ISSUE
10
November
2017

N-Wave News

in this edition:

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

One word that sums up Fiscal Year 2017 (FY17) for the N-Wave program is "milestones." There were many technical and programmatic accomplishments this year that either completed long-term projects or established the foundation for expansion of services. All efforts directly aligned with, or completed actions outlined in, NOAA's Strategic Plan for Network Optimization and Transport.

January 2017 began with the integration of N-Wave and the Silver Spring Network Operations Center. The two teams have come together under N-Wave management with the goal to optimize NOAA-wide networking services. One key objective in NOAA's network strategic plan is to unite the network operations groups that are providing multiple network services across NOAA line offices under one program. What has been gained is the incorporation of exceptional engineering staff to tackle the varying and expanding NOAA-wide networking requirements. N-Wave engineers new to the team are Michael Mankarious, Francois Ntowo, Leroy Cain, Jonathan Ricks, Tony Zhang, and David Steger. They bring many years of NOAA network engineering experience and are supporting key roles within N-Wave. A public and gracious welcome to the team!

The spring of 2017 brought to completion a year-long effort to solidify N-Wave's ten-year technical strategy and associated acquisition and investment authority. With approvals received across NOAA and DOC senior management, N-Wave will continue its partnerships within the Science, Research, and Education community. Since N-Wave's inception over seven years ago, this approach has allowed N-Wave to expand its infrastructure and services while integrating the innovation this community excels in. N-Wave will work with key regional optical network providers, Internet2, the GlobalNOC at Indiana University and within the General Services Administration (GSA) Enterprise Infrastructure Solutions program (EIS) to integrate the most flexible, widest range of networking capabilities for NOAA and all N-Wave customers.

Although there are many other accomplishments, which you can read about in the "Network Changes" article on page 11, one of the most impactful projects, and one that has truly kept the innovative edge, is the completion of all NOAA Trusted Internet Connection Access Points (TICAP). On June 26, 2017, after months of coordination across all line offices, the largest region of NOAA's historical Internet access was migrated to the new Washington, DC, metro TICAP. To add to the complexity of this effort, many N-Wave engineers from Boulder, CO, and Indiana University's GlobalNOC were in Silver Spring, MD, for team integration meetings. In order to meet deadlines amongst ever-looming critical weather delays, a war room of sorts was set up at a hotel conference room, where the out-of-town engineers were staying. At 8:00 PM Eastern, N-Wave engineers at the hotel, in conjunction with N-Wave engineers in SSMC3 and other locations, made the cutover of migrating all NOAA DC metro traffic from legacy infrastructure to the new TICAP, while over 20 line office engineers participated via phone bridge to test operations post-migration.

War Room - Left to right: Dave Mauro, Jason Iannone, Jared Brown, and Mark Mutz

Not a month later, the same approach was followed for the migration to the Denver, CO, TICAP. These actions completed the multiple year transition of all NOAA ingress and egress Internet traffic to the integrated X-Wave/N-Wave TICAP infrastructure. This not only established new highly available, scalable network infrastructures for NOAA, but also supported NOAA Cyber Security in meeting the TIC requirements of the Department of Homeland Security.

I want to thank Alex Hsia of the Boulder NOC for the Denver migration efforts, N-Wave team members Dave Mauro and Mark Mutz for months of coordination efforts for the DC metro TICAP, Dave Mauro leading the charge for war room coordination, Jason Iannone, Jared Brown, Jared Schlemmer, Michael Mankarios, and Tony Zhang for manning the routers and making the required technical cutover, along with all the line office engineers who assisted in ensuring availability of their perspective operations post-migration.

A great coordinated NOAA-wide effort!

N-Wave Cloud Update

Overview

N-Wave has been engaged with the NOAA Web Operations Center (WOC) on their project to enable Amazon Web Services (AWS) connectivity for line offices. The WOC has been working with Amazon partner JHC to understand customer requirements and migrate web resources to AWS via direct customer engagement and bi-weekly Cloud Working Group meetings. N-Wave is interested in helping provide standardized, centrally managed connectivity to AWS that addresses the needs of all customers. N-Wave is also keenly interested in engaging with line office network teams to understand cloud needs, so that we can help facilitate enterprise-wide connectivity. Please talk with us via the NOAA Networking Committee (NNC) or directly through the N-Wave service inquiry page at https://noc.nwave.noaa.gov/

To date, N-Wave has helped facilitate the WOC effort at the enterprise level by providing VPN and Direct Connectivity into AWS. VPN connectivity is over the Internet, and Direct Connect (as the name suggests) is a direct unencrypted connection between AWS and N-WAVE via a circuit that connects an N-WAVE router directly to AWS. (The equivalent Microsoft Azure product is Express Route.)

In order to maintain the Department of Homeland Security mandated Trusted Internet Connection (TIC) compliance, all NOAA public-facing content hosted by a cloud provider, including AWS, must pass through a TICAP stack. To make this possible as illustrated in the Cloud TICAP Compliance diagram below, N-Wave must draw traffic into the NOAA TICAP locations before routing it over either a VPN or a Direct Connect path. For the AWS specific projects, this is done by hosting NOAA addressing in AWS Virtual Private Clouds (VPC), either via Network Address Translator (NAT) in the VPN case or natively in VPC in the Direct Connect case. These smaller blocks hosted in AWS are part of larger supernets, which are being advertised from the X-Wave border routers to the Internet. This hairpinning of traffic is illustrated below. The following illustration shows the VPN scenario. The Direct Connect works similarly, though the AWS resources are directly connected to N-Wave.

Cloud TICAP Compliance

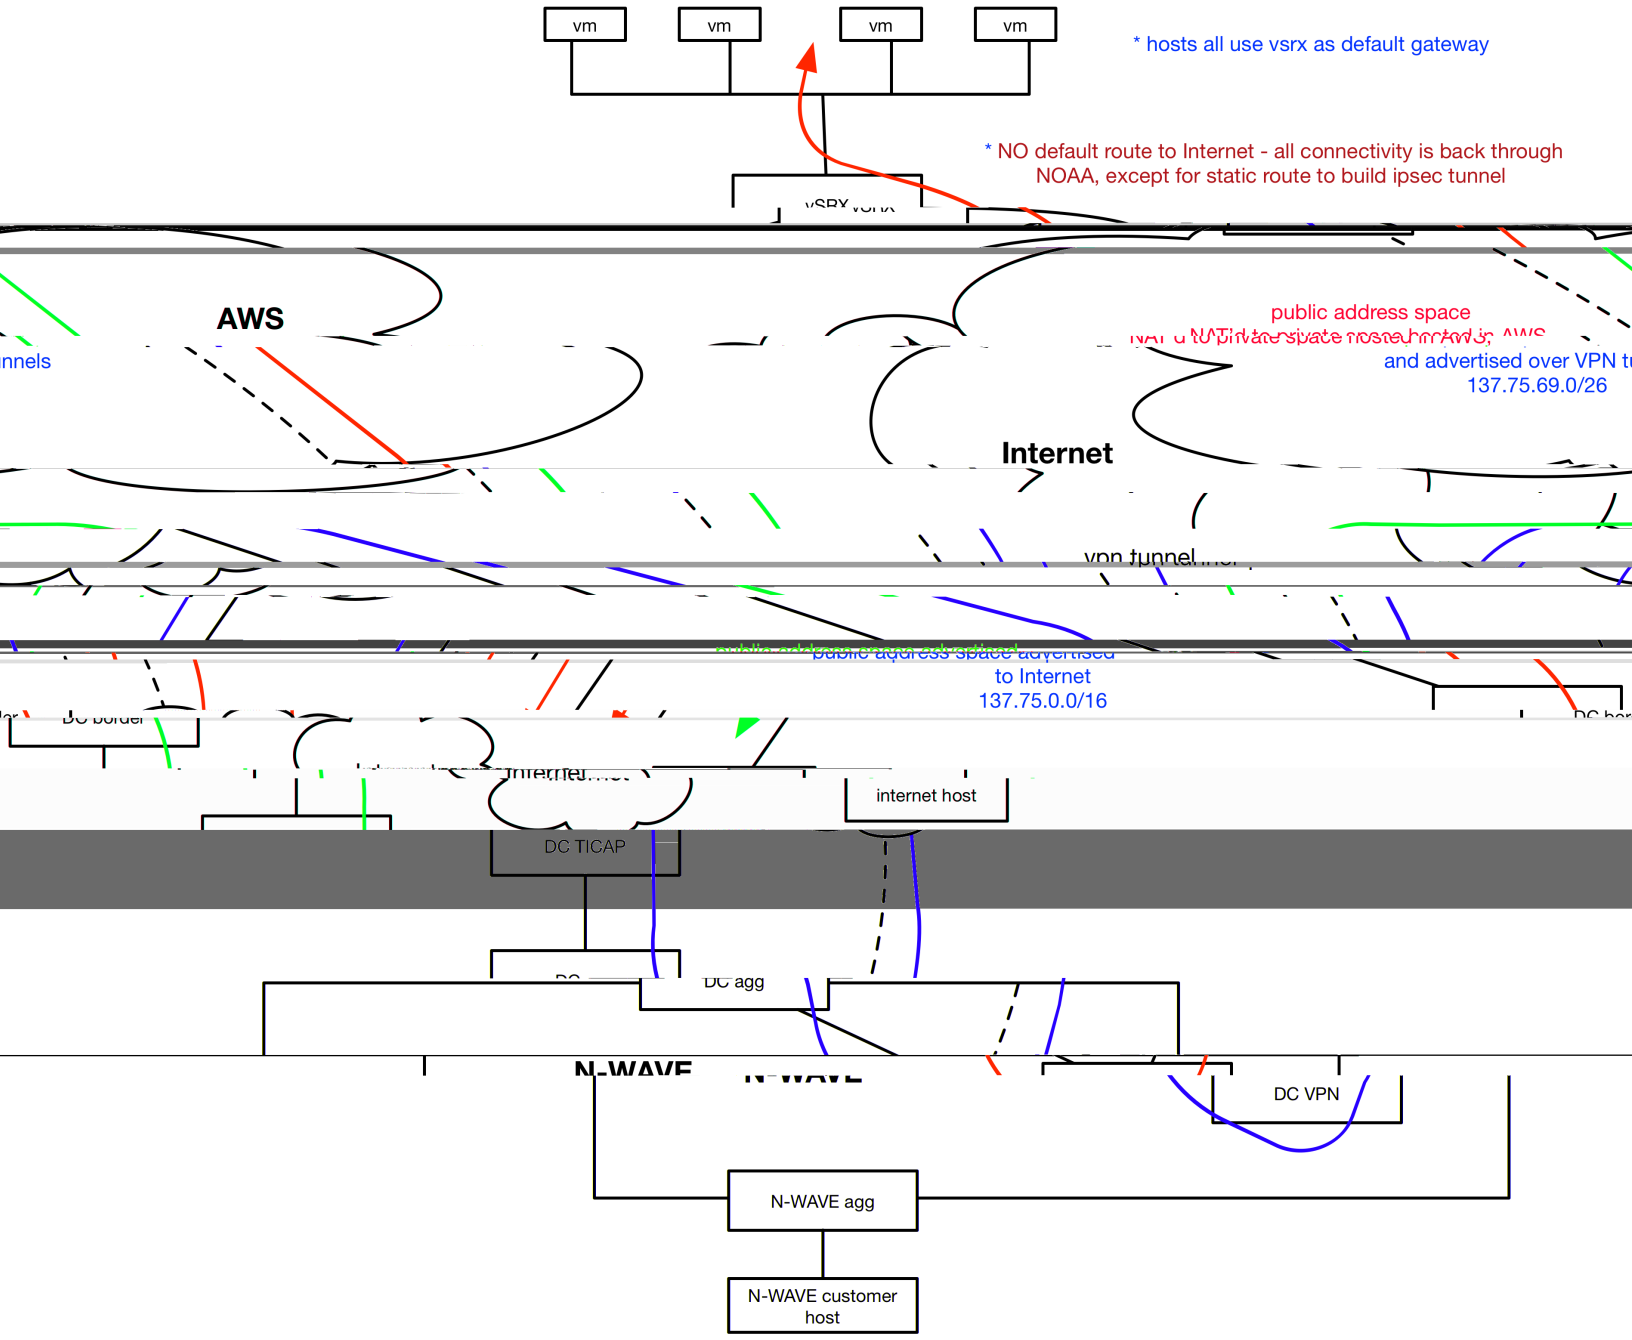

This solution mainly addresses server-based virtual machines hosted as Elastic Compute Cloud (EC2) instances in AWS Virtual Private Clouds (VPC) rather than any of the many serverless products offered by AWS. Clearly, the hairpinning is less than ideal, but NOAA and other agencies with TIC compliance requirements don't yet have an alternative. The end of this article will explore potential alternative solutions that are in the early stages of discussion.

N-Wave is aware that some line offices are independently working on cloud solutions with Amazon or Microsoft, some of which pre-date the WOC effort. It is important to remember the requirement to pass traffic bound for NOAA-owned content in the cloud through the TIC stack. N-Wave would like to work with all line offices to fulfill this requirement, as well as all other connectivity requirements.

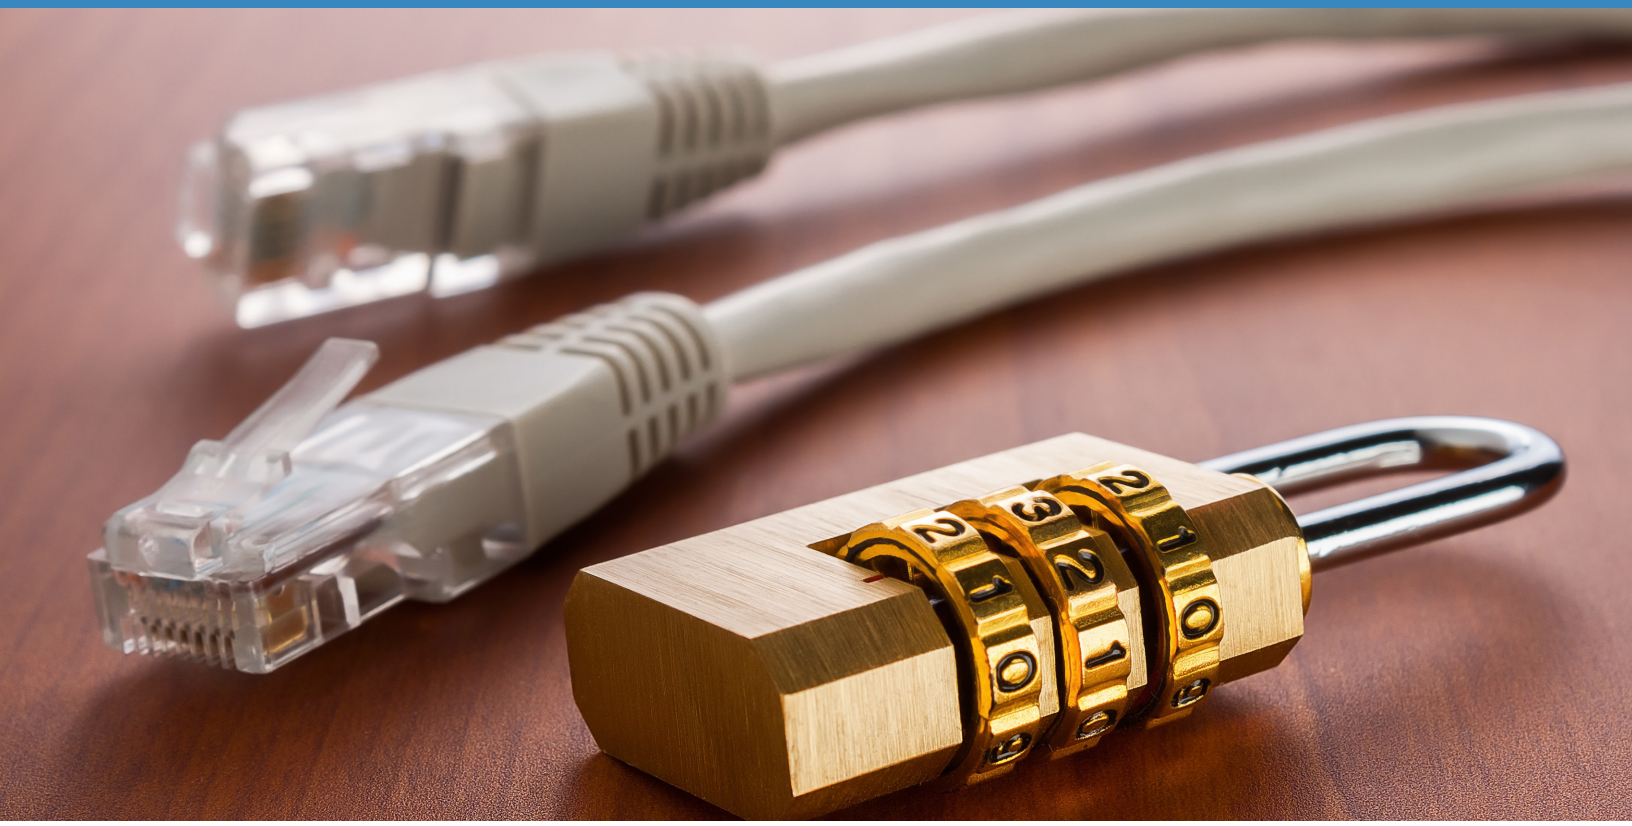

VPN

The VPN solution has been tested and deployed throughout 2017. Compared to the Direct Connect option, it is a lower-bandwidth, lower-cost solution that does not require circuit-based connectivity to AWS but instead uses the Internet. The solution encrypts all traffic from the NOAA premise to the VPC.

The standard deployment involves a virtual firewall and, in the case of the WOC projects, a Juniper vSRX hosted in a VPC. This firewall has two IP Security (IPSec) tunnels to VPN concentrators located in Washington, DC, and Denver, CO. Border Gateway Protocol (BGP) is used for failing over traffic between the two sites. All AWS VPCs for the current projects are hosted in Northern Virginia data centers, so the DC path is always preferred.

There are three kinds of subnets configured on the vSRX: OOB (out of band) for emergency access to the device; UNTRUST for building the VPN tunnels back to the NOAA premise; and one or more TRUST interfaces, which support EC2 instances behind the firewall. In order to conserve public IP address space, all VPCs are numbered with private IP addresses for their Classless Inter-Domain Routing (CIDR) block. This prevents us from having to waste /28s of public address space for the OOB and UNTRUST interfaces. The OOB and UNTRUST interfaces have AWS Elastic IP (EIP) addresses, which provides them direct access to the internet. On the UNTRUST interface, this Internet access is used only to build VPN tunnels. EC2 instances on the TRUST subnet(s) must ingress and egress via the vSRX, which routes the traffic back to NOAA over the VPN tunnels. There should be no direct access for EC2 instances to the Internet via AWS; all traffic hairpins back to the NOAA premise before routing to or from the Internet.

The vSRX is not explicitly being used as a customer firewall, but rather as a managed endpoint device in AWS. Its primary roles are IPSec manageability, BGP and routing capability, and NAT for the TIC hairpinning. It is integrated into Global NOC management tools so we can administer it like any other device for authentication and security scanning. It also gives us a standard deployment model which enhances N-WAVE's ability to deliver services quickly and in a standardized way to internal customers.

Turning up this VPN service exposed N-WAVE engineers to many of the AWS internals, which will enable us to offer consultative advice when customers are building cloud infrastructure.

Direct Connect

An alternative to VPN connectivity is having a direct peering connection from N-WAVE to Amazon over a privately provisioned circuit. The first customer using Direct Connect with N-WAVE is NESDIS for the Secure Ingest Project (SIGP) pilot, which analyzes untrusted data from external partners. There is a 10 G circuit from an N-WAVE aggregation router in McLean, VA, to an Amazon router in Ashburn, VA, which was turned up in Q3 of 2017. Traffic over this connection does not need to be encrypted and the performance is much better than what we've seen over the Internet.

As the project is a pilot, network requirements are being determined on an ongoing basis, but the information here is correct at the time of writing. It was necessary to solve the hairpinning problem for EC2 instances living inside the NESDIS FISMA boundary. There were several options available to us, but the simplest solution–hosting NOAA-owned addressed space in a VPC connected to an AWS Private Virtual Interface (VIF)–has proven successful so far. SIGP is using other AWS products like Simple Storage Service (S3, a serverless technology for web-based file storage) and Lambda (a serverless technology that runs code abstracted from having to provision any infrastructure) that are outside of the VPC connected over the new circuit, and not subject to hairpinning.

Upon successful completion of the pilot project and determination of ongoing bandwidth needs, NESDIS is open to the possibility of cost- and bandwidth-sharing of this circuit with other line offices for a more enterprise approach. N-Wave will be able to logically separate customers on different accounts and VPCs with VLANs, and enforce bandwidth levels using traffic policing or shaping.

Turning up this Direct Connect was another learning experience for N-Wave. There are two basic kinds of services offered over Direct Connect: public or private VIFs. The private VIF is generally used to extend an organization's private infrastructure and addressing into a VPC. This is what we used with SIGP, except we used NOAA public addressing. We did this to achieve the same hairpinning routing behavior used with the VPN. The VPC advertises a smaller block of NOAA-owned space into the N-Wave interior, which is then advertised to the Internet as part of a larger subnet. Using the public addressing natively in the VPC eliminated the need for the vSRX, which would have provided VPN, NAT, or BGP services, none of which was needed.

The public VIF alternative is for public peering with Amazon. When you turn up a public VIF to AWS, it is similar to a peering to any other external service provider. Amazon provides you all their internal routes. This is now over one thousand routes, which includes many blocks as large as /12s. This is understandable considering how much of the Internet is now hosted on AWS. N-Wave would advertise its own NOAA-owned addressing to AWS. If there were firewalls in the path, we would have to enforce symmetry with routing, or NAT.

Public and private VIFs can co-exist on different VLANs on the same Direct Connect circuit. Both kinds of peering to AWS require BGP with authentication. Bidirectional Forwarding Detection (BFD) is enabled by default and works immediately when you configure it. This protocol detects link-level problems along the path of the circuit in a subsecond time frame, which may not be detected by BGP as quickly, and enables BGP to be brought down. The Maximum Transmission Unit (MTU) supported by AWS on Direct Connect is 1522 bytes.

Future directions

Currently, there are two areas of interest that may relate to each other. The first is finding alternatives to the hairpinning requirement. The second involves moving N-Wave closer to cloud providers.

First, let us address alternatives to hairpinning. As described, the hairpinning solution is not ideal. Using NOAA addressing severely limits NOAA's ability to take advantage of the cloud's elasticity and agility. An area of interest within NOAA Cyber Security and N-Wave is meeting TICAP compliance by building an agency-managed security infrastructure in the cloud. AWS itself offers solutions that will help partially meet this compliance.

Second, we are focused on connecting N-Wave more directly to cloud services in order to serve all customers with enterprise-wide solutions. One possibility here is extending our network into colocation facilities like Equinix, where all cloud providers are reachable via cross-connects. (This could also take the form of connectivity in our existing TICAP locations if the appropriate connection types and carriers are available there.) If our MPLS is extended to routers colocated with direct Cloud peering, any N-Wave service could be easily offered for Internet or private Virtual Routing and Forwarding (VRFs), or for Layer 2 or Layer 3 VPN services, to any customer. Sharing the connectivity among all line offices could reduce costs. In addition to cloud providers, all other traditional telco and Internet service providers, as well as many other federal agencies, are already located in these large colocation facilities. Over time, the orientation of the network could begin to focus on these colocation facilities for things like TICAP and N-Wave core connectivity. The utility of MPLS is enormous for delivering customer services. Extending the network in this way would give us the agility to solve all foreseeable network challenges–and give us the best chance of addressing the ones we cannot foresee.

N-Wave Network Performance Metrics
Fall 2017

Welcome to the Fall 2017 edition of our network performance metrics discussion. As shown in the figure below, traffic levels on N-Wave continue to increase at a substantial rate, almost doubling over the last year.

One of our largest increases came at the end of June, when the remainder of the NOAA headquarters in Silver Spring, MD, was migrated to the new DC metro TICAP 2.0 infrastructure. Included in this migration was the National Weather Service and their dissemination of large volumes of weather data. These products account for around 75% of NOAA's Internet traffic, reaching daily peaks of 20 Gbps. To keep up with demand, N-Wave completed a project to increase the network interconnects through the TICAP infrastructure to 40 Gbps in the DC metro area and 30 Gbps in Denver, CO.

The Research and Development High Performance Computing System (RDHPCS) continues to generate the majority of internal N-Wave traffic. Transfers between RDHPCS locations in Boulder, CO, Fairmont, WV, and Princeton, NJ, average around 175 terabytes per day and have remained relatively constant over the last year.

Looking towards the future, N-Wave engineers continue to investigate options for increasing backbone and TICAP capacity to 100 Gbps to handle NOAA's increasing data needs.

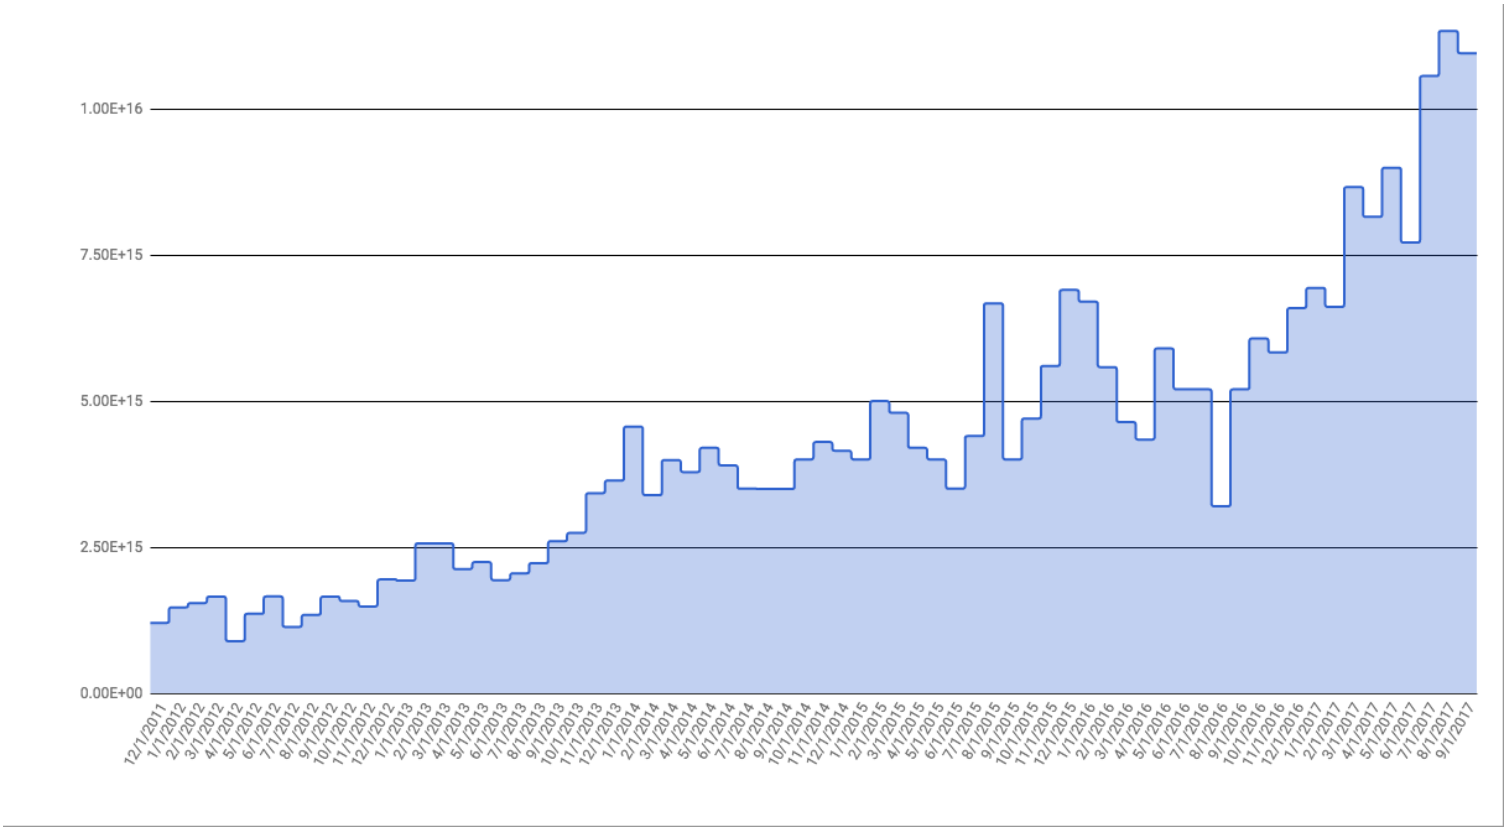

N-Wave 100 G Proof of Concept

N-Wave has partnered with Internet2 and Ciena Corporation to demonstrate a 100 Gigabit per second (Gbps) long-haul network service using next-generation Dense Wavelength Division Multiplexing (DWDM) equipment. In early 2017, N-Wave engineers enabled a proof-of-concept service capable of providing 100 Gbps of network transport between the N-Wave core nodes in Denver, CO; Chicago, IL; and McLean, VA.

This proof-of-concept network was implemented as an overlay wavelength atop Internet2's DWDM optical network. Using Ciena optical network equipment, N-Wave lit a single 100 Gbps wavelength that carried 10 individual 10 Gbps sub-channels. Each individual 10 Gbps circuit has the capability of providing 10 Gbps of bandwidth between Denver and Washington, DC. Alternatively, the optical system provides the capability to drop some or all of the 10 Gbps sub-circuits in Chicago.

Ciena, Internet2, and N-Wave engineers collaborated to test next-generation 100 Gbps DWDM linecards in the Internet2 system. These cards provided the same 100 Gbps reliability that Internet2 has enjoyed on its optical system, but the cards provide a longer distance between electrical signal regeneration than previous 100 Gbps linecards. Each individual wavelength on the Internet2 optical system will accumulate some degree of noise as the optical signal is periodically amplified along the fiber path. Once that noise reaches a certain threshold, the signal needs to be regenerated using a back-to-back pair of 100 Gbps DWDM linecards. In the industry, this is referred to as Optical-Electric-Optical (O-E-O) regeneration–and it's costly.

Ideally, 100 Gbps linecards are only placed where the signal is required to be dropped out to a high speed router or switch. The previous generation of 100 Gbps linecards would have required a costly signal regeneration in Kansas City, MO, and in Cleveland, OH. The next-generation 100 Gbps linecards in the proof-of-concept network were able to push the entire distance from Denver to Chicago to Washington, DC, without any unnecessary O-E-O signal regeneration in the middle.

This translates to an overall cost reduction to implement the 100 G transport. By pushing the signal further, increased bandwidth coupled with fewer complex pieces of equipment in the path provides cost savings and additional reliability.

Since February 2015, N-Wave has operated a 100 Gbps metro-area DWDM ring network in the DC area: Silver Spring, MD (SSMC); College Park, MD (the University of Maryland Mid-Atlantic Crossroads [MAX] co-location facility); Suitland, MD (NSOF-NOAA Satellite Operations Facility); and McLean, VA (Internet2–Level 3 co-location facility and site of an N-Wave code node). The experience gained by installing and operating this metro-area network will be valuable in jointly operating a wide-area 100 Gbps service with Internet2. The goal of this 100 G proof of concept with Internet2 and Ciena was to test and provide options for the N-Wave next-generation backbone.

These collaborations with key partners create many opportunities for N-Wave; in addition to optimizing the costs of network transport, future network devices (such as routers and switches) can operate in the optical, Ethernet transport, and Internet Protocol (IP) domains. Coupled with Software-Defined Networking (SDN), the convergence of the Open System Interconnect (OSI) model Layers 1, 2, and 3 into a single device creates opportunities for faster service provisioning, lower cost transport, and greater flexibility operating enterprise networks.

N-Wave NOC Tickets Report March-August 2017

The N-Wave Network Operations Center (NOC) is a 24/7/365 tier 1 support center for Boulder (N-Wave West), Silver Spring (N-Wave East), and N-Wave (Central). The NOC is a first point of contact for various participants, engineering groups, and vendors. It can be contacted via phone, email, or a self-serve webform ticketing structure.

The N-Wave NOC is responsible for triaging requests as well as assigning each individual request to the appropriate N-Wave team to be resolved when necessary. Staff at the N-Wave NOC are continually trained in order to handle each issue effectively. The N-Wave NOC aims to be the first contact point for all who participate in the N-Wave network.

The graph provided shows the various ticket types that the N-Wave NOC is responsible for. Administrative tickets comprise the largest amount of ticket types. These are used for various services and requests such as VPN, IPAM, Security Inquiries, and Wireless. A new type of ticket that the N-Wave NOC will utilize is the "Documentation Request." This is an attempt for the N-Wave NOC to be transparent regarding its developing processes and procedures as N-Wave continues to grow and mature.

The N-Wave Network Operations Center is available 24 hours a day, 7 days a week, 365 days a year at (812) 856-7477 and nwave-noc@noaa.gov or noc@noaa.gov.

To report a problem using the N-Wave webform: http://noc.nwave.noaa.gov/nwave/support/trouble-ticket-request-form.html

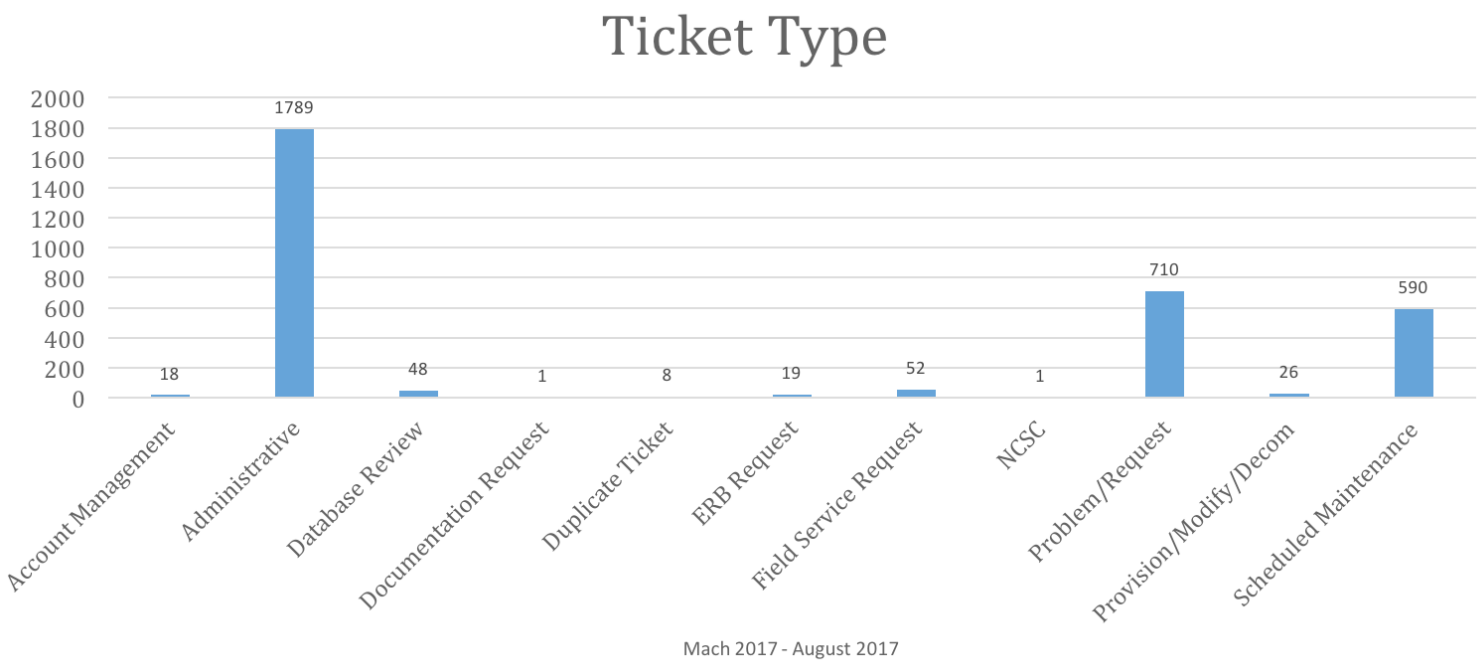

Network Changes and New Participants

NOAA Trusted Internet Connection update

- NOAA offices in the Washington, DC, metro area and Boulder, CO, along with remote programs and facilities leveraging those locations, have been migrated to the new TICAP 2.0 network infrastructures. In a phased approach to the migration, these sites have initially been migrated to "Bypass" connections, which allowed the engineers to assess any problems with the new network infrastructure, independently of the security stack. Work continues to migrate these offices to "Inlined" status in order to fully take advantage of the security suite incorporated in the TICAP 2.0 design.
- With the migration of the Silver Spring headquarters came the addition of the National Weather Service's weather data dissemination traffic. As this data alone accounts for upwards of 20 Gbps of traffic out to the Internet, N-Wave quickly upgraded the capacity of the TICAP network infrastructure. N-Wave upgraded the DC metro infrastructure from 20 Gbps to 40 Gbps and the Denver TICAP infrastructure from 20 Gbps to 30 Gbps. Work continues to further increase the capacity of these TICAP locations to account for NOAA's ever-increasing data needs.

The OMAO ACIO, U.S. Public Health Service CDR Joseph Baczkowski, and his team

N-Wave was asked by the Office of Marine and Aviation Operations (OMAO) to provide connectivity for the Aircraft Operations Center (AOC), which was moving from MacDill Air Force Base to Lakeland, FL. N-Wave within four months, designed, procured, and installed network services to support the move, which was on a tight deadline to avoid impact to flight operations. N-Wave successfully completed the network install of a single 1 Gbps transport service along with a dedicated PRI to support the OMAO VOiP system and is in the process of planning for a second, diverse 1G along with local N-Wave Wireless network service.

As part of the continuing migration of the National Ocean Service offices in California to N-Wave, the Office of National Marine Sanctuaries office in Monterey, CA, was brought online via a 1 Gbps circuit to the N-Wave POP in Sunnyvale, CA.

To provide additional redundancy to the vital satellite operations at NOAA's National Environmental Satellite, Data, and Information Service (NESDIS) office located in Suitland, MD, a second aggregation router was installed. Programs that need full redundancy can connect to both aggregation routers and no longer be dependent on a single point of failure.

To meet the upcoming projected increase in data volumes produced by the Environmental Satellite Processing and Distribution System (ESPDS), for serving out the new generation NOAA satellites (GOES-R and JPSS-1) data to the mission partners as well as the public, N-Wave has increased the capacity for ESPDS at NSOF to two 20 Gbps Link Aggregated (LAG) links, allowing both inter-agency as well as external communications (via NOAA TICAPs).

As part of the ongoing effort to deploy the JEUNO (Joint EUMETSAT and NOAA) network infrastructure for exchanging NOAA and European weather satellite data, NOAA and EUMETSAT have provisioned a private network infrastructure to support the migration of existing EUMETSAT operational data flows for Metop satellite data as well as command and control. This effort is in preparation to migrate these operational data flows from the legacy commercial 10 Mbps Transatlantic Trunk link to the new JEUNO network infrastructure. Initial testing, baselining, and fail-over capability of the private JEUNO infrastructure has been successfully conducted, and EUMETSAT is currently going through their Operational Readiness Review efforts to declare the new JEUNO connectivity as operational.

As part of NESDIS deploying a Mission Science Network (MSN), which aims to provide an IT platform for delivering enterprise services to the National Centers for Environmental Information (NCEI) and to consolidate their data current centers, N-Wave has provisioned a 1 Gbps connection and a private routing instance for MSN at two locations: Center for Satellite Applications and Research (STAR) at College Park, MD, and NCEI-North Carolina (Asheville, NC).

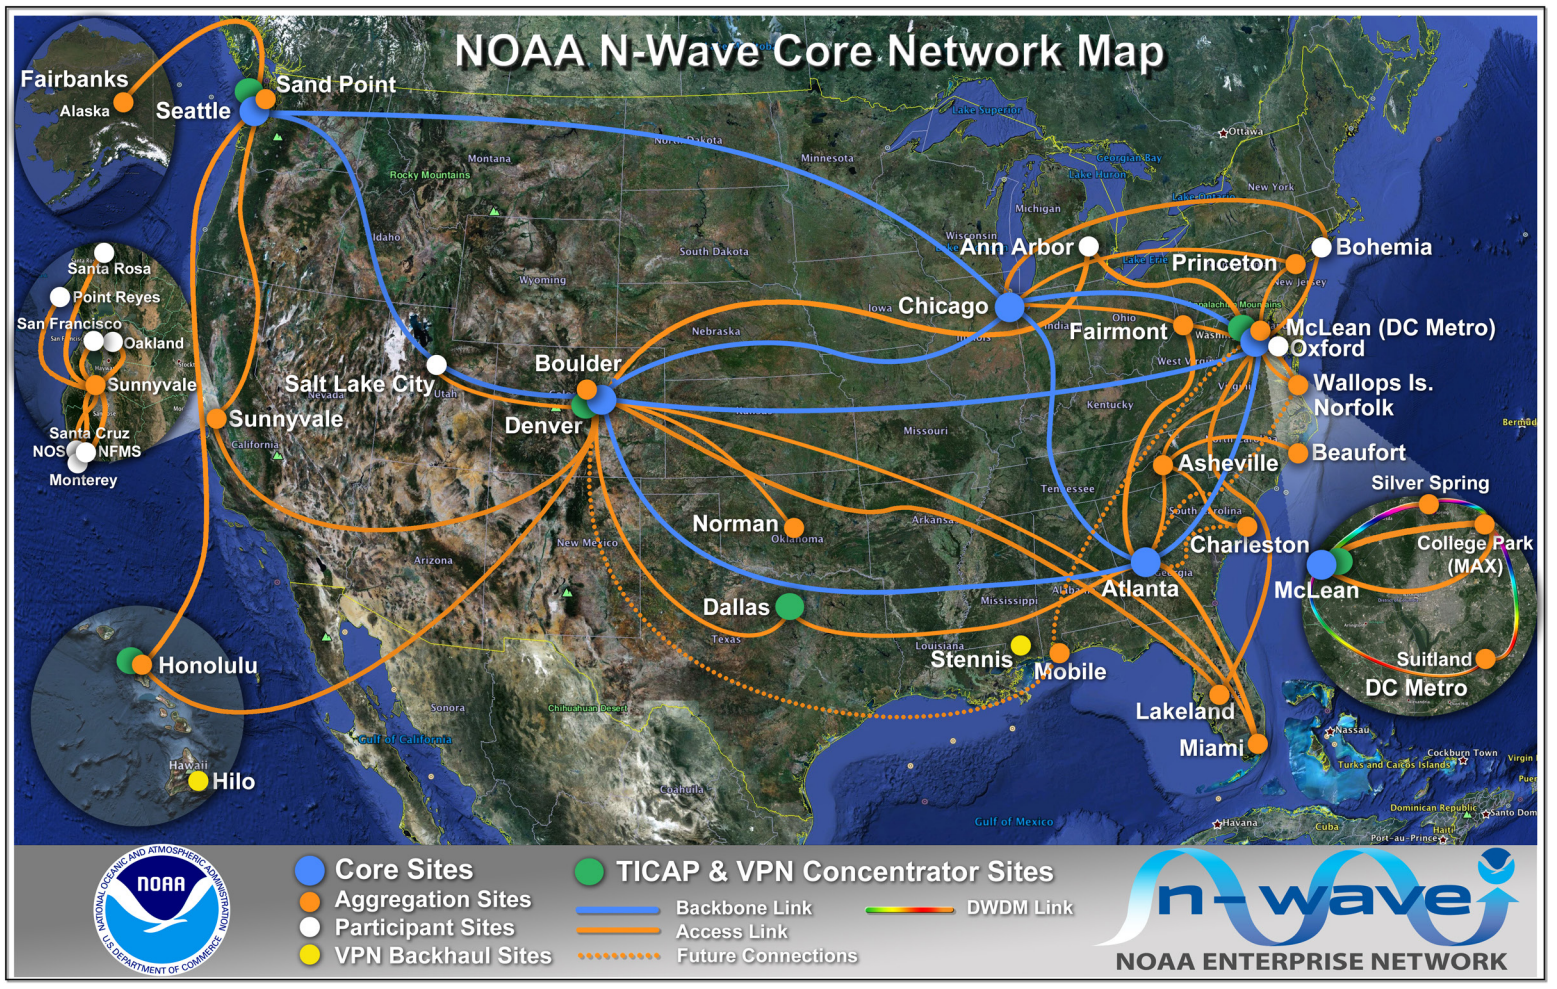

N-Wave Lab

Multiprotocol Label Switching (MPLS) service-provider networks have many customer benefits, but also much complexity. End users require high availability, so the network must be engineered at the physical and logical layers to provide resilience and fast failover.

The network's control plane is a set of complex routing protocols, distributed across geography at scale. Its forwarding plane is purpose-build silicon, potentially from different vendors that have to interoperate. MPLS provides a mechanism for an ever expanding set of services. No two service provider networks are the same, and predicting the interaction between existing features and new features can be challenging. Every network owner must manage and validate their unique combination of protocols, features, and configurations.

To manage this complexity, N-Wave uses an in-house network test lab. The lab is used to:

- Evaluate new approaches to serving customers
- Learn and test new technology
- Mitigate risk of change to the network when adding new technology
- Qualify new hardware and software
- Understand behaviors and duplicate problems outside of the production network

Most importantly, the lab is a place to model a sustainable network. Oftentimes, the urgency of work demands that an engineer make configurations or set knobs on production equipment without prior planning to get the deployment or problems of the day resolved. This can leave the network in a brittle state subject to breakage if others revisit the network, or when other changes occur. Planning in a lab enables us to avoid this and manage the complexity of a complicated internetwork. It helps ensure that, when we touch the production network, we get things working the first time we touch it, thereby minimizing the risk of network impact.

The N-Wave network lab is a model of the entire production N-Wave and X-Wave networks. (X-Wave is the external network that supports the Trusted Internet Connection Access Points [TICAP] and peers with the rest of the Internet.) At the physical layer, the network is based on four MX-series Juniper routers and two QFX-series Juniper switches. Using a virtualization feature called logical systems, each of the MX routers can support up to 15 "logical routers" virtually separate from the chassis. Using this strategy, we have built a model of the production N-Wave and X-Wave environments, including aggregate and core routers, customer connections, service provider connections, and Internet endpoints. The associated illustrations show both the physical elements of the lab as well as the network built using logical systems.

This model has been used over the past two years to support many N-Wave projects, including validation for network implementations and development of network architectures. Some of these projects include the development of the X-Wave architecture for TICAP and the MPLS-based multicast architecture to support NOAA's Okeanos Explorer's unique data flows from the ship. Any significant change or enhancement to the network should first be explored in the lab. While the lab simulation can never perfectly duplicate the production network, using the lab to develop new configurations greatly helps minimize risks of introducing new network elements and makes initial implementations run smoothly.

In general, lab testing falls into three categories: design exploration, configuration validation, and scale and performance testing. In addition to the gear used as a model of the production network, there is also "sandbox" gear for general purpose usage. These routers and switches are used for ad-hoc configuration validation and self-training.

Currently, the N-Wave lab is physically located in Boulder, CO, but we are making plans to logically extend the network presence to Silver Spring, MD, for lab availability, potentially using Virtual Private LAN Service (VPLS) on N-Wave. (For more on VPLS see the article in the November 2015 issue.) Future plans will move the lab and all distributed N-Wave engineers behind a new N-Wave LAN that spans Boulder and Silver Spring, which will make it more readily available to users. Most of the testing today is command-line oriented, but the lab should be expanded to support access to web interfaces.

As N-Wave grows, there are other future needs for the lab to promote sustainability and make it a permanent part of our engineering process. Some of these needs involve network management. As more people use the lab, it will tend towards disorganization, and there is a greater need for active lab management. This human oversight would be assisted by tools to manage and reset configurations.

There are also many equipment needs for the lab. Today, lab gear comprises spare or "demo" pool-type hardware, but there is a need for permanent devices. Upgrades are also needed for lab firewalls to get SSL-VPN functionality, and for a higher-capacity terminal server.

There is currently minimal server infrastructure in the lab, but more will be necessary for explorations into network function virtualization (NFV). NFV will both enable scaling in the lab, possibly allowing us to rely less on logical systems, and also allow us to begin validation for potential production network usage. These servers would also enable local management services for automation. Finally, there is a need for test sets to generate network traffic. This function is critical for scale and performance testing and validation.

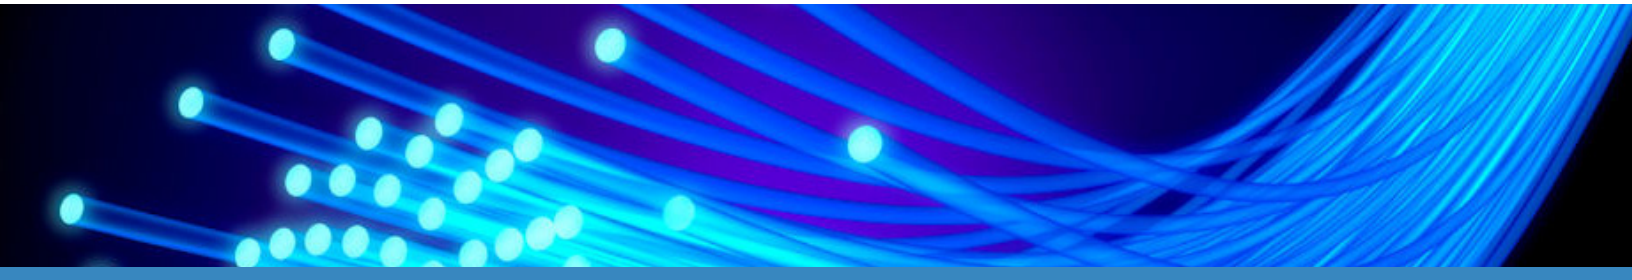

N-Wave Lab

with Logical systems

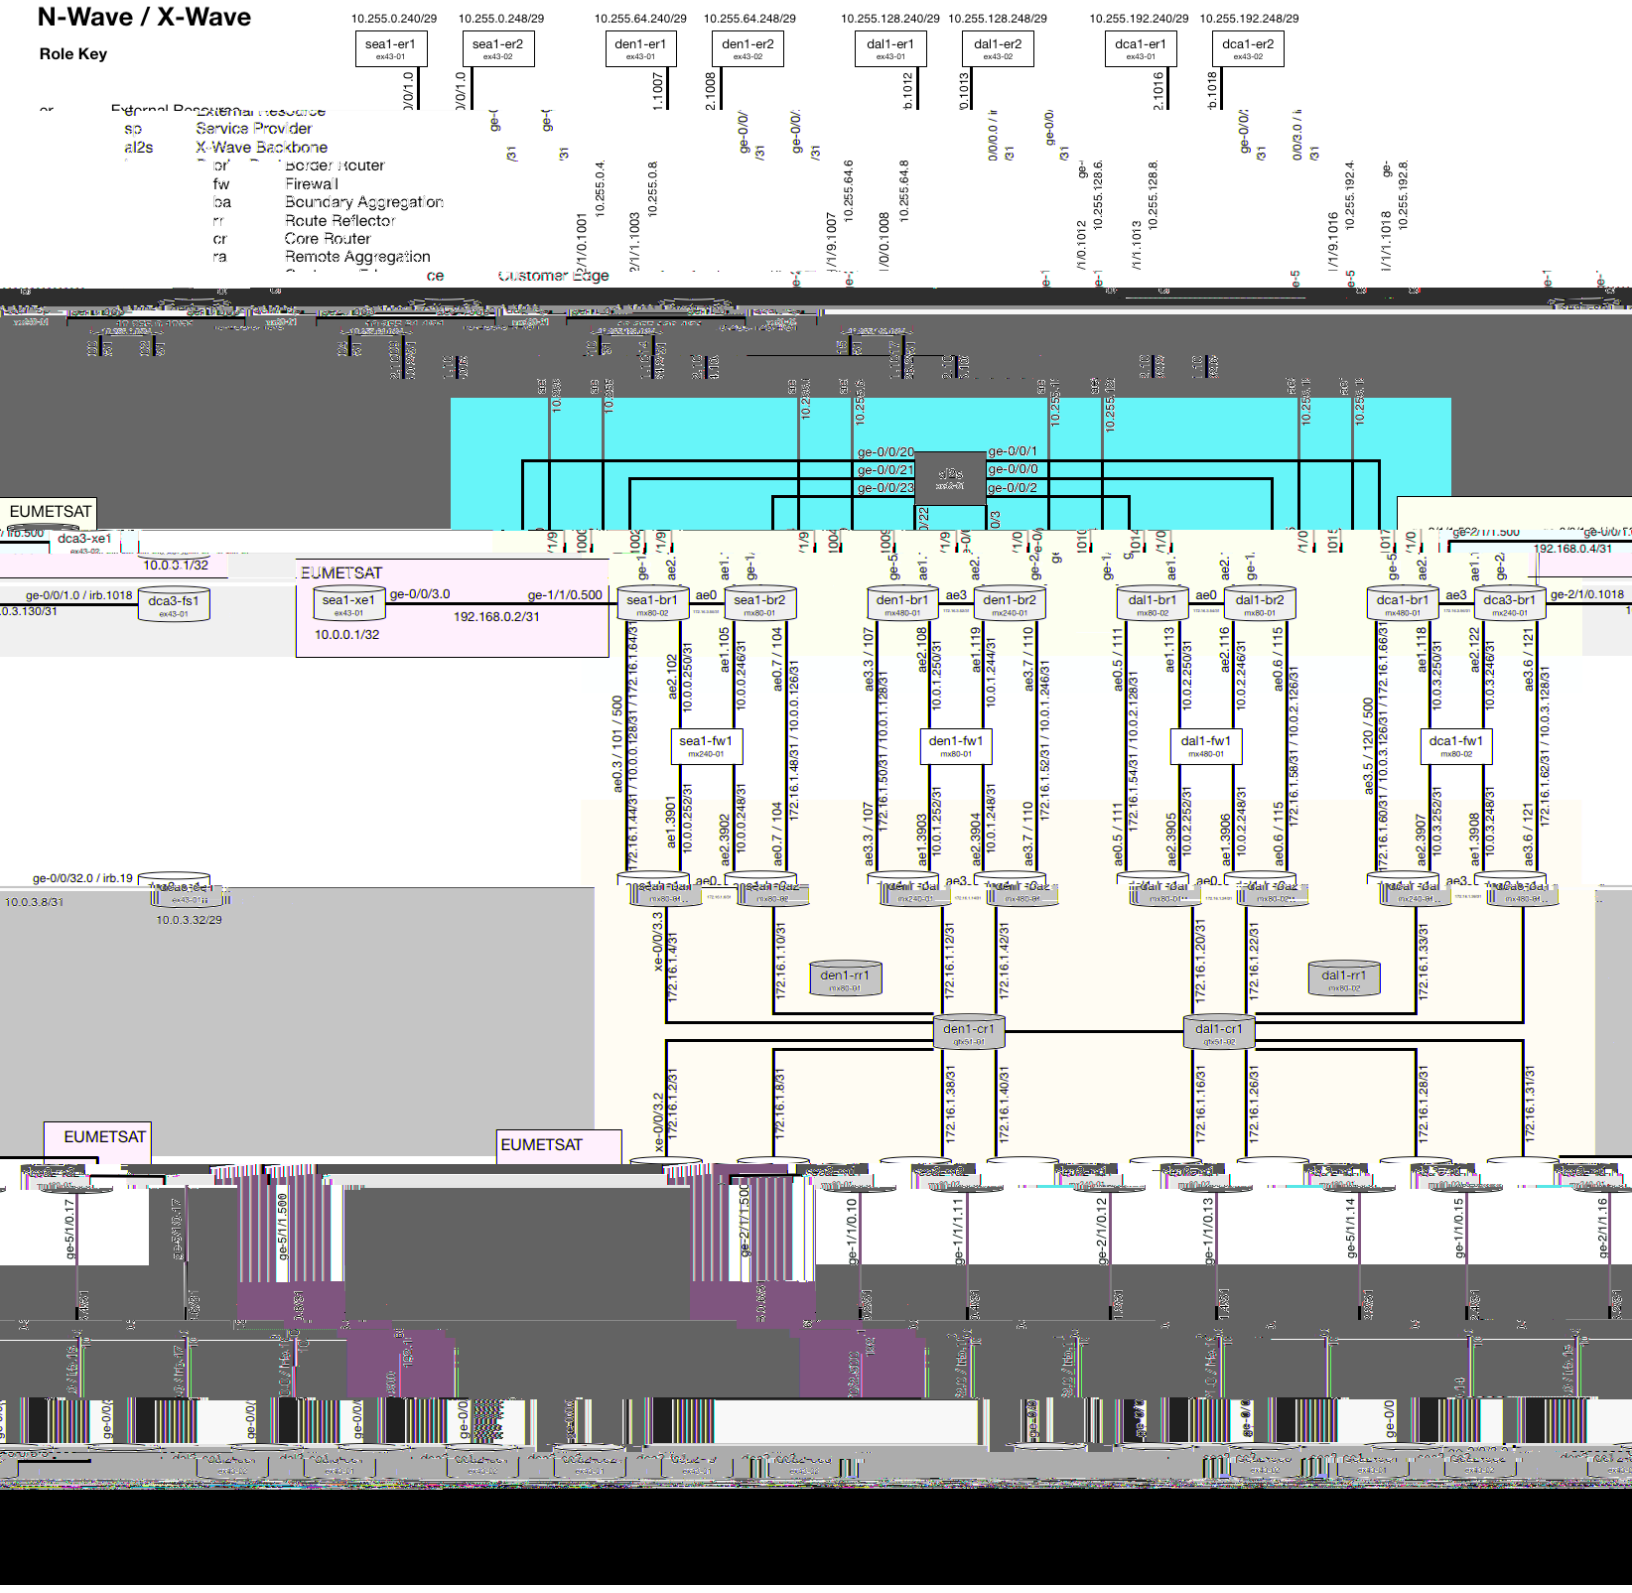

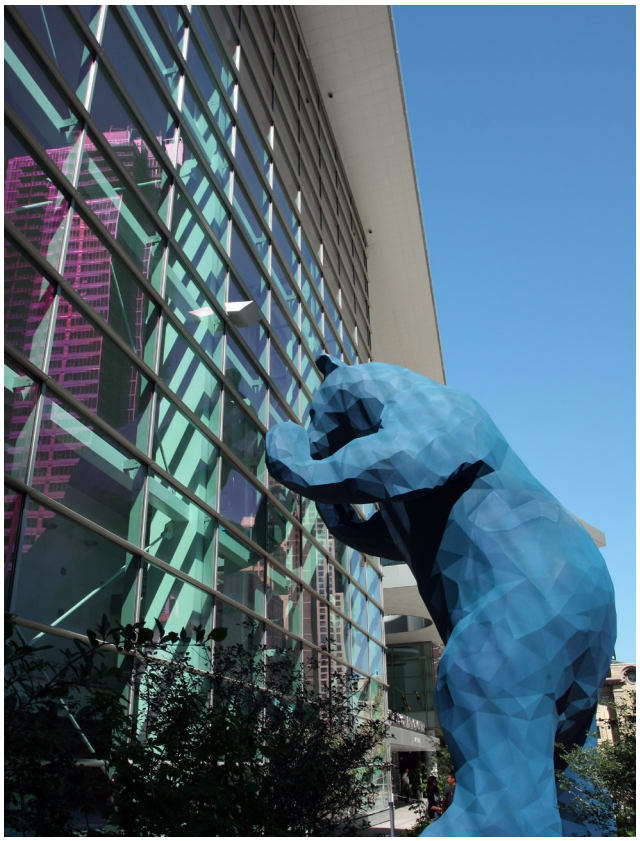

SC17 Denver, CO hpc connects.

This year SC, formerly the Supercomputing Conference, returns to the Colorado Convention Center in Denver, CO.

Once again, many of the world's leading researchers and practitioners in high performance computing, high speed networks, and large data transfers will gather to present peer reviewed papers, demonstrate their latest and best experiments and prototypes, and compare notes. The network that supports the entire conference–i.e. meeting halls, conference rooms, pervasive wireless, exhibit area, and many bandwidth intensive demonstrations–is called SCinet and will, for a few days, be the biggest in the world. At last count, the volunteer team that builds SCinet anticipates over thirty 100 Gbps circuits arriving at the conference. These include circuits from Singapore, South America, and Europe, as well as from various location within the United States.

Several members of the N-Wave team will be attending: Don Arnold, Alex Hsia, Jerry Janssen, Paul Love, Dave Mauro, and Rob Sears. Additionally this year Matt Smith is co-chairing the SCinet WAN Team.

n-wave 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
Samantha Stalion, Editor

U. S. Department of Commerce, NOAA
325 Broadway, NWAVE
DSRC - GB306
Boulder, CO 80305-3337