



in this edition:

From the N-Wave Program Manager - P.	1
N-Wave Map - P.	3
Alaska Update - P.	4
Enterprise Wireless - P.	5
IJPS JUENO - P.	7
TICAP Update - P.	9
Cloud Update - P.	10
Network Changes & New Participants - P.	14
Multicast Changes - P.	15
Network Performance Metrics - P.	16
100G Upgrades & Phased Approach - P.	16
N-Wave NOC & ServiceNow - P.	17
New Staff - P.	18
N-Wave Partnerships - P.	19
Contact Information - P.	20



NOAA N-Wave

The N-Wave Enterprise Network Services Branch, under the NOAA Office of the Chief Information Officer, Service Delivery Division, supports both operations and research, enabling NOAA's mission of science, service and stewardship through highly available, secure, high-speed network transport and services.

Mission

N-Wave is committed to providing innovative networking capabilities with integrity, transparency, and flexibility, to enable NOAA's missions through the implementation of:

- Quality, advanced high-speed connectivity both internally and externally to NOAA
- Portfolio of secure, flexible, available, high-bandwidth network services
- Retention and recruitment of exceptional operations and engineering staff.

Our Vision

To provide reliable, secure, and sustainable enterprise network services for NOAA, which enables NOAA's mission of science, service, and stewardship.

From the N-Wave Program Manager



Robert Sears

This past summer, I was very humbled when invited to provide a keynote at the Defense Research and Engineering Network (DREN) Technical Interchange Meeting (TIM) for members of the DREN community. The invitation was to share and provide insight into the characteristics of N-Wave, as another high-performance network. With a long-standing DREN/N-Wave relationship, accepting the invite was the easy part, but speaking to this highly competent technical community was something entirely different. As my technical toolbox has long been retired and traded in for contracts, budgets and those other *highly coveted* program management duties, I was steered along a different approach to address our colleagues.

The story here is more about the process of drafting the presentation rather than the presentation itself, although the DREN TIM was very successful and will result in both continued and future collaboration. What I needed to do was place a critical eye on the N-Wave program and take inventory of our successes and, perhaps more importantly, our challenges. The N-Wave program has been in an eight-year-long "all ahead full" mode. With roots in NOAA research and development high performance computing, N-Wave has branched out to multiple operational roles, including customized Line Office private national and international transport, enterprise wireless, large campus local area networking and NOAA-wide internet access via Trusted Internet Connections. Taking this programmatic inventory resulted in four key observations and lessons learned that allowed me to provide a less technical yet more fundamental view into N-Wave high-performance network operations.

Strong scientific, research and education (SR&E) network community partnerships are foundational for innovative growth. The “Have it your way” network model doesn’t scale. More bandwidth is not a panacea. Customer engagement is imperative for strategic planning.

N-Wave’s SR&E partnerships (see page 19) are foundational to the service portfolio. Additionally, this community shares in the same scientific mission challenges as NOAA, experiencing exponential data generation and dissemination coupled with regionally diverse IT assets, all while balancing finite resources. N-Wave’s continued partnerships leverage the depth and breadth of this community that is collectively working together to meet these challenges.

The other three observations mentioned above require N-Wave to make targeted and strategic changes, as they are all intertwined and have apexed over the last two years. Prior to supporting NOAA-wide services, N-Wave had what was, in retrospect, a way too flexible customer connectivity model. As transport services were mainly private intra Line Office or intra program, there were very few constraints on the ways a customer would technically connect to N-Wave. The single goal was customer satisfaction.

Fast forward to the present day when N-Wave supports NOAA-wide access to enterprise-level resources – TICAPs, data centers, campuses, cloud services and more – standardized connectivity models had to be established, restricting the more a la carte approach. Our customers surely shared in some of the growing pains that accompanied this transition. To add to these growing pains, the limited visibility into downstream NOAA networks and new IT systems (e.g., data transfer nodes, cluster compute) that all demand higher bandwidth from N-Wave requires a more strategic technical and administrative approach than just incrementally increasing bandwidth.

Pulling this all together, customer engagement is the key to meeting and mitigating these varied challenges. The customer engagement strategy goes beyond interaction with the direct customer base to include networks and programs that are behind those N-Wave customer connections. This requires dedicated engagement from technical staff, program managers and science leaders across Line Offices to gain visibility into network needs and initiate proactive changes.

At the September 2018 NOAA Networking Committee (NNC) Meeting, invited guest speaker Jason Zurawski of the Energy Sciences Network (ESnet) provided a valuable tutorial on how to conduct science engagement with directed results that not only uncover areas where network infrastructures may be bottlenecks or pain points, but also on how to improve the overall data transfer process. In FY19, N-Wave will lead a similar engagement strategy as an augmentation to the NNC to gain wider visibility into NOAA-wide network requirements.

Another strategy to better the overall customer experience is to seek more directed end user technical exchanges. In July of FY18, Mark Mohs, IT lead for the National Centers for Coastal Ocean Service, and Bill McMullen, CTO for the National Ocean Service, joined N-Wave and GlobalNOC engineers in Boulder, CO, to discuss multiple areas of networking services, from campus LAN infrastructure to enterprise firewall services. This creates a more integrated partnership – not just a service provider to customer relationship, but an experience where all parties have ownership in successful IT service delivery. These strategies will require N-Wave to supplement staff across multiple disciplines to field the growing demand for N-Wave services.

It should also go without saying that one of N-Wave’s well-standing strategies is based in our partnership with the Strategic Sourcing Acquisition Division. Every N-Wave service – from the nationally spanning backbone to a single wireless access point in the hanger of the Lakeland Aircraft Operations Center – begins with acquisitions. In FY18, the SSAD team processed over 35 task orders against existing N-Wave IDIQ awards and managed maintenance contracts for over \$7 million in hardware assets in conjunction with new hardware orders, new SR&E provider contracts and strategic planning for FY19. The SSAD team of Eboni Luck (N-Wave Contracting Officer), Ronette Pratt (N-Wave Contract Specialist), McKenzie Hunter (Branch Chief, SDD HPC) and Olivia Bradley (Director, SSAD) are all part of the N-Wave family and valued as an extension of our national team. SSAD, the expanding N-Wave engineering team and our SR&E partners all make me very proud to be part of such a NOAA mission-enabling program!



Pictured left to right: Olivia Bradley, Eboni Luck, Ronette Pratt and McKenzy Hunter of the Strategic Sourcing Acquisition Division.



Alaska Update - N-Wave establishes connectivity in northernmost U.S. city

N-Wave engineers were recently in Utqiagvik (Barrow), AK, to install hardware at the National Weather Service's Weather Service Office and the Office of Oceanic and Atmospheric Research Global Monitoring Division Barrow Observatory to support those sites, along with NESDIS. NESDIS currently has five- and four-meter POES satellite antennas. Due to lack of bandwidth, these antennas have been used only for JASON mission command and control. Previously no data could be received from the satellites due to a slow, high-latency satellite link from Utqiagvik to Fairbanks. Now with higher capacity and much lower latency (from 550 to 22 milliseconds), these limitations are removed.

Utqiagvik is connected to Fairbanks via Quintillion fiber, which is subsea from Utqiagvik to Prudhoe Bay/Deadhorse and then terrestrial to Fairbanks (see page 5 in the spring 2018 issue for more details). In Utqiagvik, the sites are interconnected with a combination of fiber and microwave links.

The circuit from Utqiagvik to Fairbanks currently terminates on N-Wave's aggregation router located at the Fairbanks Command and Data Acquisition Station (FCDAS) in Gilmore Creek, AK. The existing Gilmore Creek to Fairbanks 100 megabits per second



Less than a mile outside of the Barrow Observatory, a sign reminds passers-by what it means to be in the northernmost city in the United States.

connection is scheduled to be upgraded to 1 gigabit per second by the end of November 2018. A second, diverse 1 gbps path between Gilmore Creek and Fairbanks is on order. From Fairbanks, the traffic from Utqiagvik (as well as from FCDAS) goes to the N-Wave core node in Seattle, WA, via an Alaskan terrestrial ring that in turn connects to a subsea ring to the Northwest.



Construction is complete on the first phase of the Quintillion subsea and terrestrial fiber optic network, which extends over Alaska's North Slope and inland to Fairbanks. Map courtesy of Quintillion.

While the N-Wave team was in Utqiagvik, they worked with local NOAA Office of Atmospheric Research staff to find alternative connectivity to the new building that will replace the existing [Earth Systems Research Laboratory Global Monitoring Division Barrow Observatory](#), which will be built over the next two years. This building is located on 100 acres of NOAA-owned property adjacent to the [Barrow Environmental Observatory](#). N-Wave is exploring multiple options for fiber optic cable connectivity at this strategic NOAA resource to the Quintillion subsea cable when the new facility opens.

NOAA's new Enterprise Wireless Service is deployed at first sites

N-Wave engineers recently designed and deployed an enterprise-level wireless service with NOAA-wide capability. The service is now available at three sites: Lakeland, FL, Fairmont, WV, and Charleston, SC. These first deployments follow a year-long process of designing and testing the centrally-managed service that is capable of scaling to meet wireless needs across all of NOAA, regardless of geographic location.

The service equips NOAA employees and their guests to work wirelessly across their facilities and at other customer sites without the need for any user or IT intervention. Once an employee completes the device onboarding process, no additional setup or changes are required for three years. Employees can also sponsor guest access for individuals and groups using a simple process that reduces the IT overhead traditionally required to set it up. This provides enhanced collaboration opportunities among NOAA employees and sponsored guests, and it frees up time for IT staff to work on other mission-related activities.

The history of NOAA wireless

This N-Wave service is the first enterprise-level wireless system to be deployed across NOAA. In the past, end sites would either set up their own wireless system or not have wireless services at all. Some sites chose not to implement wireless due to the number and complexity of required security controls. Other sites battled with having multiple wireless systems at a single campus, which often interfered with one another. These methods led to sometimes poor user experiences and high reliance on hard-wired connectivity across NOAA.

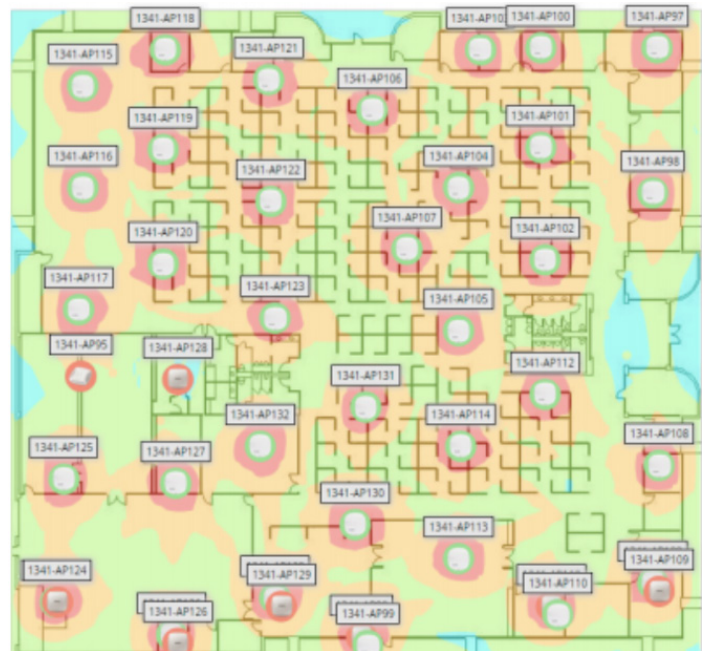
Designing for density before coverage

The N-Wave team has the expertise to provide a full wireless site design, including wireless AP layouts, to supply seamless connectivity across a given campus. These designs provide the ability for higher bandwidth customer connections. At the foundation of this deliverable is the team's design philosophy: Design by density first, and then by coverage.

Traditional designs aimed to install the minimum number of APs needed to cover all areas with WiFi

signals. Over time, this was found to be a flawed approach. Designing for coverage first and foremost left service providers scrambling during conferences and other events that created a spike in WiFi-enabled devices in a given space. The APs would become saturated and unable to service everyone.

Alternately, designing for high-density areas first – conference spaces, large meeting rooms, etc. – and then designing for coverage throughout the rest of the facility ensures continuity of service. Taking this into account, the density adequately covers for such occasions and then extends coverage to the remaining building areas that do not have high-density use.



This example AP layout shows site coverage.

A redundant architecture with east and west zones

The wireless system design also includes built-in, enhanced redundancy features. Enterprise hardware for the service was installed at the N-Wave points-of-presence (POPs) in Denver, CO, and College Park, MD. This strategy enables the wireless service to be broken up into two service areas: US-east and US-west. NOAA wireless customer end sites connect to the nearest zone, which reduces latency and improves the service performance. If there are issues in one

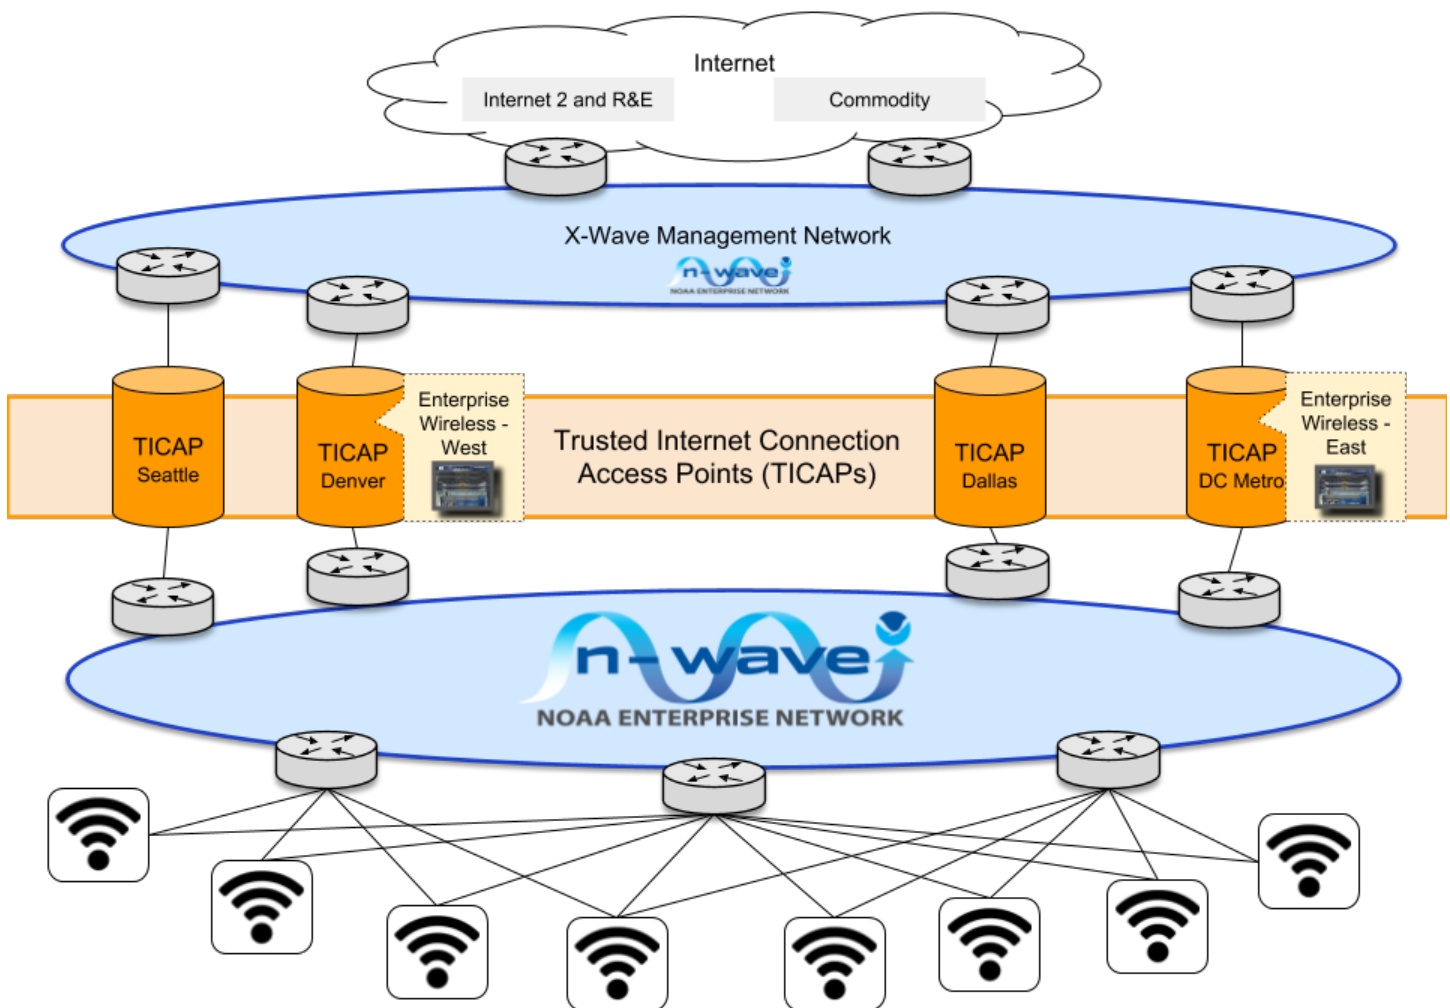
zone, all wireless customers and connectivity will fail over to the other zone, keeping wireless service available for everyone.

Security and cost-efficiency

Due to FISMA boundaries and security requirements, all wireless user traffic can only route to and from the external internet. This means wireless access to NOAA's internal networks or systems is prohibited. To get back into their internal NOAA networks when using a wireless connection, users must leverage their Line Office virtual private networks (VPN). This was a big driver in determining the locations of the enterprise infrastructure. With all traffic being internet-bound, the logical choice was to terminate

it at NOAA Trusted Internet Connection sites, also known as the TICAPs.

This new design also took into account the need to minimize the cost to NOAA end sites, while providing the same or better service than if a traditional wireless solution were deployed. To that end, customer sites that are already connected to N-Wave must purchase only the APs and the switches required to connect the APs. This is a significant advantage over traditional designs, which required redundant wireless LAN controllers, authentication servers, APs, switches and monitoring tools. This new design reduces the cost for wireless services for NOAA programs, enabling them to put those savings back into their missions.



Redundant infrastructure at the NOAA TICAPs in Denver, CO, and College Park, MD, provides enterprise wireless services to east and west zones, with the ability to fail over in the event of an outage at one of the TICAP sites.

Security features

The N-Wave wireless service also delivers a long list of security features to its users. Centralized authentication, EAP-TLS certificate-based protection, rogue AP identification and mitigation, certificate revocation and blacklisting features help protect users and the system as a whole.

- **Centralized authentication:** To configure devices to connect to the NOAA Secure SSID, users first authenticate to NOAA's Identity, Credential, and Access Management (ICAM) Single Sign On (SSO) to validate they are NOAA employees. NOAA is in the process of migrating all of its systems to ICAM, and the use of ICAM for enterprise wireless authentication is a significant step forward for this effort.
- **EAP-TLS certificate-based security:** Once authenticated, users must install service profiles and security certificates to join the NOAA Secure SSID. The enterprise wireless service uses EAP-TLS, which provides 802.1X radio frequency encryption to help secure the RF wireless connection. The EAP-TLS certificate remains valid for up to three years. This is another significant improvement over traditional systems that required users to sign in everyday.
- **Rogue AP identification and mitigation:** Rogue, unauthorized wireless access points cause interference that results in connection drops and poor performance. The enterprise wireless service is capable of detecting and locating rogue APs. If a rogue AP causes problems and does not shut down when requested, the service can employ mitigation techniques like tarpitting, which can be used

to redirect devices away from the rogue APs, toward authorized APs.

- **Certificate revocation and blacklisting:** When wireless use is in violation of DOC and NOAA policies and procedures, the enterprise service has the capability to revoke a user's certificate and blacklist the user from the network. The former prevents a user from joining the wireless network, and the latter will not allow the user to even connect to the wireless controller to get a wireless signal.

On the horizon

As N-Wave rolls out enterprise wireless to the first customer sites, engineers continue to explore the possibilities for new features that could augment the service in future maintenances. Wireless streaming capability from supported wireless devices to Apple TVs and Chromecasts is one example currently in testing. This feature would enable NOAA employees and their sponsored guests to wirelessly stream content to displays in conference spaces, meeting rooms and offices equipped with Apple TV or Chromecast devices. To connect, a user would only need to be in the same room as the device and have the device's access code. Furthermore, by detecting which devices are on which AP, the service could theoretically prompt a user to stream only to a device that is in the same room. This is incredibly useful as big campuses may have hundreds of media streaming devices across as many spaces.

Line and program offices that are interested in learning more about the N-Wave Enterprise Wireless service and how to order it for their sites can submit a new service request to the N-Wave NOC.

IJPS data flows migrate to JEUNO

The Joint EUMETSAT and NOAA (JEUNO) network infrastructure has been a topic of focus in previous editions of the N-Wave newsletter, with reports on implementation steps and progress of migrating operational data flows from the aging point-to-point transatlantic circuits to this new consolidated network architecture. JEUNO now supports wide area transport for the weather satellite data products exchange and operational data flows that keep track of

the health and safety of both NOAA and EUMETSAT spacecraft.

During the past six months, the joint team from the NOAA Office of Satellite and Products Operations (OSPO), Office of Satellite Ground Services (OSGS), the NESDIS Assistant CIO for Satellites, and N-Wave; partners ERT and Peraton; and EUMETSAT actively worked to migrate the operational Interim Joint Polar

Satellite (IJPS) data flows to JEUNO. Those flows were migrated away from the current OSPO-managed DS-3 (45 megabits per second) circuits, which routed from the National Satellite Operations Facility at Suitland, MD, and the Wallops Command and Data Acquisition Station at Wallops Island, VA, to Darmstadt, Germany.


The scheduled cutover from the DS-3 circuits to JEUNO occurred at approximately 13:00 UTC on September 19, 2018. All IJPS operational data flows for the Polar Operational Environmental Satellites (POES) and Meteorological Operational (MetOp) satellite are now migrated to JEUNO.

The notable difference: Average data rates are now over 50 mbps, with peaks of over 120 mbps, compared to about 30 mbps on the legacy transatlantic DS-3s. This results in improved timeliness transferring products to the Environmental Satellite Processing Center (ESPC) and POES. Previously, Level 1 and 2 products data files from Darmstadt were stacked as high as 50 files in the queue. Following the transition to JEUNO, the highest stack has been 3 to 4 files. Due

to the high capacity and the simplified architecture between NOAA and EUMETSAT, the latency for NOAA receiving some key EUMETSAT MetOp satellite products has decreased by 8-15 minutes.

Latency is one of the most important factors in satellite data product transmission, as old and stale data is useless for operational weather prediction systems. Near real-time satellite imagery and other satellite products delivery is always the goal, so this is a significant improvement and a win for the IJPS program.

The cutover for Joint Altimetry Satellite Oceanography Network (JASON) satellites data flows is planned for late November, which will complete the migration of all current operational data flows for both NOAA and EUMETSAT. The final set of the legacy transatlantic circuits will soon be decommissioned, saving both agencies considerable recurring costs while providing a more robust, secure, redundant, scalable and sustainable architecture for all future joint satellite missions and weather products exchange efforts.



Following the migration to JEUNO, IJPS data flow rates are now over 50 mbps, with peaks of over 120 mbps, compared to about 30 mbps on the legacy transatlantic DS-3s.

Completed Trusted Internet Connection project demonstrates engineering innovation

After several years of collaboration with the NOAA Cyber Security Division, N-Wave has completed the routing changes necessary to comply with the federally mandated Trusted Internet Connections (TIC) initiative. Internet traffic is in-lined at TIC sites across the country and must traverse the TIC 2.0 security stack.

To facilitate TIC, N-Wave engineers built a new national network called X-Wave with physically distinct hardware in four sites: Washington, D.C., Denver, CO, Seattle, WA, and Dallas, TX. This “untrust” network sits outside of the TIC firewalls, and it terminates all external peerings from service providers and other partners.

From the perspective of the N-Wave physical network, TIC compliance also represented an opportunity to build a national public internet domain on the “trust” side of the firewalls, referred to as the internet virtual routing and forwarding (VRF). Prior to this project, internet routing was a campus-based endeavor divided geographically across the country.

One of the realities of routing on the internet is asymmetry. While this reality is normal and functional on the internet, routing through distributed firewalls can make this asymmetry dangerous. For example, an outbound flow to the internet that egresses the Washington, D.C., TIC stack might return to N-Wave through the Denver TIC stack, where session state would not exist on the firewalls and the traffic would be dropped.

The new X-Wave routing design needed to account for this potential asymmetry. Using sophisticated routing policy on the Juniper MX routers, symmetry and failover are enforced both on redundant equipment pairs within Trusted Internet Connection Access Provider (TICAP) sites and also between redundant TICAP sites. In general, only the default route rather than full internet routes are distributed through N-Wave, and customers connected to remote aggregation routers are pinned to primary and backup TIC stacks.

The new design offers significant possibilities for robust and efficient routing for line and program offices and demonstrates N-Wave’s leadership as a network service provider for NOAA.

Only publicly routed IP subnets are routed on the N-Wave internet VRF. Both publicly routable and non-routable addresses continue to be routable on N-Wave private VRFs dedicated to individual customers or programs.

The X-Wave routing design also introduced some new capabilities. Generally speaking, high-bandwidth science flows struggle with firewalls that typically have shallow packet buffers. The X-Wave design includes the ability to leverage bypass links that can be used to divert mission-sensitive or high-bandwidth science traffic around the firewalls at each TICAP site. Enabled by BGP FlowSpec, N-Wave has a dynamic capability to steer traffic onto the bypass links. BGP FlowSpec, RFC 5575, can be thought of as a more elegant, protocol-signaled version of Cisco’s policy-based routing or Juniper’s filter-based forwarding.

Some elements of the original X-Wave design have not been fully implemented. The most notable of these is regional failover. Originally, N-Wave’s goal was to advertise all NOAA public addresses from all X-Wave peering routers. Due to backbone bandwidth and firewall capacity concerns, this functionality has not yet been enabled. N-Wave’s objective over the next year is to partly enable this functionality with IPv6 routes and subsets of the larger IPv4 routing table, as well as adding additional network capacity.

TIC continues to evolve, and new organizational and technical developments are on the horizon. Organizationally, the success of NOAA’s TIC deployment has attracted attention from Department of Commerce (DOC) leaders. N-Wave expects the first non-NOAA DOC customers will transition onto N-Wave’s network in 2019 to access the NOAA TIC service. (NOAA already provides internet service to other DOC agencies on the Boulder campus.)

On the technical side, the Department of Homeland Security is working with federal civilian agencies, including NOAA, to develop standards for TIC 3.0. The new requirements will likely focus more on endpoints and less on perimeterization. N-Wave engineers look forward to seeing what opportunities the new requirements will present.

The TIC 2.0 project has driven the growth of N-Wave over the past several years, as evidenced in the network traffic graph on page 16. It also enabled the N-Wave program to interact with many new

customers and deepened its relationship with existing customers. The completion of the project represents many hundreds of hours of engineering effort by NOAA engineers across N-Wave, Cyber and customer domains. The experience offered many valuable lessons about lab validation and testing, as well as the need for configuration standards and templates. While many engineers have played a part in the development and implementation of the new design, Jason Iannone of GlobalNOC is credited as its principal architect and tester.

Transporting NOAA to the cloud

Cloud computing is transforming how we think about IT infrastructure. As compute, storage and network resources become commoditized by hyperscale cloud providers, business problems and scientific analysis will be solved in the cloud as the first option whenever possible.

This migration to the cloud is already underway to varying degrees and at various levels of sophistication across NOAA. Recognizing this as a transformation in progress, N-Wave has taken on the responsibility to prepare for it and provide connectivity options that meet customer requirements for performance, flexibility, scalability and shared cost.

Today N-Wave has worked with customers to enable both virtual private network (VPN) and direct circuit-based connectivity to cloud providers. N-Wave's multiprotocol label switching (MPLS) network offers a virtual routing and forwarding (VRF) service that faces the Trusted Internet Connection Access Points (TICAPs) and carries traffic to and from the public internet (aka "internet VRF"). It also provides private VRFs to isolate line and program office traffic. This multi-tenancy capability can be extended to public cloud providers to extend FISMA boundaries or meet TIC compliance for internet-facing content. N-Wave has been engaged with many customers directly and through its work with the NOAA Web Operations Center (WOC) and Amazon Web Services (AWS) partner JHC Technology, Inc.

N-Wave looks forward to expanding its network service offerings to leverage shared resources as part of future NOAA-wide cloud procurement vehicles, such as the Cloud Integrated Project Team (IPT).


Use Cases

Two primary use cases currently exist for either VPN or direct connectivity: "private VRF extension" and "internet VRF only."

The private VRF extension use case extends a private VRF existing on the N-Wave MPLS network into a customer cloud environment. This use case connects a private VRF to a VPN tunnel running over the internet VRF. It will extend a customer's FISMA boundary into the cloud. The customer can leverage their existing firewall that is also connected to the customer's private VRF. IP addressing used in the cloud environment and routed over the private VRF can use either publicly routable or non-publicly routable addresses. Internet access can still be provided, either by routing public addresses from the customer firewall, or translating those to private addresses.

When the private VRF extension model is applied to direct connectivity rather than VPN, the connection is unencrypted. The subinterface and associated virtual local area network (VLAN) tag handed off from the direct connection are simply assigned to the correct routing instance on N-Wave. This is also true of the internet VRF only use case.

The internet VRF only use case helps meet TIC compliance. The use case connects a cloud deployment over a VPN tunnel to the interior of the internet VRF. This makes whatever is connected over the VPN appear to be behind the TIC firewalls. It is required that traffic in this use case be publicly routable. The addressing can be numbered in any N-Wave- or customer-owned IP address space being announced to the internet by N-Wave.



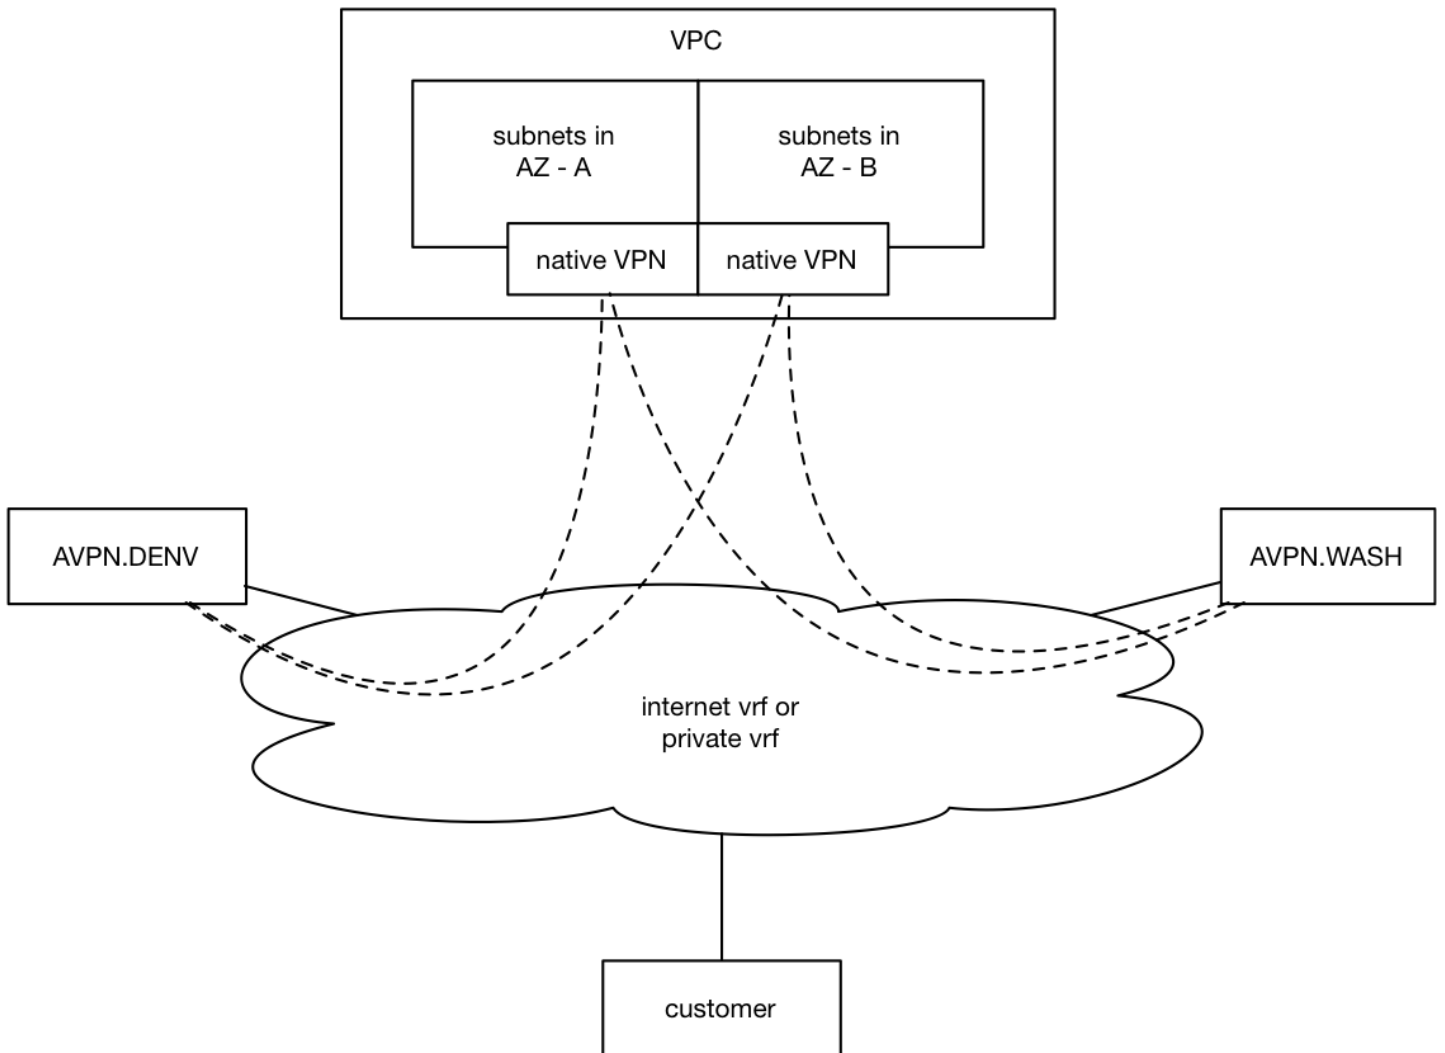

VPN Update

The initial WOC-managed AWS deployments connected through N-Wave used the Juniper vSRX as IPsec termination points for virtual private clouds (VPCs) in AWS. The vSRX offered a familiar operating system on which to run Border Gateway Protocol (BGP), IPsec and network address translation (NAT), while creating controlled demarcation points that could be managed with N-Wave's existing network management and audit tools. Over time, however, it became clear that the vendor-based model of high availability using firewall clusters did not map well to AWS. Without significant additional engineering, firewall deployments would be limited to single Availability Zone (AZ). Additionally, vendor-provided firewalls did not perform as well as expected. For example, the vSRX maxed out at 300 megabits per second of throughput per deployment in the best-case scenario.

For these reasons, as well as simplifying creation of new customer environments with automation, NOAA has moved to using the AWS natively provided VPNs. The main thing lost by not having the vSRX is a set of flexible NAT capabilities within AWS. This

has been overcome, when necessary, by configuring publicly routable addresses directly as the classless inter-domain routing ranges when the AWS VPCs are constructed.

The AWS native VPN is highly available by default, is spread across multiple AZs per VPC and provides two tunnels per customer gateway. It is capable of 1.25 gigabits per second per VPN (see also section on limitations on N-Wave VPN concentrators). AWS markets their VPN technology as CloudHub and intends for customers to use it as the hub in a hub-and-spoke routing topology over a service provider wide area network. They do not accept the same customer routes from different customer gateways. So, for example, N-Wave would not be able to advertise a default route to a customer VPC from both Washington, D.C., and Denver, CO. Initially, N-Wave engineers thought this limitation might prevent use of the native AWS VPNs. However, when the VPN tunnels and BGP session metrics have a strict active/passive configuration, failover performs as desired and the secondary routes become available within AWS after a network failure event.



N-Wave's existing generation of aggregate VPNs (AVPNs) are Juniper SRX1500s, which are bandwidth constrained. As currently configured, the aggregate of all VPNs connected to each AVPN is limited at 1 gbps. With additional configuration and tuning, it may be possible to get some more crypto bandwidth from the existing devices, but these should not be considered the platform for growth. There is a currently unfunded proposal to upgrade the AVPNs to devices that would each be capable of more than 30 gbps.

Direct Connectivity Update


Cloud providers offer direct circuit-based connections into their infrastructure. On AWS this service is referred to as Direct Connect. On Azure it is ExpressRoute. These direct interconnections offer higher speeds and lower latencies than can be achieved over the internet, as well as reduced cost for data egress fees.

Several companies now offer cloud connectivity broker services. Through a cloud broker, enterprise customers can simplify their direct connectivity by connecting only to the broker rather than each individual cloud.

Through its existing partnership with the Pacific Northwest Gigapop (PNWGP), N-Wave is partnering with Megaport to begin working on the cloud broker technical approach. Megaport offers a national layer 2 network over which N-Wave can extend its network to cloud providers.

N-Wave will provision two 10 gbps circuits to Megaport: one in Seattle, WA, and one in Washington, D.C. In Seattle, N-Wave's aggregation routers are located with PNWGP in the Westin Building, where Megaport is also located. In D.C., a 10 gbps wave through N-Wave's partnership with Internet2 will connect the aggregation router at N-Wave's McLean point-of-presence to Megaport in Ashburn, VA. Once the physical connections are in place, N-Wave will be able to provision VLANs from its aggregation routers through Megaport to a variety of cloud providers. Individual customer connections to cloud providers will be mapped to VLANs, and those connections can be rate-limited on the Megaport network.

N-Wave expects to turn up the first customer connections on this service before the end of 2018.



Network Changes and New Participants (April - October 2018)

National Ocean Service, Office of Response and Restoration, Gulf Disaster Response Center, Mobile, AL:

Working in partnership with the Florida Lambda Rail (FLR), N-Wave installed a network point-of-presence at the Office of Response and Restoration's Gulf Disaster Response Center in Mobile.

National Water Center, Tuscaloosa, AL:

N-Wave completed the installation of networking services to the National Weather Service's National Water Center (NWC) located on the University of Alabama campus in Tuscaloosa. Partnering with the University of Alabama, N-Wave was able to connect the NWC back to the network's Atlanta core location via dual, redundant 10 gbps circuits.

GOES-R LZSS (Wallops Island, VA) to NCEI-NC and CLASS (Asheville, NC):

N-Wave established a connection between the GOES-R Level Zero Storage System in Wallops Island to the National Centers for Environmental Information-North Carolina and the Comprehensive Large Array-data Stewardship System in Asheville. Completed on August 21, 2018, this new connection allows NCEI and CLASS to archive the level zero, or raw, GOES-R series satellite data in Asheville.

NOAA JASON Ground System (NJGS)

migration: N-Wave and Office of Satellite and Product Operations (OSPO) engineers migrated the NOAA JASON polar satellites ground system from OSPO-managed T1 circuits to N-Wave 1 gbps private virtual routing and forwarding, or VRF, connectivity between three major sites: Suitland, MD, Wallops Island, VA, and Fairbanks, AK. This change, completed and fully operational as of August 21, 2018, allows the National Environmental Satellite, Data, and Information Service (NESDIS) to free up and decommission legacy (and expensive) commercial T1 circuits. The JASON mission can now take advantage of increased bandwidth from 1.5 mbps up to 1 gbps, increased stability and significantly reduced latency. This also provides OSPO with significant annual cost savings after decommissioning the legacy circuits between Fairbanks and Suitland.

Partner Antenna Access Network (PAAN) TIC Connectivity at Fairbanks Command and Data Acquisition Station (FCDAS):

N-Wave established Partner Antenna Access Network TIC connectivity at the Fairbanks Command and Data Acquisition Station in Alaska. Completed August 30, 2018, this connection allows external meteorological satellite partners to access the FCDAS antenna resources.

Deep Space Climate Observatory (DSCOVR) Satellite Ground System Migration:

N-Wave and OSPO engineers completed the migration of the DSCOVR Ground System connections from slow and expensive T1s to the N-Wave 1 gbps infrastructure. This not only allows for higher speeds, stability and low latency, but also allows for better network redundancy and availability on the ground system – functionality previously not available on the ground system.

Satellite Controller Communication Services (SCCS) Connectivity:

OSPO is using N-Wave for their mission voice system replacement project. A new Voice over IP, or VoIP, system was deployed by NESDIS/OSPO to replace an aging mission voice system that ran over partial T1 lines. The new system is using N-Wave 1 gbps private IP transport infrastructure to interconnect the system between the key NESDIS sites: Suitland, MD, Wallops Island, VA, Fairmont, WV, and Fairbanks, AK.

JPSS-2, 3 and 4 Spacecraft Checkout

Activities Connectivity: On October 10, 2018, a new N-Wave point-of-presence was established at the Northrop Grumman (formally Orbital ATK) satellite manufacturing facility at Gilbert, AZ. The use cases include:

- JPSS-2, 3 and 4 Joint Confidence Test (JCT) instrumentation check-out activities at the Raytheon Factory in Riverdale, MD, and NOAA Satellite Operations Facility in Suitland, MD.
- JPSS GRAVITE to receive high rate data telemetry in their Instrumentation Validation and Verification (IV&V) environment in

Suitland, MD

- NASA Goddard Space Flight Center terminals in Greenbelt, MD

San Diego Pier - Office of Marine and Aviation Operations:

An N-Wave network presence was recently installed at the OMAO office in San Diego, CA, to support offloading data from NOAA ships returning from scientific missions. With the ever-increasing advancements in scientific data collection tools, scientists have resorted to moving data via Sneakernet (portable hard drives) to get their data off

the ships. With N-Wave connectivity to the ships, this data can be offloaded directly to the end processing site without the extra logistical steps of shipping hard drives around the country.

National Ocean Service Santa Barbara, CA:

Partnering with the Corporation for Education Network Initiatives in California (CENIC), N-Wave has completed installation of networking services to the NOS Channel Islands National Marine Sanctuary, located on the University of California Santa Barbara campus.


Multicast changes coming to N-Wave

NOAA's research vessel, the *Okeanos Explorer*, uses a special network protocol called multicast to [livestream video from the ship](#) and its remotely operated vehicle (ROV). Multicast allows the ship to use a single IP stream to share content with many viewers. This is a huge advantage, as the ship's limited bandwidth at sea would not support a single stream to each interested viewer.

The *Okeanos Explorer* currently uses Any-Source Multicast (ASM), which the networking industry is now phasing out as a second form of multicast has emerged to replace it – Source-Specific Multicast (SSM). In anticipation, the ship will migrate to SSM in early December 2018.

As its name suggests, SSM is designed for only one source, whereas ASM was designed to allow all participants to send as well as receive communications. This has made ASM complex to use and maintain. SSM, with its single source, is simpler to maintain and still meets the *Okeanos Explorer's* needs. (The ship is the only consistent user of ASM on N-Wave.)

Internet2, N-Wave's partner, will turn down ASM on December 15, 2018. This aligns with recommendations by the Internet Engineering Task Force to deprecate ASM and use SSM for all multicast scenarios. This change will allow N-Wave and Internet2 to simplify their networks' configurations and reduce maintenance efforts, ultimately providing for a more robust transport service.




ROV Deep Discoverer shines its lights and cameras on a very dense community of large, presumably very old, colonies of bamboo coral (isididae) that were found on the ridge crest of Ellis Seamount. Image courtesy of the NOAA Office of Ocean Exploration and Research, 2015 Hohonu Moana.

N-Wave Network Performance Metrics

Since the last report in the spring 2018 newsletter, network traffic rates across N-Wave have leveled off. This seems to coincide with several of the N-Wave backbone links reaching full capacity. Traffic to the Fairmont data center from the new National Centers for Environmental Prediction (NCEP) supercomputers, MARS and VENUS, combined with legacy NCEP and Research and Development High Perform-

mance Computing systems has saturated the links from Washington, D.C., to Fairmont, WV. In FY19, N-Wave will roll out multiple 100+ gigabits per second upgrades across the backbone, D.C. Metro ring and the Fairmont data center. Following those upgrades, N-Wave anticipates a return to the growth trend seen in the past.

Network Traffic (December 2011 - September 2018)



100 gigabit upgrades coming to N-Wave

After eight years of operation and multiple augmentations to the original 10 gigabits per second backbone, N-Wave will soon upgrade to 100+ gbps network speeds.

Multiple new customer sites, big data demands and new services – including Trusted Internet Connections and NOAA supercomputing expansions – have stressed the existing backbone and shared data center circuits. Recognizing the limitations of continued incremental 10 gbps expansions, N-Wave concluded the next magnitude in upgrades is required. In Q4 of FY18, N-Wave purchased the necessary routing and dense wavelength division multiplexing

(DWDM) hardware to facilitate the upgrades.

During FY19, N-Wave will complete the upgrade in three phases. Phase one will involve implementing the upgraded core between Denver, CO, Chicago, IL, Atlanta, GA, and Washington, D.C., as well as the connections from the core to the NESCC shared data center in Fairmont, WV. Working with Ciena for phase two, N-Wave will update the D.C. Metro DWDM ring to 400 gbps of capacity utilizing the new Waveserver Ai platform. Finally, phase three will involve 100 gbps upgrades from the core to the David Skaggs Research Center in Boulder, CO.


New N-Wave ticket management system offers improved metrics and visualizations

N-Wave recently transitioned to a new ticket management system called ServiceNow. The official cutover date for this change was September 24, 2018.

ServiceNow enables the N-Wave Network Operations Center (NOC) to collect more granular task data, which will offer new views and opportunities for analysis that will ultimately benefit N-Wave processes and services. ServiceNow additionally displays ticketing data in a dynamically updated dashboard.

The example dashboard view below shows data related to catalog tasks, which are tasks needed to complete a given service request. When a customer submits a service request, the NOC will create a request ticket and then assign associated catalog tasks to the team or

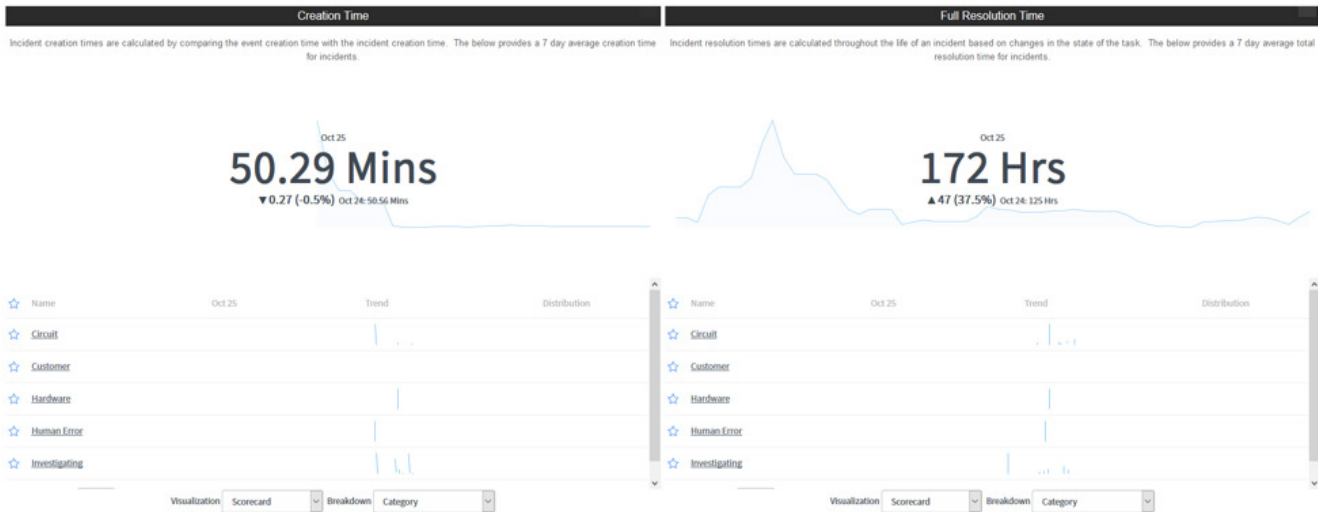
teams who can complete the request. The largest graph shows two datasets: the blue line represents catalog tasks submitted on a given date, and the orange line represents the trend for that particular type of task. The donut chart at the bottom left shows assignment groups, with each color representing an N-Wave group responsible for a different type of task. The line graph at the bottom right represents the number of catalog tasks assigned to an individual within N-Wave. Moving forward, the NOC will have these types of dashboard views for standard tasks, incidents, incident tasks, catalog tasks, and problems, with the goal of adhering more closely to ITIL IT Service Management standards. Dynamic data views like these will offer new insights into N-Wave services and customer support needs.



**Note: All images and data presented in this article should be viewed as examples. The sampling period since the September 24, 2018, transition was too brief to represent an accurate data set.*

The second example dashboard below displays incident creation time compared to full incident resolution time. With this information, the NOC will be able to look at specific dates to examine creation versus completion times, as well as the overall averages for those metrics. This information also will show

trends as to the types of incidents the NOC processes. The ability to collect and view this information in this way will allow the NOC to analyze areas of excellence and areas for improvement, which will lead to better processes and procedures.



N-Wave adds five new staff members

N-Wave’s expanding team includes five new network engineers: Mike Warner, Robert Webb, Peyman Paysarvihoseini, William Connors and Jason Wilsey.




Mike Warner joined the N-Wave team in June and is based at the Boulder, CO, campus. His areas of focus within N-Wave include Core switched infrastructure for the Boulder campus, lab management and Microsoft Azure

integration. Before joining N-Wave, Mike worked as a network engineer for Pacific Pulmonary Services.



William Connors joined N-Wave in October and is based at the Western Region Center (WRC), Sand Point, WA. He is focused on N-Wave services including enterprise wireless and WRC campus networking. William has a degree in network administration, and his career

background includes recent roles in network security.



Robert Webb joined N-Wave in August. He is based at the Fairmont, WV, NESCC site and is initially focused on providing network support for the Commerce Business Systems (CBS) Core Financial System (CFS), which is provided by the NOAA Corporate Services

Financial Systems Division. Before joining the N-Wave team, Robert worked as a senior engineer for the NOAA Cyber Security Operations Center.



Jason Wilsey joined N-Wave in October and is based at the Boulder campus. His area of focus is on enterprise wireless and campus network services. Before joining N-Wave, he was a systems administrator for the NOAA Earth System Research Laboratory Director’s Office, Office of the

Executive Director and the NOAA Boulder Library.



Peyman Paysarvihoseini joined N-Wave in October. He is based out of Largo, MD, and his role with N-Wave is focused on supporting the Commerce Business Systems (CBS) at the NOAA ITC. Peyman’s background includes a Master of Science degree.

Welcome to the N-Wave team, Mike, Robert, Peyman, William and Jason!

N-Wave is built on partnerships

N-Wave's partnerships with other organizations in the scientific, research and education and (SR&E) networking communities are essential for delivering the innovative underpinnings of the N-Wave network. While many of these relationships have existed since N-Wave's early years of operation, we take this opportunity to reiterate the value of our partners.



Celebrating 30 years of SC - SC18 in Dallas

This year marks the 30th anniversary of the SC Conference, with special events to celebrate the evolution of SC and progress of the HPC community. The conference returns to Dallas, TX, November 11–16, 2018.

SC18 General Chair Ralph McElDowney announced the theme of the year's conference, "HPC Inspires," as a gathering to "inspire the next generation of scientists and engineers to find solutions to the world's greatest challenges." The event includes technical presentations, papers, workshops, tutorials, research posters, Birds of a Feather sessions, and exhibits showcasing the latest innovations in HPC and related fields.

SCinet, SC's dedicated high-capacity network infrastructure introduced to the conference in 1991, will again emerge as the fastest and most powerful network in the world. The network is anticipated to deliver 4.02 terabits per second of wide area capacity to the Kay Bailey Hutchison Convention Center Dallas for SC18.

N-Wave team members attending SC18: Jerry Janssen, Paul Love, Amber Rasche and Robert Sears.



N-WaveNews

Issue 12

November 2018



NOAA ENTERPRISE NETWORK

For more information contact:
NOAA N-Wave Program
<http://noc.nwave.noaa.gov>
Office of the Chief Information Officer
<http://cio.noaa.gov>
Robert Sears, Network Manager
Paul Love, Newsletter Coordinator
Holly Palm, Design and Layout
Amber Rasche, Editor

U. S. Department of Commerce, NOAA
N-WAVE
325 Broadway
Boulder, CO 80305-3337