ISSUE
08
May
2016

N-Wave News

NOAA N-Wave

N-Wave is a program and Enterprise Science Network that supports both operations and research, enabling NOAA's mission of science, stewardship and service through highly available, secure, high speed networking services.

Mission

N-Wave is committed to providing innovative networking capabilities with integrity, transparency, and flexibility, to enable NOAA's missions through the implementaton of:
-Quality, advanced high-speed connectivity both internally and externally to NOAA.
-Secure, private, flexible, hight-bandwith virtual circuit capabilities.
-Retention and recruitment of exceptional operations and engineering staff.

Our Vision

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

this issue

From the N-Wave Science Network Manager P. 1
Engineered for Availability P. 2
Enterprise N-Wave Security P. 6
Update - Network Upgrade for NOAA's Mauna Loa Observatory P. 9
Network Changes & New Participants P.12
N-Wave Service Reliability P.13
X-Wave & TICAPS : Status Update P.15

From the N-Wave Network Manager

Robert Sears

To quote Tennessee Williams, "time is the longest distance between two places", and for N-Wave time has been an ally. It has allowed for the collection of statistical data including environmental, monitoring, measurement, traffic analysis and other performance metrics, all substantiating the posture of a highly available, reliable and secure infrastructure spanning the spectrum from engineering and deployment to full operations and enhanced service delivery.

This edition of the N-Wave newsletter focuses on the engineering and operations that deliver a highly robust network infrastructure and associated networking services, and you won't just have to take our word for it. Pulling from the myriad of data metrics collected, you will see how the N-Wave team develops "service level actuals" to provide its customer base with service level agreements backed by historical data. Be it hardware, circuits, underlying optical transport or support systems, N-Wave has a highly visible and monitored infrastructure from customer edge to core to national transport, TICAP, and provider services.

As NOAA is an agency of science, stewardship and service, N-Wave's approach to service level agreements through service level actuals provides more in-depth, fully visible options to the N-Wave customer over traditional provider agreements that may not contain the historical and in-depth metrics of the end-to-end service.

When we proudly state that the N-Wave backbone has had 100% availability since inception in late 2010, this is a testament to the engineering, management and operations of the infrastructure, which once again, can be visualized through collected data. The N-Wave network operations center (NOC) is the team that binds together and brings full circle the highly available, reliable and visible posture of the network through true 24x7x365 active monitoring, measurement, and service support with tier-1 through tier-3 engineering, visualization, standardized troubleshooting tool sets and change tracking.

Although the focus of this newsletter is service-centric, traditional articles like project updates and the service desk ticket report are included. Please enjoy!

Team picture from the N-Wave Spring 2016 engineering summit held at the Indiana University Bloomington Campus. Team members shown (not all in attendance) represent NOAA engineers; GlobalNOC service desk, tier-2 and tier-3 engineering, systems, monitoring, measurement and visualization.

Engineered for Availability

The N-Wave network provides network transport for some of the most demanding science needs within NOAA. In addition to requiring low-latency, high-bandwidth capability between endpoints, N-Wave users also require a robust path that provides greater resource availability than they would find on the commercial Internet. From the beginning, the N-Wave network has been engineered with these high-resiliency capabilities in mind. These capabilities are supported by three primary pillars of network design: redundant architecture, high-access monitoring and engineering, and strong change management processes.

Architecture

The easiest way to avoid a disaster is to prepare for it ahead of time. This design tenet pervades all layers of the N-Wave network and has driven the network design from the start. Where feasible, all aspects of the network have been pre-provisioned to support individual component failure and obscure outages from the network users. The diagram below illustrates the philosophy behind the design of each of N-Wave's major core nodes. Their locations and redundant interconnections can be seen in the current N-Wave map on page 8.

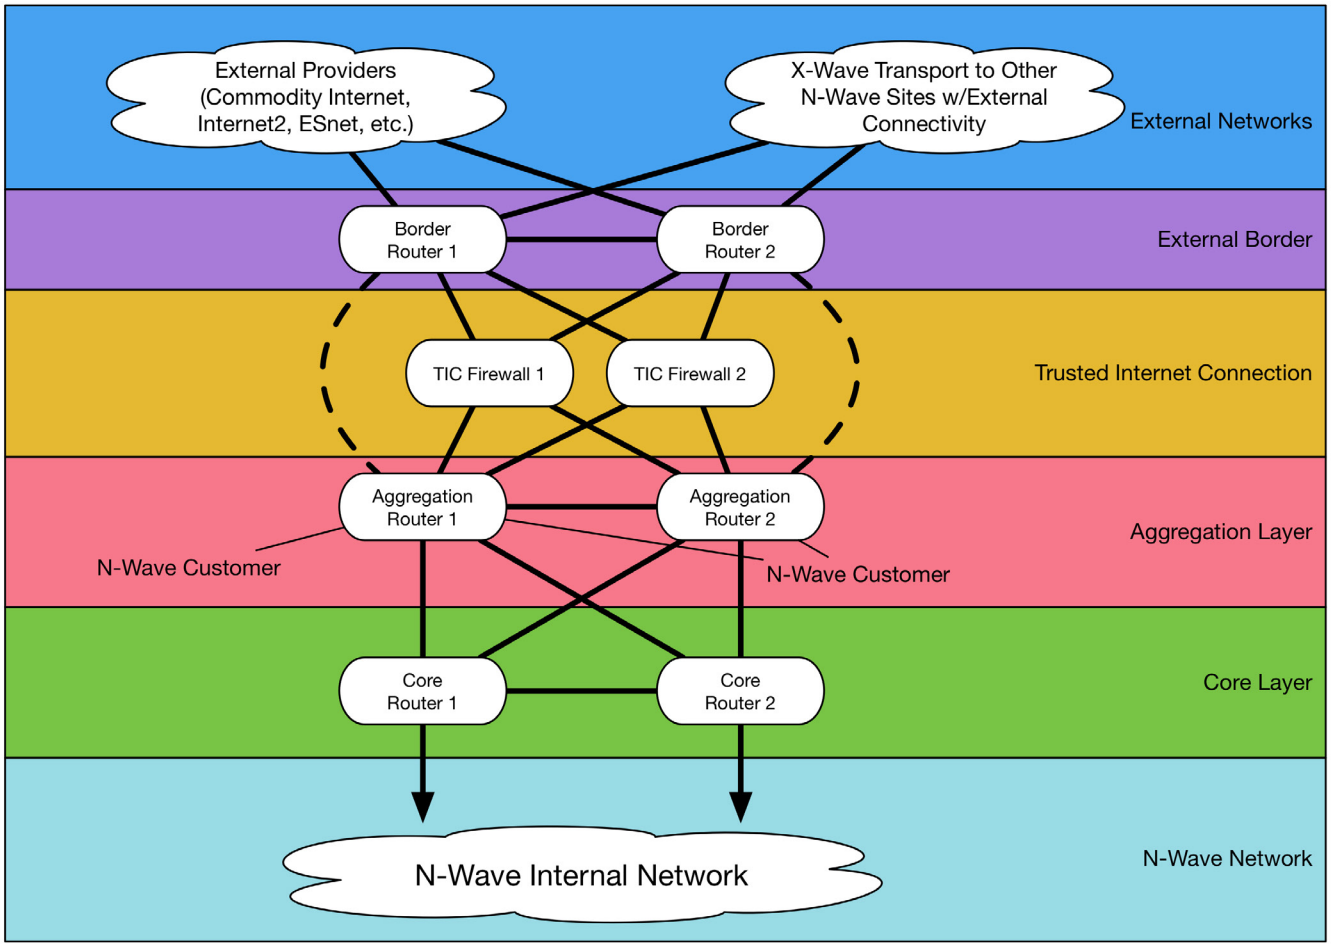

Each layer of the network, from the internal N-Wave network interconnects all the way to the external providers, has been designed with redundancy in mind. When possible, each network device in the diagram is outfitted with redundant power supplies and network management systems within each box. This first line of protection helps prevent individual device failure. Should a device fail, the system is designed to automatically failover to a second box to restore services within seconds. N-Wave participants have the option of interconnecting with multiple network devices, in different cities where

feasible, to provide even more redundancy along geographical boundaries. This helps ensure that the N-Wave users at a site will never be down for longer than a few seconds in the worst case. All of the network components in each city work in concert to maintain constant connectivity to each N-Wave participant network.

Monitoring and Engineering

A robust infrastructure is nothing without a system that periodically checks for proper function and reachability. The N-Wave network operations center (NOC) systems constantly monitor all network devices for service degradation or failure. These monitoring checks are custom-built to meet the high availability demands of the N-Wave network, and are often more persistent than network monitoring systems in commercial networks. The N-Wave NOC monitoring systems provide the first indication of a device malfunction, allowing the dedicated N-Wave engineering staff to react immediately to an issue. This constant attention supports the goal of a highly-available network.

Change Management

Networks change and evolve over time to meet the needs of the network users. To support this necessary evolution, the N-Wave network has embraced a very deliberate and open change management policy. Each change to the network that will impact a user is discussed in a weekly call that is open to all interested parties. Once changes are approved, they are tracked and reported, to maintain an auditable history of network changes. These processes maintain network stability as the network grows over time.

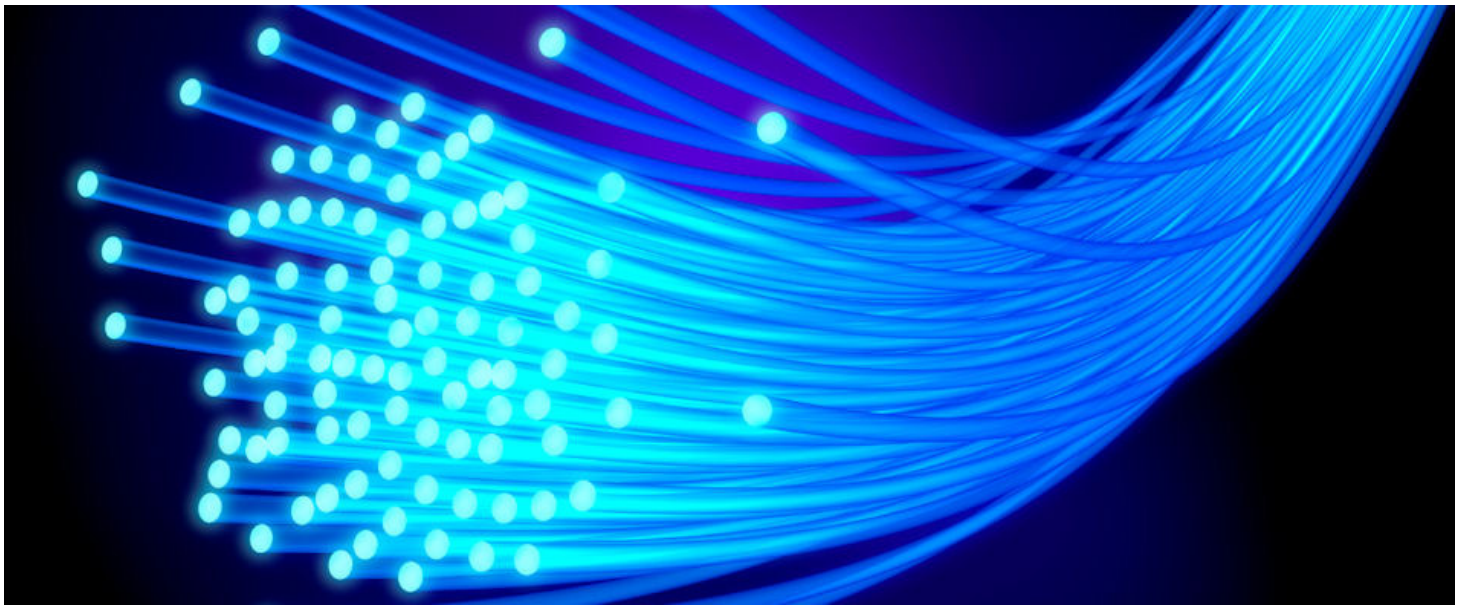

N-Wave Network and Performance Metrics

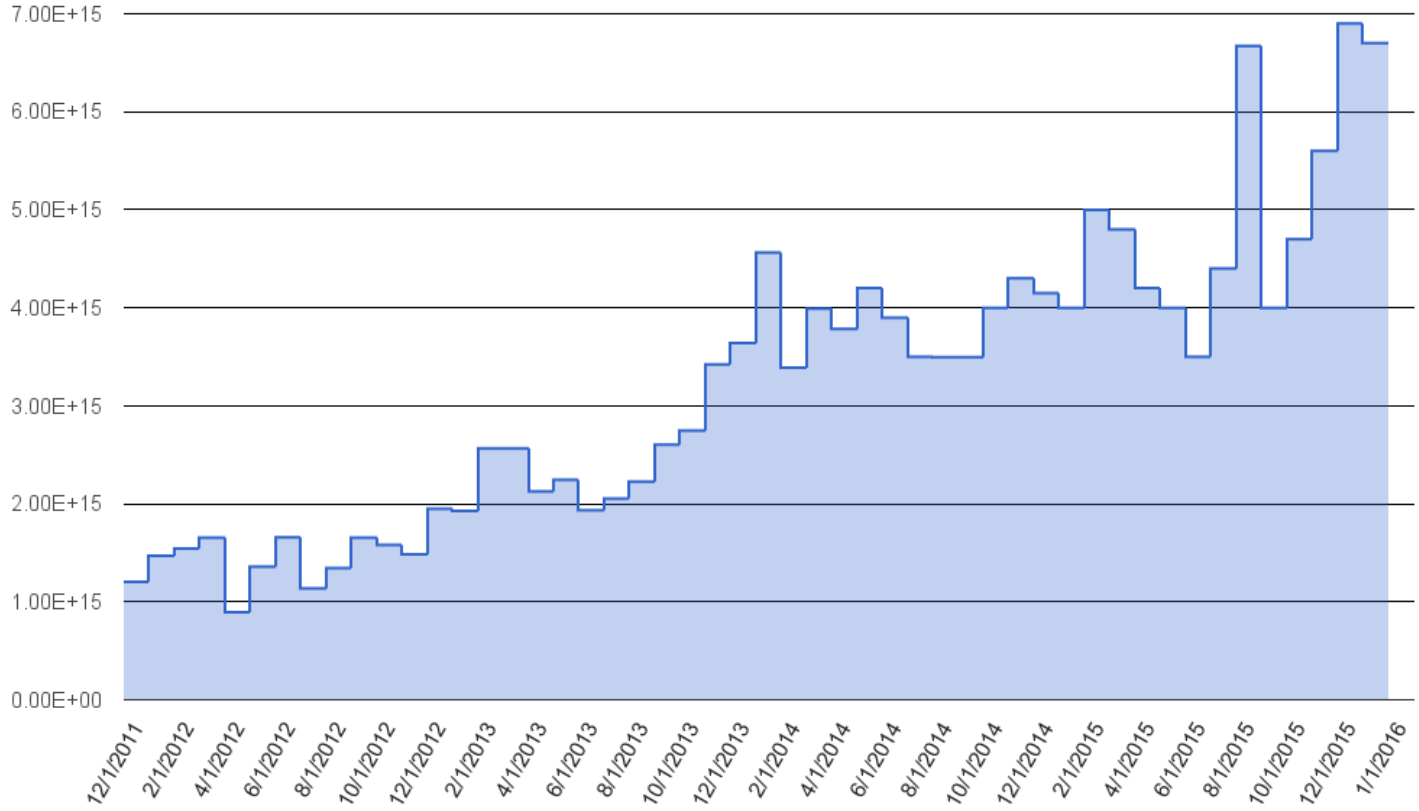

Welcome to the spring 2016 edition of our network performance metrics discussion. As shown in the figure above (total N-Wave traffic volumes in TBytes/month), N-Wave continues to transport more traffic, due to the addition of new users and increased high-performance computing (HPC) utilization.

The Research and Development High Performance Computing System (RDHPCS) and the Weather and Climate Operational Supercomputing System (WCOSS) continue to be high volume users, especially between College Park, MD; Fairmont, WV; Princeton, NJ and Boulder, CO.

To meet the needs of HPC efforts, upgraded network capacity is being put into place in the Washington, D.C. area specifically for WCOSS traffic to the Fairmont archive. The N-Wave D.C.-area metro network offers capacity "on demand" with minimal effort needed to turn-up 10 Gigabit per second (Gbps) channels. In order to meet additional programmatic demands, N-Wave and RDHPCS are also increasing their peering interfaces from 10 Gbps to 20 Gbps at Fairmont and College Park.
The D.C.-area network will be expanding to provide capacity to various NOAA missions, such as acting as the backbone for TICAP 2.0. The metro fiber network NOAA deployed will play a crucial role in integration and optimization of the NOAA networks in the D.C. area.

The N-Wave team is currently studying options for increasing capacity over the wide-area backbone. Today, the backbone utilizes 10 Gbps links between the various core nodes. Multiple 10 Gbps links, as well as 40 or 100 Gbps links, are on the list of options. These upgrades, as well as new routing technologies, will meet the growing demand of NOAA's networking needs for the next few years, and bring new services such as Traffic Engineering and Software Defined Networking to N-Wave.

IdeaData TOTAL NETWORK VISIBILITY

N-Wave Monthly CostCentre Destination Traffic

Period From: Wed 2016-01-06 00:00:00 UTC
To: Fri 2016-02-05 23:59:59 UTC

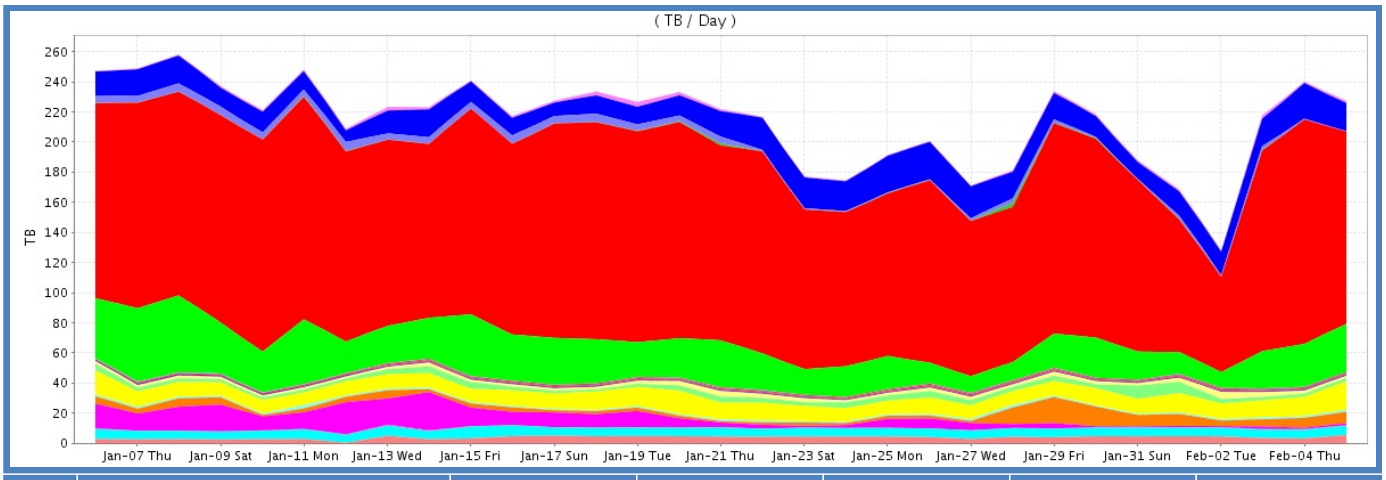

No.	Destination CostCentre	TB sum	TB avg	TB max	G.Pkts sum	M.Flows sum
1	RDHPCS-FAIR	3,879.084	125.132	149.007	567.533	8.602
2	RDHPCS-BOUL	825.770	26.638	50.937	185.985	2.826
3	Undefined	505.942	16.321	24.890	418.254	44.174
4	NOAA-Boulder	340.348	10.979	18.917	258.884	12.544
5	NCEP-WCOSS-ORLANDO	242.708	7.829	25.354	117.799	2.950
6	CLASS-BLDR	181.504	5.855	7.860	123.325	1.773
7	NCEP-WCOSS-RESTON	136.686	4.409	17.278	79.075	1.271
8	CLASS-ASH	120.867	3.899	5.439	82.220	0.964
9	NOAA-DC	99.014	3.194	8.812	112.567	9.341
10	RDHPCS-ORNL	91.475	2.951	5.962	109.163	0.839
11	NSSL	63.728	2.056	4.609	47.555	4.148
12	WISC-EDU	36.821	1.188	3.195	26.717	0.638
13	NESDIS-NCDC	32.424	1.046	2.251	95.508	10.624
14	NWAVE-CORE	26.019	0.839	1.118	35.765	4.950
15	NWAVE-CHI	15.857	0.512	0.669	15.561	0.118
16	NWAVE-DEN	15.693	0.506	0.697	15.294	0.124
17	RDHPCS-GFDL	11.707	0.378	2.741	4.912	0.118
18	NWAVE-ATL	11.610	0.375	0.452	11.486	0.174
19	NWS-ROC	8.235	0.266	0.428	8.692	1.529
20	GOES-R NSOF	6.039	0.195	0.538	10.173	0.346
	Total (39)	6,651.531	214.568	258.092	2,326.468	108.053

Note: TB = 1000^4 Bytes, G.Pkts = 1000^3 Packets, M.Flows = 1000^2 Flows

If you would like to subscribe to the outbound traffic information, please go to: noc.nwave.noaa.gov to submit a service inquiry and fill in a request.

N-Wave Enterprise Security

N-Wave is designed, engineered, deployed and managed from customer edges to the N-Wave provider core, including all associated transport services. This allows for complete visibility and makes N-Wave NOAA's only end-to-end fully FISMA-certified Network Infrastructure.

N-Wave is managed end-to-end for the following enterprise-level security features:

- Direct visibility into layer-1 transport infrastructures, including fiber maps for fiber path diversity and high availability assurance.
- Direct engineering relationships and membership on technical and administrative boards for partnering regional, and national optical networks.
- All server and network device configurations across the national infrastructure are managed centrally via replicated repositories.
- Changes are detected through automated processes and tied to Configuration Management tickets in the N-Wave ticketing system.
- All devices are continuously monitored for potential vulnerabilities and unauthorized changes. Any risks identified are immediately escalated to N-Wave management.
- All devices are provisioned from a known-secure template. After provisioning, configuration settings are regularly checked against CIS recommended baselines for compliance. Any deviation from established baselines requires justification and mitigation.

Interconnections: N-Wave performs thorough engineering analysis for any new connections to customer sites. A detailed Interface Control Document (ICD) describes the technical details of the connection, as N-Wave engineers work with the customer to refine requirements and plan for future needs. The ICD serves as the technical basis of our Interconnection Security Agreement (ISA), which defines responsibilities and authorizes the interconnection.

Trusted Internet Connections (TIC): N-Wave provides 5 TIC Access Points (TICAPs):

- Dallas, TX
- Denver, CO
- Honolulu, HI
- Seattle, WA
- Washington, D.C.

These sites are authorized by the Department of Homeland Security (DHS) to serve as NOAA's only official Internet access path for all traffic entering or leaving NOAA's internal networks. The aggregation of all Internet traffic through managed access points provides NOAA with comprehensive traffic inspection and attack response, as well as agency-wide oversight of all public network traffic.

Monitoring: N-Wave continuously monitors its network and interconnections for connectivity, stability, and performance issues, including TICAP network infrastructures. N-Wave maintains a detailed history of traffic usage for both capacity planning and anomaly detection.

N-Wave NOC Tickets Report: September 2015-February 2016

The N-Wave ticket report tracks various ticket types that are used to support the N-Wave network. In this issue, we discuss one of these ticket types and its associated workflows: the Problem/Request ticket.

The Problem/Request ticket category is used by the N-Wave NOC as a general catch-all to track customer and site inquiries regarding things that do not impact network service or represent an outage or maintenance event. These general tickets are non-specific and do not conform to the standard Outage, Maintenance, Service Inquiry, Account Request, or Security ticket types.

Problem/Request tickets may be the result of a suspected issue or service inquiry that was submitted to the N-Wave NOC by email or phone call, rather than the Service Inquiry or Trouble Ticket web forms. The N-Wave NOC recommends using the Trouble Ticket web form, and including the requested delivery date so that an appropriate priority can be assigned to facilitate timely resolution. Problem/Request tickets are triaged and processed with a default “Elevated” priority, and are assigned for review and action within 3 business days, unless otherwise specified.

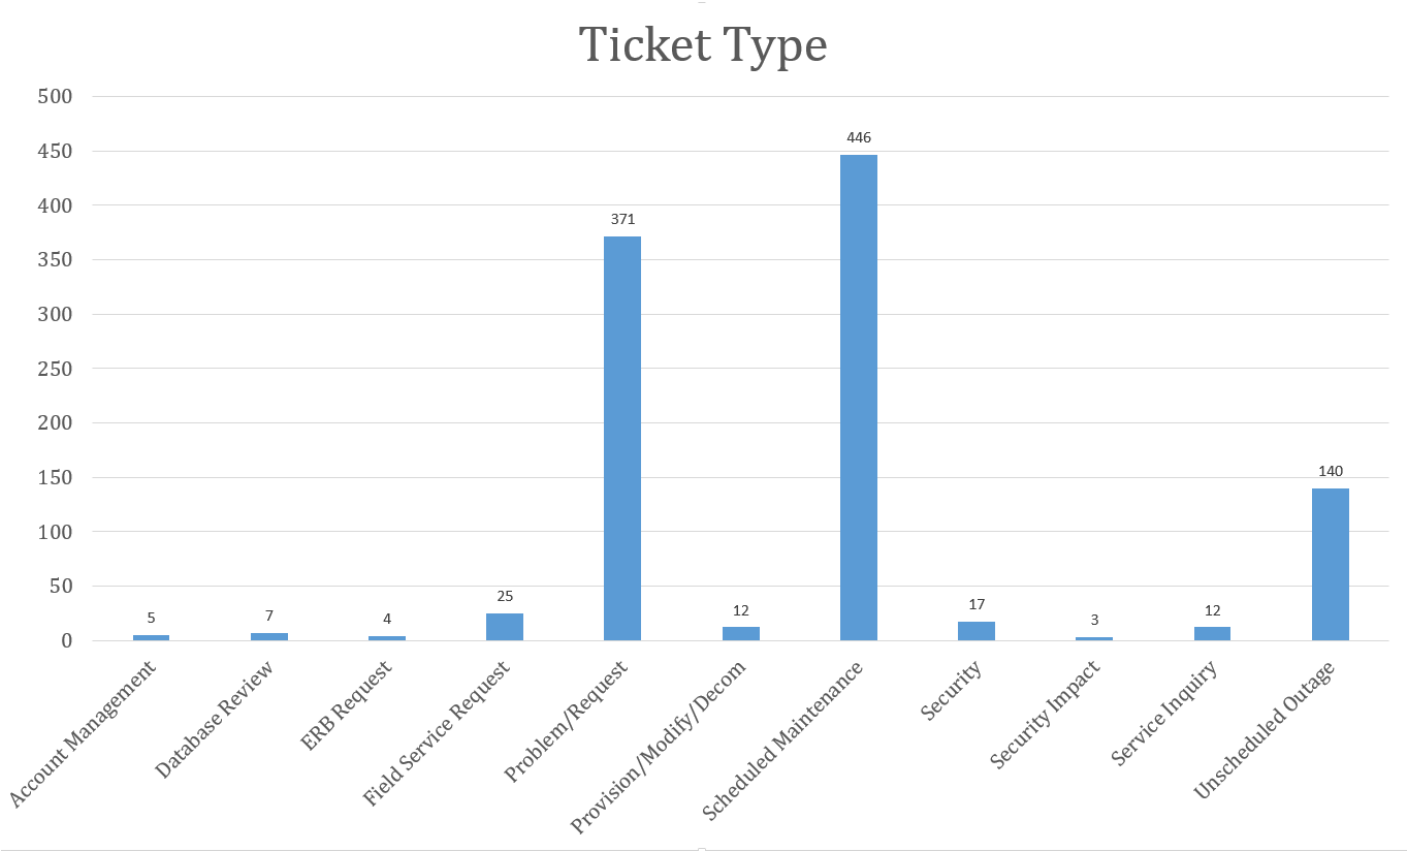

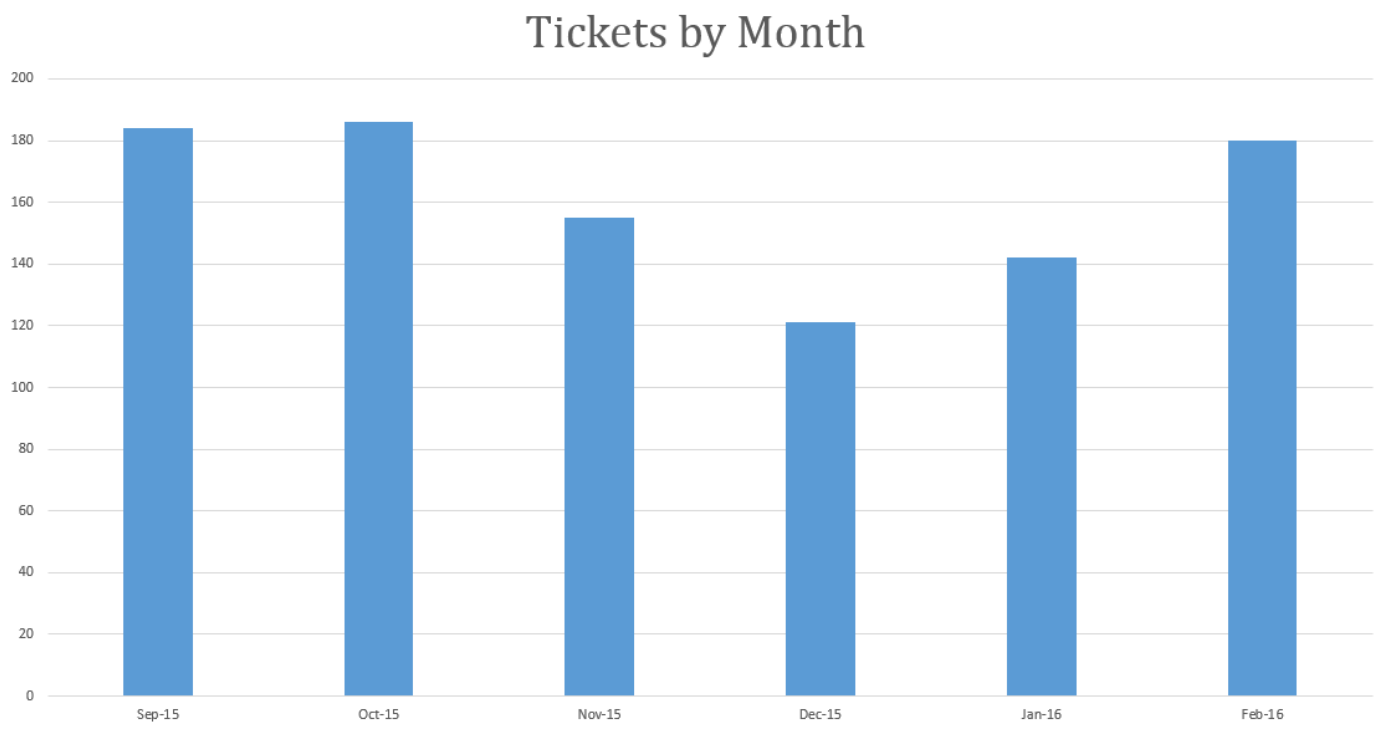
Tickets by Month
200
180
160
140
120
100
80
60
40
20
0
Sep-15
Oct-15
Nov-15
Dec-15
Jan-16
Feb-16

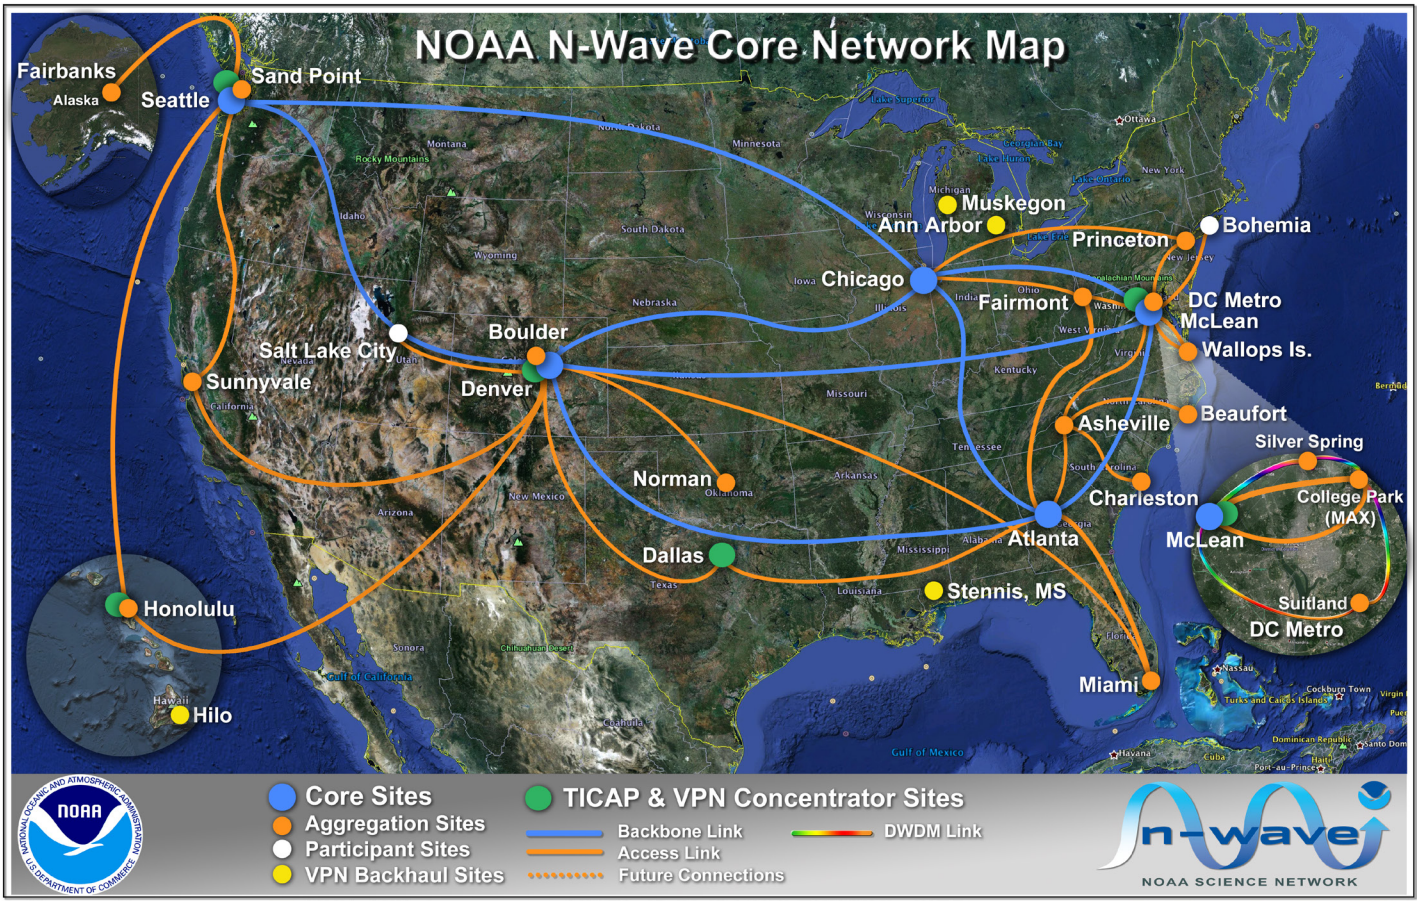
NOAA N-Wave Core Network Map
Fairbanks
Alaska
Seattle
Sand Point
Muskegon
Ann Arbor
Princeton
Bohemia
Chicago
Fairmont
DC Metro
McLean
Wallops Is.
Salt Lake City
Boulder
Denver
Sunnyvale
Asheville
Beaufort
Silver Spring
Norman
Charleston
College Park
(MAX)
McLean
Atlanta
Dallas
Stennis, MS
Suitland
DC Metro
Honolulu
Miami
Hilo
Core Sites
Aggregation Sites
Participant Sites
VPN Backhaul Sites
TICAP & VPN Concentrator Sites
Backbone Link
Access Link
Future Connections
DWDM Link
NOAA SCIENCE NETWORK

Update - Network Upgrade for NOAA's Mauna Loa Observatory

The Mauna Loa Observatory (MLO) facility is located at 3440m elevation on the north flank of the island of Hawaii's Mauna Loa volcano. Established in 1957, MLO has grown to become a premier long-term atmospheric monitoring facility, and is the site where the steadily-increasing concentrations of global atmospheric carbon dioxide were first determined. The facility currently relies on a 10 Mbps commercial internet service provided by Hawaiian Telcom via a microwave connection. This connection was upgraded last year from the previous 5 Mbps connection after OAR's Global Monitoring Division (GMD) took over financial responsibility for the connection from the Academia Sinica Institute of Astronomy and Astrophysics (ASIAA).

The Mauna Loa facility is getting a network upgrade as part of the National Science Foundation's (NSF) solicitation for "Campus Cyberinfrastructure – Data, Networking, and Innovation (CC*DNI) - Instrument Networking". A successful proposal for the MLO Science Instruments was submitted by the University of Hawaii's Institute for Astronomy (IfA) with support from the Global Monitoring Division and the National Center for Atmospheric Research (NCAR).

The proposal will increase aggregate bandwidth to the site from 160 Mbps to 310 Mbps, and finally to 378 Mbps by 2017. It will also provide redundant network routes to Internet and Research and Education (R&E) networks via UHNet. To do this, two carrier-grade Aviat IRU600 radios will be bonded together at each end of the link. One pair of radios and antennas will be located at MLO while the other pair, along with a router, will be located at a line of sight distance of 24 kilometers on Humuula, a peak in the saddle between Mauna Loa and Mauna Kea, where UHNet has a microwave tower for their statewide network. UXA6-U57AC2HG-C six-foot antennas are required at each end to meet the necessary signal levels. Each antenna has a 6 GHz 30 MHz channel at 189 Mbps, yielding a maximum of 378 Mbps between MLO and Humuula. The radios at MLO will connect to a Cisco 3560 Layer 3 switch. NOAA will install new fiber to the Uchida/Shop building, connecting to the 3560 to pass traffic through the switch to reach Humuula. The radios at Humuula will connect to a new Cisco ASR1001 router, then connect to an existing UHNet OC-3 microwave backhaul link to UHNet, Internet and R&E networks.

The existing IfA 50 Mbps 802.11n Wi-Fi path will remain in operation and serve as a backup route independent of the OC-3 digital microwave network. When UH is able to extend fiber to the Humuula microwave tower, the OC-3 will become the backup route and the fiber the primary route to UHNet, thereby enabling the use of full radio bandwidth of the 378 Mbps.

UH Information Technology Service (ITS) and IfA staff will provide network management and monitoring on a 24/7/365 basis. They will rely on "smart network" tools such as Nagios, Cacti and perfSONAR to support continuous monitoring.

The current and proposed network diagrams for this proposal (seen below) help explain the infrastructure changes.

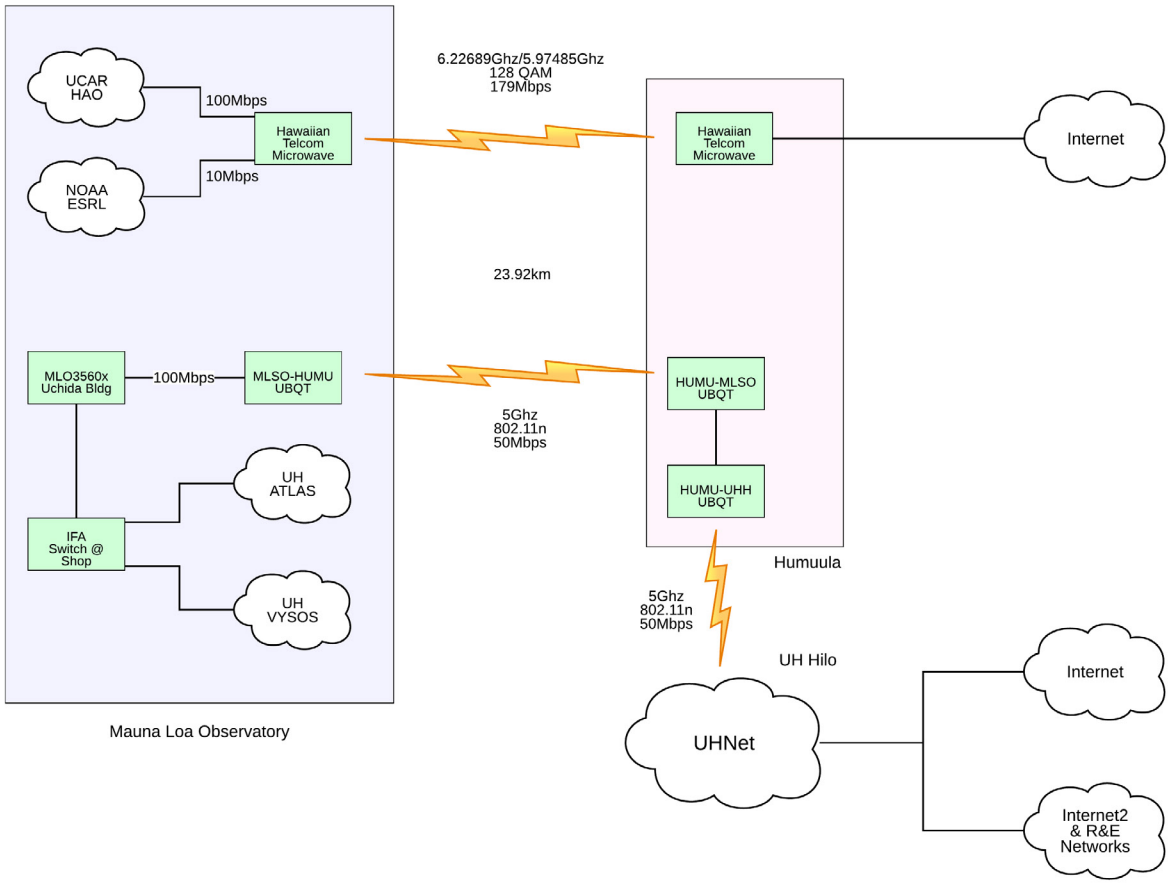

Figure 1 UH/IfA MLO Current

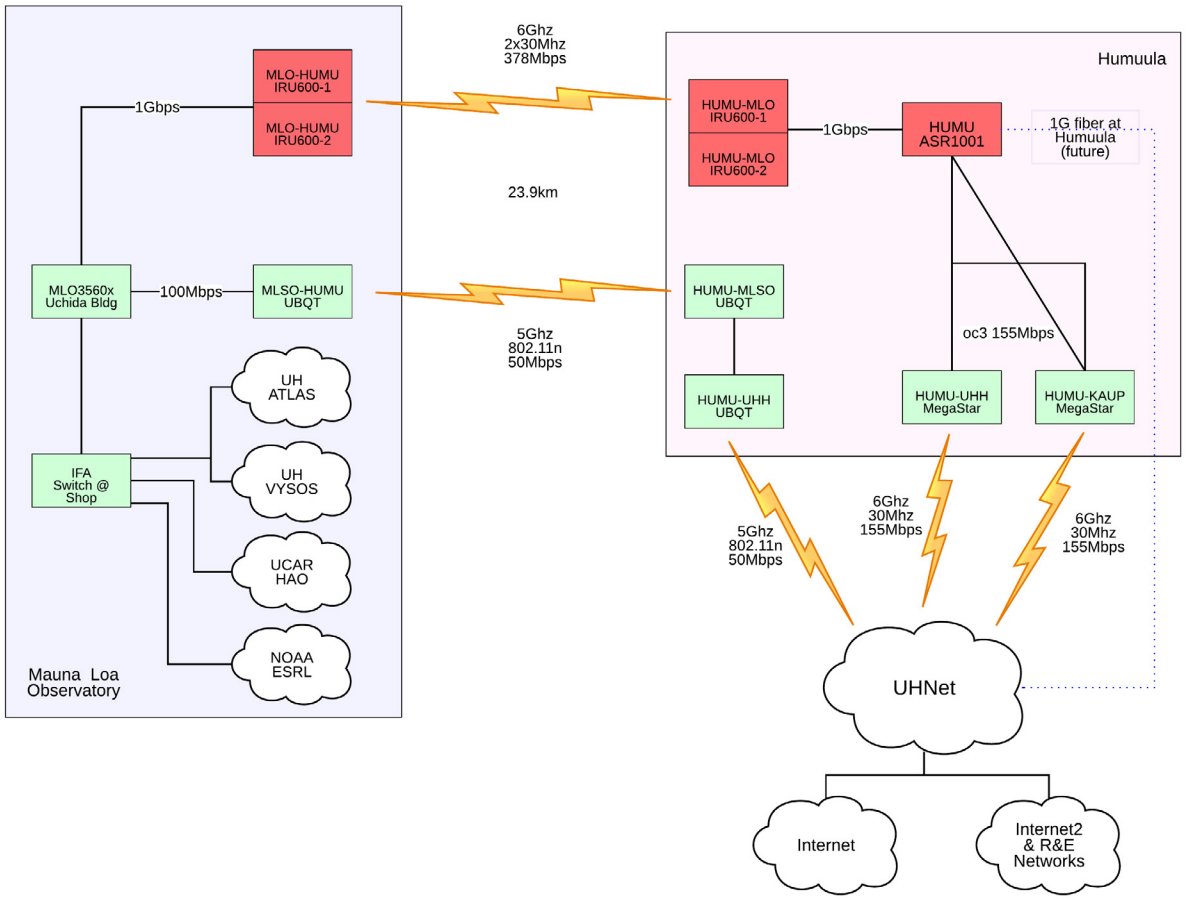

Figure 2 UH/IfA Proposed

Brian Vasel, Director of Observatories for OAR/GMD said “This upgrade will benefit a number of important scientific programs and services at the site, including the delivery of new real- and near-real-time solar data products from the Mauna Loa Solar Observatory (MLSO) to improve space weather forecasts.”

The IfA received the NSF award* September 1 2015, with GMD, IfA and the University Corporation for Atmospheric Research (which manages NCAR) as the primary implementation partners. The first Mauna Loa Network (MLONET) Policy/Technical Committee (MPC/MTC) meeting was held January 15, 2016 with representation from all three organizations. The radios, network equipment and installation services are currently being procured by IfA with the intent to install in May 2016. As the infrastructure in the saddle is upgraded to support higher bandwidth, the full capabilities of the new link will be realized.

*This network upgrade is made possible by NSF award ACI-1541471 CC*DNI Instrument: High Performance Reliable Network Access to Mauna Loa Observatory Science Instruments.

Network Changes and New Participants

N-Wave expands to California: In partnership with the Corporation for Education Network Initiatives in California (CENIC), N-Wave has expanded its reach to California. A new aggregation router has been installed in Sunnyvale, and work continues to connect several National Ocean Service and National Marine Fisheries Service sites back to N-Wave, leveraging CENIC's partnership with AT&T for ASE (metro-Ethernet) service. N-Wave has been able to utilize our membership with the Western Region Network - a multi-state partnership of CENIC, PNWGP, FRGP, the University of Hawaii and the University of New Mexico - for backbone connections to N-Wave core sites in Seattle and Denver.

Hawaii bandwidth upgrades: Working with the University of Hawaii, N-Wave has upgraded our backbone connection to Hawaii from a single OC3 to two redundant 1 Gigabit/s Ethernet connections. Hawaii is now dual-homed to the Seattle and Denver Core POPs.

Seattle and Washington, D.C. TICAP/X-Wave infrastructure: NOAA's third and fourth planned TICAP locations have been installed. N-Wave engineers are working with engineers from the NOAA Cyber Security Center (NCSC) to bring these TICAPs online and prepare them for customer migrations.

NOAA Search and Rescue Satellite Aided Tracking (SARSAT): A new VRF on N-Wave has been provisioned for the NOAA SARSAT program, providing connectivity between satellite download facilities in Fairbanks, Wallops Island and Suitland.

NESDIS and SWPC DSCOVR Connections: With the recent and long-awaited launch of the nation's first operational climate satellite in deep space - the Deep Space Climate Observatory (DSCOVR) - N-Wave took part in providing network connectivity for the DSCOVR satellite ground system from the Fairbanks Control and Data Acquisition System (FCDAS) in Alaska. N-Wave helps deliver DSCOVR telemetry data to the main ingest point and user of this data - the Space Weather Prediction Center (SWPC). N-Wave delivers the DSCOVR data from FCDAS to SWPC at its primary location in Boulder, CO and at the Alternate Processing Site in College Park, MD.

Ball Aerospace to JPSS GRAVITE Connection: Just before the end of 2015, N-Wave provisioned a unique connection for the Joint Polar Satellite System (JPSS). It connects Ball Aerospace in Boulder, CO to the JPSS systems in Suitland, MD via the NOAA Denver TICAP. This connection is used to transfer JPSS satellite instrument characterization data for the JPSS-1 bus and instruments from the Ball Aerospace Boulder facility to the JPSS Government Resource for Algorithm Verification, Independent Test, and Evaluation (GRAVITE). This new and temporary N-Wave link carries science scenario data, satellite-related test data, and sensor characterization data during the JPSS spacecraft's "check-out" time frame - before its delivery to the launch site.

NESDIS / OSPO / LRIT Connections: As part of a joint effort, NESDIS has been working with N-Wave to consolidate expensive, low-bandwidth commercial point-to-point circuits out of Fairbanks and Wallops Command and Data Acquisition Systems over the N-Wave high-speed circuits. The NESDIS Low-Rate Information Transmission (LRIT) System was one of the first operational satellite systems to migrate and be fully tested for speed, stability and failover (N-Wave to the current operational T1s). N-Wave and the Office of Satellite and Product Operations (OSPO) network engineers worked hard on a solid implementation plan to not only encrypt the data passing over the N-Wave private VRF for LRIT, but also to make a seamless transition. This was achieved by first migrating the existing LRIT backup T1 routers to the N-Wave Ethernet links, and then by changing the routing metrics to make the LRIT backup routers run as primary over the N-Wave links while keeping the dedicated T1s for backup.

NESDIS Admin LAN: A private VRF is now operational, connecting the NESDIS National Centers for Environmental Information (NCEI) data centers in Stennis, Boulder, Asheville and Silver Spring.

N-Wave Service Reliability

Now entering its sixth year, N-Wave has delivered high-reliability service to its customers from inception. The N-Wave core network has consistently provided better than five 9s (+99.999%) of network service availability.

Designed with high availability in mind, the N-Wave core network consists of carrier-class Cisco and Juniper routers. Each core site in Atlanta, Denver, Chicago, Seattle and Washington, D.C. hosts two core node Juniper or Cisco routers. Internet2 provides the 10G optical wave services between these core nodes. Internet2 has contracted with Level 3 for optical fiber and for joint management of the optical network infrastructure.

For servicing customers at the edge of the network, and where high availability is necessary, N-Wave recommends and deploys two diverse paths, with a target of four 9s (99.99%) availability. For example, at the National Centers for Environmental Information (NCEI) in Asheville, NC, two fully-diverse 10G circuits connect the Asheville facility to both the N-Wave Atlanta core node (via a southern route) and the N-Wave McLean, VA core node (via a northern route). Over the past year, NCEI-Asheville has had a measured availability of 99.99943% utilizing this network topology.

Another high-availability site for N-Wave is the NOAA Environmental Security Computing Center (NESCC) in Fairmont, WV. Currently, N-Wave has five commercial 10G circuits going into NESCC over three diverse paths and landing at three different N-Wave core nodes - one in Chicago, another in Atlanta and the third in Washington, D.C. For additional reliability there are two N-Wave Juniper routers diversely terminating these services at Fairmont. Over the past year, NESCC has been available over 99.999% of the time.

The table below shows other selected edge sites and their layer 3 measured network availability. [2]

Silver Spring (SSMC-2)	99.999%
College Park (MAX)	99.994%
Suitland (NSOF)	100%
Wallops (WCDAS)	99.999%
Fairbanks (FCDAS)	99.999%
Boulder (DSRC)	99.999%

[1] See article on N-Wave Service Level Agreements
[2] More sites and our weekly report available at:
http://noc.nwave.noaa.gov/nwave/support/weekly-reports/n-wave-noc-weekly-reports.html

N-Wave Service Level Agreements

N-Wave can provide NOAA customers with a network Service Level Agreement (SLA) based on an aggregation of several technical approaches:

- Carrier and/or partner SLA: When developing a solution for customers, N-Wave can establish an SLA based on the SLAs available from the contracted underlying circuits.
- Historical measurements of service availability: The N-Wave NOC actively manages and monitors all N-Wave circuits, routers, switches and supporting hardware and can use historical performance to estimate service reliability when developing an SLA.
- Engineering: N-Wave can develop an SLA through established engineering expertise, industry and community best practices, and coordination with customers, carriers and network partners. For example, telecommunications equipment is often rated by third parties for mean-time between failures (MTBF) and service availability ratings.

N-Wave measures service availability based on uptime. This is expressed as the amount of time (usually a percentage) that a site was available, i.e. able to reliably pass traffic, over a given period of time.

Some sites, specifically dual-homed sites, will have a *measured availability* (as measured by the N-Wave NOC through common network monitoring practices) and a calculated *aggregate availability*, based on a calculation done with the historical measured availability of the two circuits providing connectivity. The calculated aggregate availability number can be used as a probability of future uptime availability. In practice, the actual measured availability of the site is often better than calculated, since rarely are both circuits to a site down at the same time.

X-Wave & TICAPs: Status Update

In order to meet the requirements of the DHS TICAP 2.0 standard, N-Wave is working with other members of the NOAA Office of the Chief Information Officer (OCIO) and NOAA Cyber Security Center (NCSC) to implement a brand-new network for NOAA, known as X-Wave. X-Wave is an externally-facing network that complements N-Wave, providing a resilient border network outside of the TICAP 2.0 stack, facing the Internet. X-Wave enables TICAP forwarding symmetry and TICAP failover. It consolidates NOAA's external peers into a single, high-availability security zone. Other agencies such as NASA and the Department of the Interior are implementing a similar strategy for supporting TICAP requirements.

X-Wave routers are deployed with the TICAP services at the following locations: McLean, VA (Level 3 Communications), University of Maryland, College Park (MAX Gigapop facilities), Denver, CO (Level 3 Communications), Seattle, WA (Westin colocation facilities) and Dallas, TX (Level 3 Communications). These locations provide rich connectivity to external partners such as Internet2, academic Research and Education (R&E) networks, other agency networks, international R&E networks, and large commercial Internet service providers. As designed, X-Wave fits well within the "DHS TICAP Reference Architecture".

The existing N-Wave high-performance network will continue to provide NOAA-to-NOAA transport, as well as providing transport from NOAA users to the Internet and external partners (via the TICAPs). Existing N-Wave contracts and relationships are being leveraged to implement X-Wave.

The mission of X-Wave is to provide redundancy for the TICAPs; a failure of one X-Wave site or TICAP site at any of the above locations will automatically failover to another pre-designated X-Wave/ TICAP site. This resiliency will provide NOAA with high-availability routing to external partners, in order to meet our mission. To ensure this, the N-Wave and X-Wave architectures have been thoroughly modeled in a Juniper virtual environment to fine-tune the IP routing policies.

As of March 2016, the X-Wave routers have been deployed at all of the above locations. The Dallas and Seattle sites are available and ready for customers, and Denver is scheduled to be available in Q4 FY16. The Honolulu TICAP site is also available to customers now, but due to its location it is not connected to X-Wave.

The target availability for the D.C.-area X-Wave and TICAP 2.0 stack is Q4 FY16. NOAA users in the D.C. area are currently supported by the Silver Spring NOC. The N-Wave team is working with the Silver Spring NOC to develop transition plans to migrate NOAA users from the existing TICAP 1.0 stack and legacy Internet connections to TICAP 2.0 and the new, upgraded Internet connections, which will include 100 Gbps connections to the MAX Gigapop and Internet2, as well as 10 Gbps to the commodity Internet services of Level 3 and MAX. The intent is to make the transition as smooth and seamless as possible. More information will be provided from the NOAA OCIO soon.

The following figure provides an overview of the regional N-Wave to X-Wave relationship through the TICAP infrastructure.

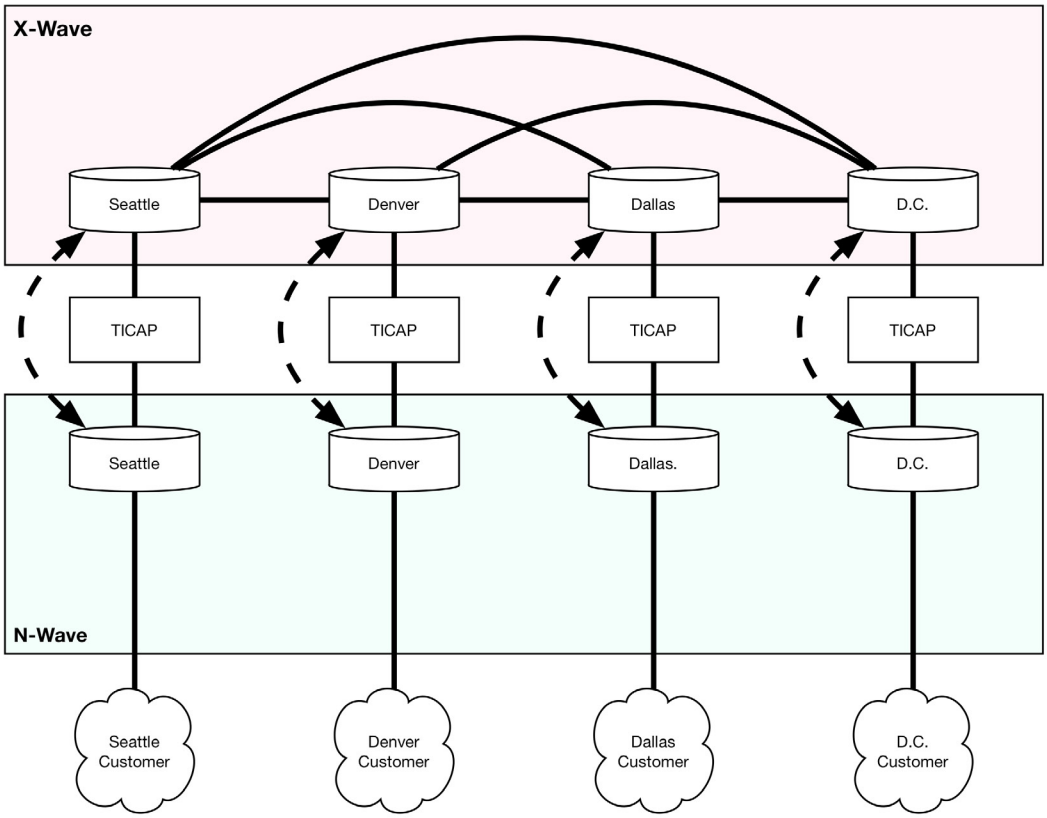

Each regional TICAP requires forwarding symmetry to provide proper security functionality. N-Wave's external network, X-Wave, uses asymmetric forwarding for sources from external peers. The X-Wave network ensures forwarding symmetry by assigning N-Wave participants to specific TICAP sites. The following figure illustrates X-Wave's role for ensuring forwarding symmetry.

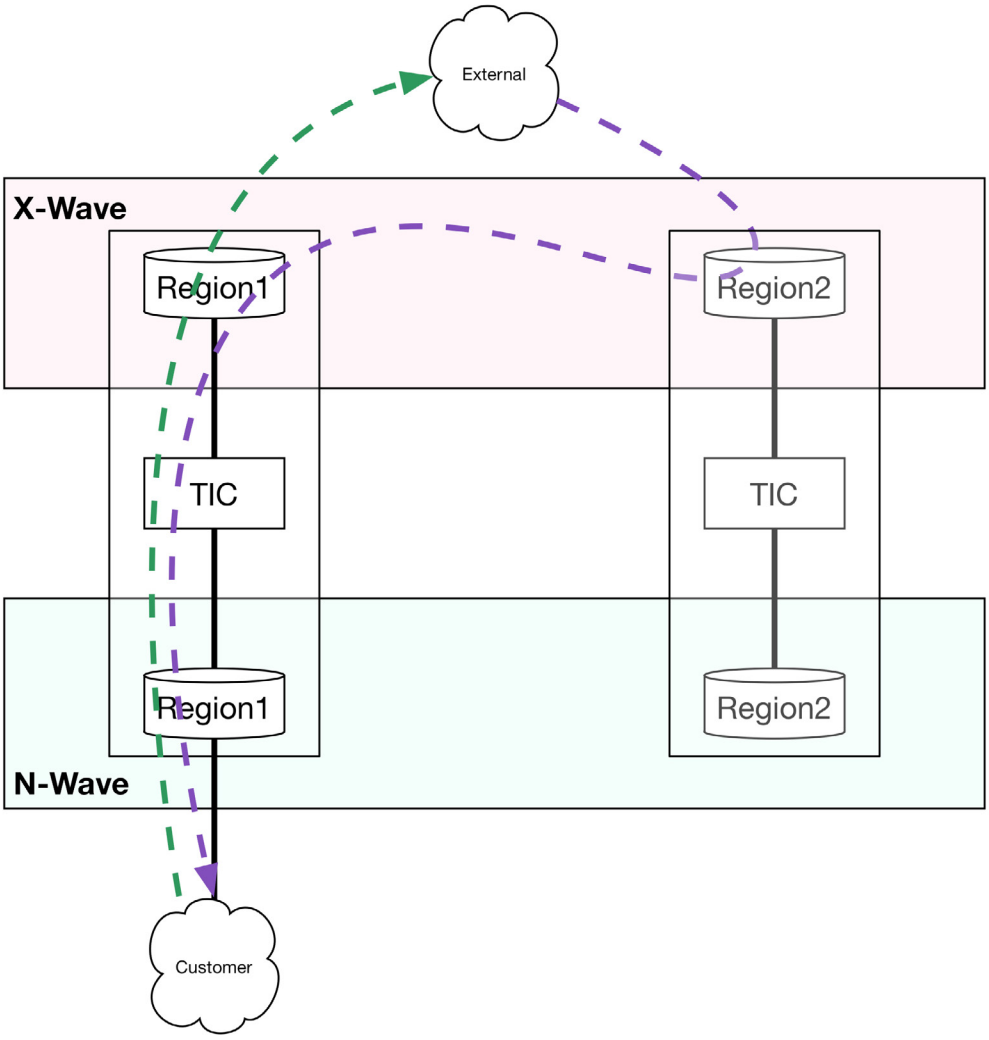

Additionally, X-Wave and N-Wave mitigate site failures by dynamically forwarding participant traffic toward designated backup TICAP sites. The figure below illustrates regional resiliency upon failure of a primary TICAP facility.

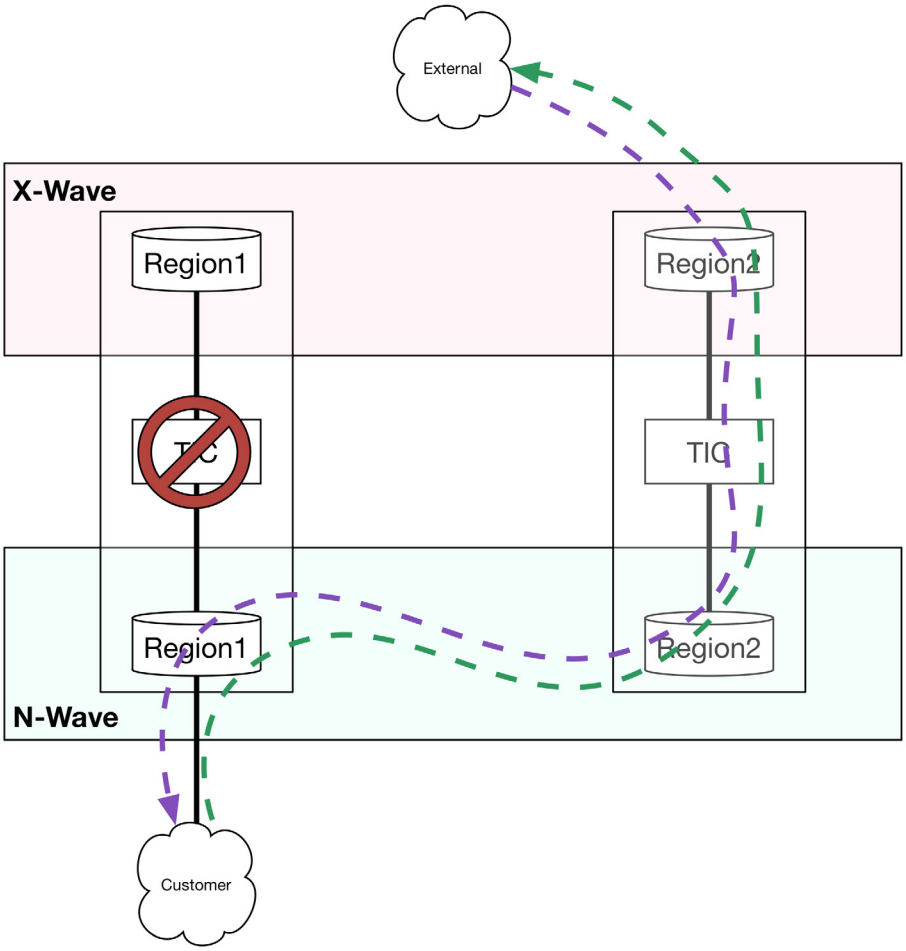
External
X-Wave
Region1
Region2
TIC
TIC
Region1
Region2
N-Wave
Customer

Internet2 is holding its annual meeting of members, the 2016 Global Summit, 15–18 May in Chicago, IL. 2016 marks the 20th anniversary of the founding of Internet2, and this year's Global Summit will showcase the community's rich history and bright future.

The Global Summit brings together the researchers, network engineers, presidents, and CIOs from Internet2 members, collaborating international Research and Education (R&E) networks, and federal agencies that use high-performance networks to support mission applications and disciplines. Members of Internet2 include academic, R&E network, affiliate and industry representatives. (NOAA became a member in 2000 in order to access Internet2's high performance network, then called "Abilene").

NOAA attendees Robert Sears, David Hartzell, Mark Mutz & Jerry Janssen will join various programs, sessions and side meetings including:
David Hartzell, Mark Mutz & Robert Sears: Joint Engineering Team side meeting
Robert Sears & Jerry Janssen: Executive Track

N-WaveNews Issue 07 May 2016

NOAA SCIENCE NETWORK

For more information contact:
NOAA N-Wave Science Network
http://noc.nwave.noaa.gov
Office of the chief Information Officer
http://cio.noaa.gov
Robert Sears, Network Manager
Holly Palm, Design and Layout
Tiffany Fourment, Editor

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