



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 16, 2024, 12:00 – 2:00 p.m. ET

This meeting was held virtually

Participants

Hans Addleman, Indiana University
Dale Carder, ESnet
Basil Decina, NRL
Ann Keane, NOAA
Michael Lambert, PSC/3ROX/ACCESS
Theo Lavis, MCNC

Paul Love, NCO/NITRD
Joe Mambretti, StarLight/MREN
Aruna Muppalla, NASA/GSFC
Kate Robinson, ESnet
Michael Sinatra, ESnet
David Sinn, Pacific Wave/PNWGP

Proceeding: This meeting was coordinated by Paul Love (NCO).

I. **Action Items:** (none pending)

II. **Review of the Minutes of the November 2023 meeting:** The final minutes were received – no corrections were received.

II. **IPv6 at SCinet: An Update & Lessons Learned – Hans Addleman and Kate Robinson**

A. What is SCinet?

- a. It's a pair of networks that support the SC (formerly Supercomputing) conference – one for “production” support of the meeting rooms, conference area Wi-Fi, exhibitors and the streaming of many sessions. The second is the research network supporting the Networked Research Exhibition (NRE) and SCinet's experimental network projects (XNET).
- b. It commonly takes a year to design, a month to build, a week to operate and one day to tear down and pack up.
- c. For 2023 there were 205 volunteers making up SCinet. They were broken into 14 teams covering everything from routing, help desk, NRE, XNET, DevOps, and architecture to finance and physical security.

B. SCinet by the numbers (with thanks to Hans)

- a. 6.71Tbps delivered to the Colorado Convention Center
 - i. 16 - 400Gbps connections.
 - ii. Some experimentation with 800G, but internal only.
- b. 33kW power draw.
- c. 205 SCinet Volunteers!
 - i. Volunteers from 9 Countries, 31 states, and 113 institutions!

- d. \$40.5 million in donated/loaned hardware, software, and services from 30 contributors.
 - e. 12+ miles of Fiber deployed.
 - i. 2178+ Fiber patches.
 - f. 415+ wireless access points.
- C. Some of the neat things SCinet implemented at SC23 (again thanks to Hans)
- a. IPv6 including transitional helpers such as DHCP option 108, NAT64 and DNS64.
 - b. Verizon 5G Wireless on the show floor.
 - c. 400G ZR+ and tapped the Internet2 commodity circuit using 400G ZR+.
 - d. NSO Orchestration and Automation with EVPN/VXLAN.
 - e. Migrated most management and commodity nets behind the Juniper firewall.
 - f. New custom written database for physical and configuration data - iNext.
 - g. Network security team tapped 8.04 Tb of traffic- LAN and WAN.
 - h. 2x400GE (800GE) optics between vendors.
 - i. Multi-vendor zero touch provisioning.
 - j. RPKI (Alerting, not dropping).
 - k. FABRIC in SCinet.
 - l. Automated black hole routing with FlowSpec via Gravwell and Corero.
 - m. InfluxDB and Grafana.
- D. IPv6 and SCinet at SC23
- a. The SCinet chair picks what will be the technical reach for the year – for SC23 it was IPv6.
 - b. About 20-25% of SCinet volunteers are from federal institutions. That, combined with the federal IPv6-only mandate (80% of all network devices before October 2025), made a good combination and reinforced the choice of IPv6.
 - c. SCinet is a great place to try new things and large changes as, in many ways, it's a green field network every year. (In early October 2023 SCinet was all in boxes.)
 - d. There was some discussion of this focus on IPv6 was still a useful exercise, still timely. IPv6 has been around for many years and the last block of IPv4 space was handed out in 2015. But IPv6-only is still complicated and challenging so very timely.
 - e. The goal was to see what worked and what didn't in the very condensed timeframe and to come out of it with good lessons learned and documentation.
 - f. Complicating things even more than the usual compressed time frame of assembling SCinet in the week+ in October, in that same period Denver had a metro fiber cut and the convention Center suffered a power outage. Finally, there were several significant software issues SCinet needed to work with vendors on – some IPv6 related, some not.
- E. Planning for SCinet at SC23
- a. SCinet holds three planning meetings each year to discuss architecture and design. At these it was decided on two areas for IPv6-only: the management and wireless networks.

- b. Where possible some ideas were tested virtually in a lab as all the loaned pieces for SCinet (hardware/software/etc.) arrive onsite in early Oct and aren't available beforehand.
 - c. The planning process identified these crucial building blocks for the two areas:
 - i. Management: NAT64 and DNS64, SLAAC and/or DHCPv6, and DHCP Option 108.
 - ii. Wireless: NTP, Syslog, SNMP and streaming telemetry, configuration backup, NSO/automation tools/configuration tools, and ZTP.
 - iii. NAT64 and DNS64 transition tools were the 2 most critical, what was tested throughout the year and what SCinet had operational as soon as they could in Oct.
- F. A critical piece: DHCP Option 108 for the wireless network
- a. This was one of the most interesting things SCinet explored for supporting IPv6-only in an environment where some devices would be IPv6-only, some able to operate with dual stack and some needing IPv4.
 - b. SCinet weighed a variety of options: SSIDs specific to IPv6-only, dual stack, EDUROAM dual stack, etc.
 - i. Considerations included allowing for users to self-select which SSID to use.
 - ii. Confusion if the selected SSID didn't work and how to back out.
 - iii. Trying to reduce the amount of Help Desk visits (plus the Help Desk doesn't open until the exhibits open, but the rest of the conference starts a couple of days before that).
 - iv. Wireless performance (5-10k wireless devices were anticipated).
 - c. SCinet elected to use a wireless SSID that utilized DHCP option 108.
 - i. One of the biggest challenges to deploying IPv6-only was the sheer number of wireless devices. Some were fine with IPv6-only, but many had IPv4 dependencies. Option 108 seemed to be just what was needed – DHCPv4 with option 108 should only turn off IPv4 if both sides agree to use IPv6-only - the devices would self-select IPv6-only. With Option 108 this single SSID could support both dual stack and IPv6-only reducing complexity for the user.
 - ii. DHCPv4 with Option 108 is another transitional tool. If a network is only supporting IPv6 then probably Option 108 wouldn't be configured. As a transitional tool it was really interesting to see how many devices would come up as IPv6-only at SC23.
 - iii. But not all OSs support Option 108: most mobile devices do, MAC-OS does, many Linux distributions do, but Windows 10 and 11 do not.
 - d. Another example of all OSs not being equal: SCinet chose to use SLAAC rather than DHCPv6 to hand out IPv6 addresses. This means that DNS is provided by router advertisements - which not all OSs are prepared for. In the end different clients were doing different things to get their addresses. Confusing for SCinet in debugging and helping users, but worked for almost all.

- G. The biggest technical issue on the wireless network involved multicast
 - a. By default router advertisements are sent by multicast. But there many issues with multicast on the wireless. With Option 108 IPv4 was being dropped before the IPv6 was usable. A route advertisement would take a minute or more to resolve. As more and more devices were connecting over the wireless things just got more backed up, severely impacting wireless performance. What should have been a quick reply became as bad as 100 seconds for reply.
 - b. Router advertisements be switch to unicast made a great improvement.
- H. Client issues reveled due to the use of Option 108
 - a. Centrally managed devices:
 - i. IPv6 stack removed.
 - ii. Firewalls had several problems such as filtering ICMP.
 - iii. VPNs going back to the home institution.
 - b. Older Android devices and non-Google (pixel) Androids, had limitations with IPv6 and certain networks.
- I. Lessons Learned (with thanks to Kate)
 - a. Communication is key! Within SCinet but also with the SC23 attendees about was being tried.
 - b. Change can be difficult (as we all know).
 - c. It required the effort of all teams to produce a working environment, including devops, network security, wireless, routing, help desk and communications.
 - d. So many things at the conference are dependent on the conference wireless.
 - e. SC23v6 SSID (using option 108) was an opt-in wireless SSID this year instead of the default SSID. This resulted in about 10% of the wireless users working under Option 108 for IPv6-only.
 - f. It's either too much [IPv6] or not enough [IPv6] depending on who you talk to
 - g. There is always the propensity to blame the "new" thing even if that's not the issue.
- J. SC24
 - a. Working to carry SC23's IPv6-only lessons forward.
 - b. Every SCinet has some new volunteers joining and some old ones stepping back so, if you'd like to join SCinet for 2024 please use the following link (SC24 will be in Atlanta):
 - <https://sc24.supercomputing.org/scinet/participate-scinet/>
 - c. If you'd like to volunteer with a focus on IPv6 please say so.
 - d. SC24 SCinet Chair, Angie Asmus, Colorado State University.
 - e. SC24 SCinet Technical Directory, Brenda Meade, Indiana University.

III. Women in IT Networking at SC – Kate Robinson

- A. Women in IT Networking at SC (WINS) started off as an NSF award in 2015 to enhance gender diversity in SCinet by providing the funding for to participate in SCinet who don't have the funding for the travel, lodging and meals to do so.

- B. It's open to US network engineers in government, academia or private industry and gives great opportunities to work on a wide variety of cutting-edge equipment alongside some of the best hardware and software engineers.
- C. I was an awardee in 2019 and it has been great for the trajectory of my career. The SCinet 24 chair and technical directory are also past awardees.
- D. While applications for SC24 closed the end of January, there will be the WINS program in 2025. *(Note: In 2025 WINS will also be open to international network engineers.)*

IV. Operational Security Round Table:

- A. SC23 (Kate Robinson): SCinet had extensive discussions with its Network Security (NS) team over possible issues if SCinet used SLAAC rather than DHCPv6 for IPv6 addresses. In the past NS had used the DHCP server to match the IP address with a MAC address of a bad actor, then track down the user via the wireless access points. In the end SCinet went ahead and used SLAAC. NS was able to track down users, it just wasn't as convenient as before.
- B. PSC/3ROX/ACCESS (Michael Lambert): PSC is locking down SSL certificates by using CAA in DNS to restrict who can generate them.

V. Network roundtable

- A. ESnet (Dale Carder, Michael Sinatra): ESnet has been focused on getting ready for the LHC data challenge. It's scheduled for mid to late February and is a dry run to test the readiness for the data rates required for the LHC's upcoming High Luminosity runs. ESnet has had discussions with end sites and RONS that will be involved. To exacerbate these preparations a trans-Atlantic cable had an outage on the first of January. A repair ship is on station and making repairs.
- B. MCNC (Theo Lavis): Nothing for today.
- C. NOAA/N-Wave (Ann Keane):
 - a. N-Wave's DC area DWDM ring is being upgraded with new nodes in Ashburn/Equinix and the NIST Gaithersburg campus. The latter will bring a lot more partners on at the Gaithersburg site. The ring's new capacity is 800G.
 - b. For the National Weather Service (NWS) N-Wave is working to bring the Strom Prediction Center in Norman, OK, and the National Hurricane Center in Miami, FL, on board. These are expected to be finished within 6 months.
 - c. N-Wave is also connecting NWS' Hawaiian forecast office onboard.
 - d. N-Wave has a backbone upgrade underway to Alaska along with a major build within AK to bring the forecast offices there onto N-Wave.
- D. NRL (Basil Decina): NRL is reconfiguring some its links as it returns equipment borrowed for demos at SC23.
- E. Pacific Wave (David Sinn):
 - a. Pacific Wave (PW) supported a number of the NRE demos at SC23 including a number at 400G. The latter leverage the new NA-REX 400G connections to Seattle and Los Angeles.
 - b. PW is in the middle of rolling out its new PTXs at PW's exchanges. These will support local 400G. Concurrently PW is working through upgrading its west coast

backbone so that inter-city 400G will also be possible. Intercity ETA is at least three months out.

- c. Route server participation continues to grow.
- d. Uptake of using the router servers is a little slow. Most likely a result of a mature exchange where most participants have already established their bilateral agreements.
- e. PW is planning to add perfSONAR (pS) nodes that will learn routes from the router servers. This will let PW be a pS test location.
- f. PW is investigating extending the exchange up to Fairbanks. It's looking positive at this point. The University of Alaska is highly likely to be a good partner for the expansion and host the exchange.

Question: Are you using BIRD for your router server or something else?

Answer: Something else – vMX which has just gone end of life (and complicated by HPE's probable acquisition of Juniper.) PW chose vMX as it easily provided a virtual router server for each VLAN in PW.

F. PSC/3ROX/ACCESS (Michael Lambert):

- a. PSC and 3ROX have completed an upgrade to their DWDM ring to more easily support 100G. This has reduced their Cisco 15454s to transponder shelves with new Ekinops gear to drive the ring with a capacity of up to 800G. The existing 10G lags remain in use while the Ekinops gets stable. The lags will then be migrated to 100Gs. PSC/3ROX has a few stories from this. If you're interested in Ekinops deployments please contact Michael directly.
- b. No updates for ACCESS this month.

VI. Exchange Points Round Table

A. NASA Ames/NGIX (Bobby Cates via email)

- a. DREN has moved regional peering to Equinix San Jose although they still support Department of Defense (DOD) components on campus from there via leased circuits.
- b. All NIPR, SIPR and DSN services are provided remotely via a circuit to a DOD facility.
- c. USGS has no local peering, although it intends to set up a path to Lumen's facility at 1380 Kifer, Sunnyvale, CA. USGS has taken over the three campus buildings it were leased. USGS' RF tower has been completed. They are just breaking ground on construction of a new lab on the campus.
- d. UC Berkeley (UCB) may break ground this year on its new campus at Moffett Field. UCB intends to connect to CENIC – bandwidth still to be determined. Construction won't be completed before 2025.
- e. NREN still has peering and transit in CENIC cage at Sunnyvale. They have local cross connect at that location to peer with ESnet. (It's believed that NASA JPL also has CENIC connectivity in Los Angeles and that 100G path has been provisioned between them.)
- f. NASA's WAN peers at Equinix San Jose (as well as Dallas, Chicago and Virginia). They have local cross connect to DREN in San Jose.

- B. PNWGP (David Sinn): PNWGP is in the middle of its PTX deployment. Initially it will be offering 400G to its transport providers. 400G to its participants will follow. The deployment is anticipated to take the rest of CY2024.
- C. StarLight (Joe Mambretti):
 - a. In mid-December StarLight (SL) received a pallet full of gear for a terabit FABRIC node. The node was assembled 2 January, configured, tested and is now up and running. The node will be connected to a similar implementation in McLean, VA, where the Joint Big Data Testbed (JBDT) with NRL and NASA/GSFC is located. The result will be a Tbps testbed for data intensive science.
 - b. SL is planning on prototyping Tbps flows between Chicago and the JBDT using 3rd Gen 400G DTNs at each end.
 - c. In San Diego, CA, in March:
 - i. SL is also getting ready to do demos at both the KNIT 8 workshop and the 5NRP meeting at UCSD.
 - ii. Other SL demos will be at OFC the following week in San Diego. They'll include 400G and 800G services. Also anticipated are quantum demos.
 - d. SL is organizing a panel about the SC23 NRE demos for Internet2's CommEX.
 - e. SL is part of a group working with the LHC networking community of packet marking experiments. At SC23 Yatish Kumar from ESnet was doing marking experiments. His results are expected in a few days.

Question: What is SL expecting with its quantum co-propagation experiments at this year's OFC?

Answer: At OFC last year SL measured 95% fidelity over about 2 miles on the show floor and about 91% on a 16 and a half run to a downtown POP. These used the O band for the quantum and C band for classical. (Later SL switched the bands using its metro Chicago testbed and increased the classical channel to 400G with no difference in results.)

Question: How was the 91% measured? Error rate with and without coexistence?

Answer: Yes. SL used a photon detector made by Quantum Opus a small company in Michigan. It uses cryocooled nanowires to do individual photon counts.

Meetings of Interest 2024

Jan 21-24	PTC'24 , Honolulu, HI
Jan 29 – Feb 2	APAN57 , Thailand
Jan 30-31	HIC , Kauai, HI
Feb 5-7	NANOG 90 , Charlotte, NC
Feb 19-22	SupercomputingAsia 2024 , Sydney, Australia
Mar 4-7	Internet2 Community Exchange , Chicago, IL
Mar 12-14	The Quilt Winter Meeting , virtual
Mar 16-22	IETF 119 , Brisbane, Australia
Mar 19-21	KNIT 8 , San Diego, CA
Mar 20-22	Fifth National Research Platform Workshop , San Diego, CA
Mar 24-28	OFC , San Diego, CA
Mar 25-27	CENIC 2024 , Monterey, VA

Apr 14-17 [ARIN 53](#), Bridgetown, Barbados
Apr 30 – May 2 [ESCC](#), Berkeley, CA
Jun 10-12 [NANOG 91](#), Kansas City, MO
Jun 10-14 [TNC24](#), Rennes, France
Jul 20-26 [IETF 120](#), Vancouver, B.C. Canada
Aug 26-20 [APAN58](#), Islamabad, Pakistan
Sep 16-19 [Fifth Global Research Platform Workshop](#) at [IEEE eScience](#), Osaka Japan
Sep 17-19 [The Quilt Fall Meeting](#), location/format TBA
Oct 21-23 [NANOG 92](#), Toronto, ON Canada
Oct 24-25 [ARIN 54](#), Toronto, ON Canada

Next JET meetings

Note: It is anticipated that most JET meetings will remain virtual for the foreseeable future

Mar 6, 2024 12-1:15 p.m. CT, hybrid
Note: This meeting is collocated with Internet2's Community Exchange in the Michigan B room of the Sheraton Grand Riverwalk Chicago, 301 E North Water St, Chicago, IL 60611

Apr 16, 2024 12-2 p.m. ET
May 21, 2024 12-2 p.m. ET