



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
November 16, 2021, 12:00-2:00 p.m. ET
This meeting was held virtually

Participants

Jeff Bartig, Internet2	Paul Love, NCO/NITRD
Jamin Becker, Dynamite Analytics	Joe Mambretti, StarLight/MREN
Nick Buraglio, ESnet	Linden Mercer, NRL
Rich Carlson, DOE/SC	Edward Moynihan, Indiana University
Basil Decina, NRL	Aruna Muppalla, NASA/GSFC
Bill Fink, NASA/GSFC	Dave Nevin, Oregon State University
Kevin Kranacs, NASA/GSFC – EOS	Adam Pumphrey, Dynamite Analytics
Padma Krishnaswamy, FCC	Glenn Ricart, US Ignite
Yatish Kumar, ESnet	Oleg Sinitsin, Dynamite Analytics
Michael Lambert, PSC/3ROX	Kevin Thompson, NSF

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

I. Action Items:

- Internet2 and ESnet updates on their respective new networks.

II. **Review of the Minutes of the October 2021 meeting:** The review was delayed and will be done via email due to the draft minutes not being available before this meeting.

Note: They were circulated and have now been posted to the JET's web site.

III. PCAP Analytics using PacketTotal – Jamin Becker & Adam Pumphrey

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

<https://www.nitrd.gov/nitrdgroups/images/2021/JET-Jamin-Becker-11162021.pdf>

- A. Dynamite Analytics was founded in 2015 in Atlanta, GA. It's received several Federal R&E awards primarily through the SBIR program. It's partnered with the Department of Energy (DOE) and the Department of Defense (DOE). Their network analytics programs are in both the open source and commercial product spaces.
- B. PacketTotal (PT) was acquired recently by Dynamite Analytics. PT has been in operation since 2016 as a free community service. PT is a tool to process PCAP files and give a security researcher quick insights into a captured PCAP file that may contain malicious packets. Over 100k PCAP files have been shared (uploaded) for analysis. (PCAP is a standard way to capture network traffic that can then be looked at offline by a variety of tools such as Wireshark and PT.)

- C. PT was designed with three goals:
 - a. Simplify the process of PCAP analysis.
 - b. Provide a platform for sharing network traffic samples.
 - c. Prioritize the importance of search.

It was aimed at aiding the security researcher who doesn't need the verbose information provided by a tool such as Wireshark. It gives the security research a mechanism to find PCAPs that have the issue the researcher is studying by including a search language designed with PCAP files in mind.
- D. Demos were given of:
 - a. Analyzing a PCAP file
 - b. Locating captures via search
 - c. Locating captures via similar search
- E. In using PT most new PCAP files will take several seconds or more to be processed. PT uses two network analysis frameworks – Zeke and Suricata. Zeke provides a log similar to NetFlow but richer. Suricata highlights known malicious activity. For a stateful protocols, such as TCP, PT displays the flow as a single entry vs individual transactions. Any application layer protocol that PT has an engine for will be broken out as a log that can be displayed. It is possible to download things of interest for additional, local analysis. But, caveat emptor, these are files with malicious content and some may still be active.
- F. In using PT be careful of what you upload as it will be publicly visible since all uploads are shared:
 - a. Avoid uploading PCAPS of home or work networks.
 - b. Avoid capturing traffic containing information about your environment.
 - c. Consider using capture filters, only capture the traffic you wish to analyze/share.
 - d. Consider capturing PCAPs from a sandbox environment (this is strongly recommended).
- G. PT's search capabilities are based on Elasticsearch. Simple items can be searched for – IP address, text, domain names, etc. But more complex searches can be done using Boolean operators or wild cards for partial matches. A search can also be done looking for a certain Suricata signature plus another item such as a text match.
- H. An additional tool is the ability to upload a PCAP and then have PT search for similar PCAPs in its database.
- I. PT has a malware archive. It contains a description of the malware along with five PCAPs that contain examples of each entry.
- J. Question: Does PT support IPv6? Answer: Yes.
Question: Does PT run in an IPv6-only environment. Answer: It should but that's theoretical as Dynamite Analytics is unaware of any instances.
- K. Related tools from Dynamite Analytics:
 - a. PCAP Anonymizer: Helps promote sharing PCAPs. Research for this tool was supported by a DOE Cyber Phase 1 R&D Award. Among other features it allows payloads to be overwritten with random data. When the Phase 1 project is completed (it's currently in beta) the tool will be available for free download

from the dynamit.ai website. The beta is closed - contact Dynamite Analytics if you're interested in joining.

- b. Dynamite Agent: Network traffic sensor for AWS users built on Zeke and Suricata. A commercial product to look at live streams. Easily integrated into an existing set of security tools. It is available on the AWS Marketplace. A starter edition is available for free. Larger instances require a fee.
- c. Dynamite NSM: An open source Network Security Monitor. Full network monitoring stack plus a SDK. Simple Zeek and Suricata management. Full data ingest pipeline. Available on GitHub.

IV. SC21 NRE Brief – Joe Mambretti, Basil Decina, Linden Mercer and Bill Fink

Note: The slides for these talks are posted on the JET's web page at:

<https://www.nitrd.gov/nitrdgroups/images/2021/JET-Joe-Mambretti-11162021.pdf>

<https://www.nitrd.gov/nitrdgroups/images/2021/JET-Decina-Mercer-11162021.pdf>

An overall note regarding these demos, for a variety of reasons work was able to start in the summer at the new, permanent site in the MAX suite in McLean, VA, and in StarLight. This greatly enhanced what was able to be brought together for SC21.

- A. SC21 Advanced Networking For Data Intensive Science Experiments and Demonstrations (Joe Mambretti, StarLight)
 - a. StarLight (SL) is working with collaborators around the world to build the Global Research Platform (GRP) – basically a worldwide DMZ. It's focused on the needs of data intensive science, the needed capacity and the programmability of that capacity.
 - b. Part of what is being showcased at SC21 are 400G and 600G channels (just as 100G WAN circuits was first at SC08).
 - c. A core component is the StarLight Exchange with over 80 100G circuits terminating there. With funding from the National Science Foundation the exchange is a SDX used to support various test beds and new technologies. Using SDN "slices" can be provided for multiple virtual environments including the NSF's IRNC SDX and the network research GENI SDX.
 - d. Currently SL is supporting about two dozen test beds. Working with the SCinet WAN team to build a national testbed with multiple 100G circuits from coast to coast and to the show floor via StarLight. There are 2x400G circuits from SL to McLean, VA, supporting demonstrations by NRL and NASA GSFC. From SL to the show floor there is 1.2T of capacity.
 - e. A big building block of what is being shown is the AutoGOLE fabric of open R&E exchanges located around the world. These currently do dynamic layer 2 provisioning. Work is under way to add layer 3.
 - f. Another major portion are demonstrations of DTN-as-a-Service linking DTNs worldwide.
 - g. Other demonstration support by SL include (see slide deck for more details):
 - i. SC21 NRE 002 Network Optimized Transport for Experimental Data (NOTED) AI/ML Driven WAN Network Orchestration
 - ii. SC21 NRE 004 Global Research Platform (GRP)

- iii. SC21 NRE 005 IRNC SDX 1.2 Tbps WAN Services
 - iv. SC21 NRE 006 IRNC SDX E2E 400 Gbps WAN Services
 - v. SC21 NRE 007 IRNC SDX International Testbed Integration
 - vi. SC21 NRE 008 StarLight SDX for Petascale Science
 - vii. SC21 NRE 009 DTN as a Service
 - viii. SC21 NRE 010 P4 Integration With Kubernetes
 - ix. SC21 NRE 011 PetaTrans Services Based on NVMe Over Fabric
 - x. SC21 NASA Goddard Space Flight Center Demonstrations
 - xi. SC21 NRL Demonstrations
 - xii. SC21 PRP/NRP Demonstrations
 - xiii. SC21 Open Science Grid Demonstrations
 - xiv. SC21 N DISE Demonstrations
 - xv. SAGE3: Smart Amplified Group Environment.
- h. Other efforts being undertaken at SL include:
- i. The integration of Chameleon and FABRIC. An initial demonstration will be shown at SC if possible
 - ii. Supporting the Data Mover Challenge at SupercomputingAsia 2022
 - iii. Compostable Platform-as-a-Service
 - iv. Quantum computing test beds.
- B. Resilient Distributed Processing (Basil Decina and Linden Mercer, NRL)
- a. For DOD one of the big interests in high performance computing and networking is the need to support hundreds of thousands of human entities involved along with the related computers and other devices. This quickly leads to millions of things to interconnect scattered around the world and move as they need to. These tend to have relative low data rates 10s-100s of Mbps. At the other end there are DOD's supercomputer centers which don't move but have data rates of expected of supercomputers. The challenge is connecting an entity on the field to a supercomputing resource. Some intermediate nodes – planes, ships, etc. - are possible with some control of their locations to facilitate the interconnection and provide dynamic distributed processing. A lot of the current effort focuses on how to have high efficient bandwidth to the low bandwidth users. Current effort also focuses on being able to dynamically reconfigure the intervening networks to take advantages of low latency paths, etc. The goal is faster and better decision cycles – improved OODA (Observe, Orient, Decide, Act) loops – through resilient distributed processing and immediate global data access.
 - b. To demonstrate these efforts NRL is showing dynamic arrangement and re-arrangement of widely distributed processing of large volumes of data across a set of compute and network resources organized in response to resource availability and changing application demands. A real-time video processing pipeline will be demonstrated from SC21 to the NRL assets in Washington, DC, McLean, VA, Chicago, IL, Berkeley, CA and back to SC21. High volume bulk data will be transferred concurrently across the same data paths. A software-controlled network will be assembled using a number of switches and multiple SCinet 100G/400G connections. This will show rapid deployment and

redeployment, real-time monitoring and QoS management of these application data flows with very different network demands. Technologies leveraged include SDN, RDMA, RoCE, NVMe, GPU acceleration and others.

NRL has two major thrusts for its SC21 demo/tests: Mission Oriented Reconfigurable Networking (MORN) and rapid terabyte data movement. It has five sites (NRL/Washington, DC; McLean, VA; StarLight/Chicago, IL; NERSC/Berkeley CA; SC21/St. Louis, MO) in the 100G network monitored and controlled by MORN. Four of those sites have 400G wide area network connections supporting fast delivery of critical data sets and massive distributed computing problems.

The specific goals are:

- i. Fast fault (near real time) detection and location using an active probe.
 - ii. Dynamic shifting of processing and network resources from one location/path/system to another (in response to demand and availability).
 - iii. Leverage improved (restored) RDMA/distance performance for timely Terabyte bulk data transfers (goal < 1 min TB transfer on 400G network).
 - iv. Network data flows are protected by IP and Ethernet Traffic Flow Security.
- c. Besides many 100G and some 400G WAN connections some servers have 200G interfaces.
 - d. As part of a Mission Oriented Reconfigurable Networking (MORN) system you have a mission's needs and priorities. These are the inputs to control the mission's data. A MORN isn't a best effort network, it's prioritized effort network. A MORN provides its resilient services by using rapid fault detection and location and then prioritizing the restoration of services using standard industry traffic engineering and SDN techniques.
 - e. The fault detection and locations are done by:
 - i. Path knowledge
 - ii. Limited by time synchronization, variations in network delay and accuracy of the detailed knowledge of the path.
 - f. Restoration also requires determination of the fault being in the network or in NRL equipment.
 - g. MORN orchestration then uses the fault detection information and its knowledge of the available networking resources to reroute/reprioritize to meet the various mission needs.
 - h. The goal of moving 1 TB of data in under a minute is driven by DOD's need to move the very detailed imagery. It will also be of value in 3D additive manufacturing and medical imaging. A single 100G can't achieve it, nor can 2x100G currently.
 - i. Last year using RDMA over Ethernet on 100G in a lab environment NRL achieved full data rate out to about 12k km. This year at SC one data point has been measure (at the time of the JET meeting) over 400G circuits. On 400G circuits, using devices with 200G interfaces 146 Gbps in a single stream was measured

over 5,842 miles (~9,400 km). For next year the hope is for the next generation of hardware – PCIe Gen 5 with 400G interfaces.

- j. Traffic Flow Security (TFS) has historically used bulk encryption. Network encryption is now replacing link encryption for wireline systems. Work is underway in the IEEE and the IETF to update the standards:
 - i. IEEE MACsec 802.1 Aedk - Publication approval process underway
<https://1.ieee802.org/security/802-1aedk/>
 - ii. IETF IPsecME WG TFS - Core protocol accepted, working through process. Publication approval process requested
<https://tools.ietf.org/html/draft-ietf-ipsecme-iptfs-03>
- C. GSFC demos (Bill Fink)
 - a. GSFC built two Supermicro AMD Epic servers each with four 200G interfaces. The original plan was to have one in McLean, VA, and the other on the show floor at SC. Due to COVID travel restrictions both were installed at McLean. Using the infrastructure provided by ESnet and Internet2 there were 2x400G connections to SL with 500G into the booth at SC. There are also 3x100G into the booth.
 - b. The plan is do disk-to-disk from one server in McLean to SC and then back to the other server in McLean. Each server has 16 NVMe drives. The transfers will be NVMe over Fabrics (NVMe-oF) over TCP. The goal is achieve at least an aggregate of 400 Gbps disk to disk.

V. 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.

The related, live map is at: <https://www.globalresearchmap.org/>

- A. Many thanks to Internet2, the Indiana University (IU) GRNOC, and the community in their continued support of the project.
- B. The IU GRNOC has increased resources on the virtual machine hosting the Internet2 time series database so that it can continue to take streaming metrics from the new I2 NGI network and also accommodate the I2 Community metrics.
- C. Work continues to integrate LEARN. The LEARN topology is up with partial live data available via the map:
https://globalresearchmap.org/?map_group=US%20Regional%20R%26E%20Networks&network_name=LEARN
To highlight LEARN on North America, click on the link and then hide the other regional networks.
- D. Work continues to integrate MERIT. A couple of endpoints are now on the map with live data. Another higher priority MERIT project is competing for people's cycles at the moment. The project is synergistic with the community collection, though, and should show results in a few weeks for both.
https://globalresearchmap.org/?map_group=US%20Regional%20R%26E%20Networks&network_name=MERIT

- E. Work is progressing to show the CAAREN exchange point. The results will mimic what displays on the current Quilt map
<https://www.thequilt.net/wp-content/uploads/QUILT-4-3-18-11x6.pdf>.
Two of its circuits are now live on the IU GRNOC map.
https://globalresearchmap.org/?map_group=US%20Regional%20R%26E%20Networks&network_name=CAAREN%20Geographic
NOTE: Work still continues on scaling and you may have to zoom in a bit if coming from country scale.
- F. Members of the team in Brazil have reached out in the past week. They are assembling SNMP information for at least a portion of the network there. They are hoping to send over the results within a few weeks.
- G. SunCorridor has also agreed to share their metrics with the community efforts as part of their work the EPOC project.
- H. In October, the community project did a presentation to the LSN broadband team and fielded questions regarding how the project was currently taking data and how it was working with the different community members.
- I. As an aside, the ESnet NetBeam API will be transitioning to ESnet's new Stardust API. Community tools which rely on that API will be transitioning as well. Andy Lake is reaching out to groups on this topic.
- J. In recent weeks there have also been discussions on how the EPOC project, NetSage project and community project efforts can collaborate and respect people's data. Data ownership is very important and the projects want to make no assumptions. For entities willing to share across projects, there are great synergies for operations and for research.
- K. Again, a lot of thanks to those involved in the efforts.
- L. 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_PNVpeT42nAxa3bW4QostMxcHTXkWSPbZOplFwE/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_hlnrjL8WpTbHmQSW112zzvKPBp6fx8Gw/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:Uhl, George D. <george.d.uhl@nasa.gov>)

VI. Operational network security roundtable No updates were received.

VII. Network roundtable

- A. Internet2 (Jeff Bartig): Internet2's (I2) focus the last several months has been migrating services onto its NGI network which is nearing completion. I2 has started decommissioning its old Juniper MX network and the associated 100G backbone. It's anticipated the decommissioning will take two to three months.
- B. International Networks – Indiana University (Ed Moynihan):
 - a. Good news on the fiber cuts mentioned last month:
 - i. NEA³R's 100G circuit that was cut is repaired and is now operational. The cut was not too far off the coast of Ireland and caused by a fishing boat's anchor. Repair was delayed due to bad weather. ANA worked well balancing the trans-Atlantic traffic over the consortium's remaining circuits. With 900 Gbps total in ANA with its agreement for mutual backup there isn't a cost justification for carrier provided protection circuits.
 - ii. TransPAC's new 100G Guam<>Singapore is anticipated to be repaired by 24 November with traffic flowing after Thanksgiving. The cut is off the coast of Indonesia
 - b. In the Pacific region International Networks – Indiana University is working with others on the Asia Pacific Oceania Network Collaboration (APONet). This will be

the same sort of consortium as ANA. It's expected to be in use as the last of the member's circuits come online.

- c. The UbuntuNet Alliance meeting last week discussed its new investments in eastern Africa.
- C. NASA EOS (Kevin Kranacs): No update today.
- D. NRL (Basil Decina): No additional updates.
- E. Pacific Wave (Jonah Keough – via email):
 - a. Pacific Wave is supporting multiple demos at SC21, with 600Gbps of capacity and various activities, including dynamically provisioned circuits directly from UCSD to São Paul.
 - b. The optical system upgrade on the west coast continues to progress.
- F. Oregon State University (Dave Nevin): Nothing to report today.
- G. PSC/3ROC/XSEDE (Michael Lambert): The recently turned up 100G between PSC's two machine rooms (to provide 100G to PSC's Bridges-2 supercomputer) has some continuing issues. It's built with Ekinops boxes running alien waves over a pair of Cisco 15154s.
- H. US Ignite (Glenn Ricart):
 - a. US Ignite is continuing to work seven communities on connecting the unconnected to make them neutral host carriage facilities on which multiple services can be operated. US Ignite is a big believer in neutral host because it allows for things like adding public safety and other high priority services on the same kind of network and then being able to use priority to adjust network traffic.
 - b. US Ignite is also working with number of innovation districts where they're putting in new fiber and/or new wireless to then becoming neutral host carriage facilities as well.
 - c. There are also a host of issues with neutral host networks that we're grappling with at the moment.

VIII. Exchange Points Round Table

- A. MAN LAN/WIX (Jeff Bartig): MAN LAN was migrated to its new cabinet last week.

IX. Update from the LSN's Annual Planning Meeting including tasking to the JET for CY2022 – Rich Carlson

At the LSN's Annual Planning Meeting last month there were no questions on the JET's proposed tasking for CY2022 so they are adopted as submitted.

X. Discussion

- A. Kevin Kranacs: What happened to the GLIFF maps of worldwide R&E circuits?
Ed Moynihan: The GLIFF merged with GNA to form GNA-G which has a working group on mapping. Ed is unaware of an updated map.
Note: A subsequent email from Joe Mambretti confirmed that GNA-G's mapping WG hasn't produced an update.

Meetings of Interest 2021-2022

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

Nov 14-19 [SC21](#), St. Louis, MO, hybrid meeting
Jan 16-19, 2022 [PTC'22](#), Honolulu, HI
Jan, 26-27 [Hawaiian Intranet Consortium](#), virtual
Feb 9-10 [The Quilt Winter Meeting](#), virtual
Mar 1-3 [SupercomputingAsia 2022](#), Singapore, hybrid
Mar 7-11 [APAN53](#), Bangladesh, virtual
Mar 8-10 [The Quilt Winter Meeting](#), virtual
Mar 19-25 [IETF 113](#), Vienna, Austria
Apr 24-27 [ARIN 49](#), Nashville, TN
Jun 6-8 [NANOG 84](#), Montréal, Québec, Canada
Jun 13-17 [TNC22](#), Trieste, Italy
Jul 23-29 [IETF 114](#), Philadelphia, PA

Next JET meetings

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

Dec 21, 2021, 12-2 p.m. ET *n.b. Will be held only if needed*

Jan 18, 2022, 12-2 p.m. ET

Feb 15, 2022, 12-2 p.m. ET