGen is a top-to-bottom redo for Internet2. Started discussions in 2016 and concluded with an "Expectations 2022" paper (see: [internet2.box.com/v/NGI2022](http://internet2.box.com/v/NGI2022)).
- B. Anticipated layer 1 (1-2 years):
  - a. Waves up to 800G.
  - b. Support for 3<sup>rd</sup> party transponders.
  - c. Multi-tenancy – sharing fiber & systems.
  - d. Flex spectrum 50+ GHz.
  - e. Lower fees – based on spectrum & segments.
- C. Anticipated layer 2/3:
  - a. 10, 100, nx100 and 400G connections.
  - b. Re-examine connector platforms - backhaul & "local" routers and the impact this can have on the fee structure.
  - c. On-demand layer 2 & 3 services including cloud connect.
  - d. Enhanced self-service.
- D. Enhanced tools for peering:
  - a. Cloud services with perfSONAR at the cloud edge.
  - b. Rapid deployment of low cost 10 and 100G dedicated ports at peering point for use by connectors.
  - c. Options for connectors to deploy virtual or physical routers at remote peering points.
  - d. Virtual machines available at peering points for use by connectors.
- E. Active efforts in past 12 months:
  - a. Cloud Connect: on demand layer 3 to the big 3 providers; over 50 campuses already in pilot service.
  - b. TR-CPS "cap" raised to 50G. The related exchange point parings are being upgraded to 100G.
  - c. NGI optical RFP released, received and being evaluated.
  - d. BAFO is targeted in a few weeks.
- F. Optical photonics investment in 2019:
  - a. This will support nationwide community wavelength sharing.

- b. It will support: >50 GHz channels, 200-800G transponders, inter-network optical connections and foreign waves. The open line system (OLS) is Internet2's focus.
  - c. With maintenance windows the goal is to upgrade without impacting performance.
  - d. Expecting to run their legacy transponders on the new OLS for 6-12 months.
  - e. They are "mindful" of SC coming in Nov and hope the new system will be able to assist.
- G. Packet services:
- a. A high level RFI was due April 15.
  - b. Responders were asked to envision a 2022 R&E network.
  - c. Community collaboration workshop to help shape an RFP.
  - d. RFP targeted for late summer, evaluation during the 3<sup>rd</sup> quarter with award in the 4<sup>th</sup>.
- H. Interconnection upgrades: Two large sites were migrated to new hardware with much higher port density and with much lower per port and maintenance costs. Recovered hardware was allocated to the other 5 peering sites.

### **Meetings of Interest 2019**

June 10-12	<a href="#">NANOG76</a> , Washington, DC
June 16-20	<a href="#">TNC19</a> , Tallinn, Estonia
July 20-26	<a href="#">IETF 105</a> , Montreal, Quebec, Canada
July 22-26	<a href="#">APAN48</a> , Kuala Lumpur, Malaysia
Sep 10-12	DREN Technical Interchange Meeting, Moffett Field, CA
Sep 17	GLIF Americas Workshop, San Diego, CA
Sep 17-18	<a href="#">Global Research Platform Workshop</a> , San Diego, CA
Sep 19-20	GLIF community/GNA meeting, San Diego, CA
Sep 24-25	<a href="#">National Research Platform</a> , Minneapolis, MN
Oct 28-30	<a href="#">NANOG77</a> , Austin, TX
Oct 30 – Nov 1	<a href="#">ARIN 44</a> , Austin, TX

### **Next JET meetings**

- May 21, 2019 12-2 p.m. ET, NCO/NITRD, Washington, DC
- June 18, 2019 12-2 p.m. ET, *n.b.*: This will be a virtual meeting