PACIFIC WAVE

PACIFIC WAVE IS A PROJECT OF: CENIC PACIFIC NORTHWEST GIGAPOP

Year 2 Annual Report (Sept 1, 2016 — August 31, 2017)

NSF Award ACI-1451050

IRNC: RXP – Pacific Wave Expansion Supporting SDX & Experimentation

Submitted: July 19, 2017

Louis Fox (lfox@cenic.org)

Year Two Goals & Accomplishments

Project goals

This award funds the development of a distributed multi-domain SDX (Software Defined eXchange) to support collaboration among international research and education networks experimenting with Software Defined Networks throughout the Pacific Rim region. The Pacific Wave exchange will support and enhance collaborative research into astronomy, digital media and cinematic arts, biology and medicine, computation sciences, earth sciences, grid computing, high-energy physics, and oceanography.

The overall objectives of the Pacific Wave Expansion Supporting SDX and Experimentation are to enhance, upgrade, and evolve the production Pacific Wave distributed open exchange fabric to support multiple 100 gigabit-per-second connections from countries through the Asia-Pacific area and to develop a set of parallel facilities supporting experimentation in interconnecting multi-domain Software Defined Networks through a distributed exchange. Pacific Wave will achieve these objectives through engagement with the 17 major international research and education networks currently connected to the Pacific Wave exchange and through information and technology exchanged with other IRNC awardees and the US R&E community.

Activities & Accomplishments:

Software Defined Networking/Software Defined Exchange

- Acquired and deployed an SDX node in Sunnyvale [Completed May 2017]
    - Deployed SDX switch at Sunnyvale
    - Connected Sunnyvale SDX to Los Angeles SDX and Sunnyvale PacWave at 100 Gbps
- Continue engagement with SDN applications of Pacific Wave participants
- Aim to engage and support testbeds, experimentation, and SDX control frameworks
- ONOS SDN software [Completed March 2017]
    - Installed ONOS software on Los Angeles Hypervisor server

    - Logically connected ONOS controller to both Sunnyvale and Los Angeles SDX switches
- Castor SDX software [Completed May 2017]
    - Collaborated with Castor developers from Australia and installed Castor software onto Los Angeles Hypervisor server
    - Integrated Castor into ONOS platform and performed initial data-plane testing using Los Angeles SDX and Los Angeles PacWave switches
    - Discussed collaboration with Starlight and AARNet
    - Developed integration plans for Castor connectivity to other entities
    - Made plans to identify and explain integration plans in Year3 Goals

Free-Range Routing (FRR) software [Completed June 2017]

- Installed FRRouting software on Los Angeles Hypervisor server that will function as the BGP route-server connected to the Castor software and also peer with other IRNC entities

Specific Objectives:

100 Gbps connected perfSONAR nodes

One goal of Pacific Wave's perfSONAR deployment has been to develop perfSONAR nodes capable of speeds approaching 100 Gbps. Pacific Wave has completed the deployment of two 100 Gbps connected perfSONAR nodes in Seattle and Los Angeles. As part of upgrading the perfSONAR toolkit from version 3.5 to 4.0, we re-imaged the 100 Gbps connected nodes with CentOS7 to take advantage of new features, TCP enhancements, and tuning capabilities present in the 3.x Linux kernel.

In *ad hoc* testing between the two 100 Gbps connected nodes we obtained a peak throughput of ~45 Gbps; in mesh-scheduled testing, we observed an average throughput of ~40 Gbps. In both scenarios, we employed moderate system and NIC tuning and used iperf3, TCP, single-stream over the ~25 ms latency between the two endpoints.

In collaboration with AARNet (Tim Rayner, Warrick Mitchell, et al.) we set up an environment to enable more aggressive, multi-stream testing between the Pacific Wave 100 Gbps connected nodes in Seattle and Los Angeles and an AARNet 100 Gbps connected node in Canberra, Australia. In this scenario, each node had three separate instances of iperf3 running in server mode on distinct TCP ports, with each iperf3 process pinned to a separate CPU core. The theory was that multiple iperf3 instances should be able to achieve a higher aggregate throughput than a single iperf3 process — particularly on 40 Gbps or 100 Gbps connected devices. Resultsincluded a peak aggregate throughput of 53 Gbps from Canberra to Seattle and of 49 Gbps from Seattle to Canberra.

We are in the process of acquiring a 100 Gbps connected perfSONAR node for Sunnyvale as part of our Year 2 deliverables, and expect to deploy it in the summer of 2017, although it is possible the work will extend into Year 3.

PACIFIC WAVE

Increased instrumentation of Pacific Wave infrastructure

At the invitation of ESnet and Internet2, one of Pacific Wave's 10 Gbps connected perfSONAR nodes in Los Angeles participated in a beta program of pScheduler — the replacement scheduler under perfSONAR 4.0 for the separate bwctld and owampd scheduler services. This node has since been upgraded to the perfSONAR 4.0 test point bundle running on CentOS7.

Pacific Wave also provisioned a 10 Gbps connected perfSONAR node for Pacific Wave Sunnyvale to replace a server borrowed from CENIC's CalREN network. The motivation for the equipment loan was to bridge a gap in visibility and support during SC16 activities. The IRNC-funded Sunnyvale 10 Gbps connected perfSONAR node has since been integrated into the Pacific Wave perfSONAR MaDDAsh.

Staff also worked with Internet2 to set up regular testing between the Pacific Wave 10 Gbps connected perfSONAR nodes in Seattle and Los Angeles and a selection of four Internet2 Performance Assurance Service (PAS) perfSONAR nodes across the US topology.

We continued to explore potential frameworks by which to enable the dynamic scheduling/provisioning of 10 Gbps connected perfSONAR nodes for a user-specified circuit (VLAN). As part of our participation in the GLIF AutoGOLE activities (see below), 10 Gbps connected perfSONAR nodes in Seattle, Sunnyvale, and Los Angeles have been made available via NSI-based orchestration/reservation.

perfSONAR MaDDash

During September, 2016, we migrated the Pacific Wave perfSONAR MaDDash from v1.3 to a beta of v2.0. We worked with the developers to identify and resolve apparent interoperability and configuration challenges in a subset of MaDDash and perfSONAR toolkit interactions, including the time-series graphing functionality.

Regular testing among Pacific Wave perfSONAR nodes across public, inter-exchange VLANs was established to provide additional visibility of performance. Pacific Wave has also continued to explore the potential for establishing regular testing among Pacific Wave IRNC-funded perfSONAR nodes and those of interested Pacific Wave participants and partner networks.

Virtual Customer Equipment (VCE) Software

Begin development of Virtual Customer Equipment (VCE) with GlobalNOC [In progress]

- VCE software connected to switch in CENIC HQ Lab for initial testing
- Basic functionality (show commands) confirmed working
- More advanced functionality (interface/service configuration) is being tested and debugged; development is ongoing

Currently working on moving VCE software from IU-controlled VM's and onto CENIC-controlled VM's

Significant Results:

AutoGOLE / NSI

We continued collaboration with the GLIF - AutoGOLE working group, exploring service and delivery models for providing dynamic, interdomain circuit services based on NSI-signaling. As part of a long-term trial of providing dynamic circuit services on Pacific Wave, we established separate NSI domains for the Pacific Wave Seattle, Sunnyvale, and Los Angeles exchanges, with these elements:

- NORDUnet-developed OpenNSA, providing the abstraction layer between the Pacific Wave traditional Layer2 switches and NSI orchestration elements.
- Control-plane peering with the NSI-aggregators managed by ESnet, Netherlight, and StarLight.
- Data-plane peering with ESnet, StarLight, and SINET, with interest in developing further peering with other Pacific Wave participants.

- A 10 Gbps connected perfSONAR node available as a NSI-assignable resource in each domain:
    - urn:ogf:network:sttlwa.pacificwave.net:2016:topology:irnc-10g02.sttlwa
    - urn:ogf:network:snvaca.pacificwave.net:2016:topology:irnc-10g01.snvaca
    - urn:ogf:network:lsanca.pacificwave.net:2016:topology:irnc-10g02.lsanca
- perfSONAR nodes, participating in the AutoGOLE 'demo' range of VLANs IDs, 1779-1799.

- Participation in a pilot leading toward broader adoption of the RNP-developed MEICAN webUI for dynamic circuit provisioning:
    - https://wiki.rnp.br/display/secipo/AutoGOLE+MEICAN+Pilot

- Expanded staff involvement in AutoGOLE activities, with contributors from the engineering groups of PNWGP and CENIC as well as representation from the CENIC NOC, and with a goal of gaining familiarity with MEICAN webUI and the underlying elements.

Key Outcomes & Accomplishments:

Data Transfer Nodes (DTN)

As part of the collaboration with the PRP (Pacific Research Platform), we reprovisioned a contributed FIONA (Flash I/O Network Appliance) with 96TB (raw) storage for deployment as a DTN on Pacific Wave Los Angeles. The storage subsystem is based on 16x8GB SAS2 spindles managed as a ZFS pool. The DTN was initially connected at 20 Gbps (organized as a 2x10 Gbps LACP-bundle), and we have since upgraded the connectivity to 1x10 Gbps + 1x100 Gbps. The DTN is running CentOS7, under which we can take advantage of TCP pacing of multiple streams to reach higher effective (aggregate) throughput.

We are in the process of acquiring 100 Gbps connected DTNs for Pacific Wave Seattle, Sunnyvale, and Los Angeles. Goals include taking advantage of recent advances in flash-based storage. Our intent is to provide flexible access to a modest amount of very fast NVMe-based storage and a larger amount of storage based on SAS3 spindles.

As of the end of June, 2017, we are in the final stage of vetting the hardware specifications for the new DTNs with equipment providers. Due to the challenges of optimizing the component selection, the vetting process has been longer than initially expected. Full deployment and integration of the DTNs may slip into Year 3 activities.

We continue to explore data-transfer toolsets and supporting technologies in the interest of providing researchers a range of options from which to choose the data-movement / data-placement methodology which best fits their science workflow. In addition to GridFTP and as a Globus endpoint (w/Globus Provider - 'network' subscription), we have explored making the PRP-contributed PWave DTN in Los Angeles available using Apache - wget, Squid proxy with OAuth 2.0, and ssh (scp, ssh-ftp/sftp):

- Apache - wget & Squid proxy - OAuth 2.0: Working with Andrew Howard in support of researchers moving genomic data to Singapore's NSCC, we have set up access to the DTN via Apache/wget, using local accounts. On our roadmap, we intend to stand up an authenticated Squid proxy/cache, with OAuth 2.0 authentication.

- ssh (scp, sftp): Working with with NCMIR - UCSD and the Kisailus lab at UCR, as part of exploring movement of imaging data between the two US facilities and across the 'Microscopy Wormhole' to researchers in Singapore.

The DTN participates in the PRP GridFTP MaDDash, which represents disk-to-disk performance data as throughput using the same tools as for visualizing iperf / memory-to-memory performance. Under the current test regime, a 10 GB reference file is transferred among participating DTNs. Depending upon the other endpoint in the pair-wise mesh testing, we have observed the DTN to have an average disk-to-disk performance of ~14 Gbps.

Working with Globus, we established the DTN as a (public) Globus Managed Endpoint, part of a trial of Globus' Provider Subscription Service. This is part of our exploration of making the DTN available as a resource for 'live' research data movement. In this context, the write-access to the DTN is obtained through using Globus' AWS-based orchestration for federated authentication. The DTN is also available as a public anonymous test-point for read-only transfers of a subset of the reference data files provided by ESnet.

The DTN also participates in our AutoGOLE / NSI activities as a resource in the Pacific Wave Los Angeles NSI domain:

urn:ogf:network:lsanca.pacificwave.net:2016:topology:dtn0.lsanca

The DTN participates in the AutoGOLE 'demo' range of VLANs IDs, 1779-1799. In this context, the DTN is primarily useful as a test-point for *ad hoc* disk-to-disk performance testing.

PACIFIC WAVE

In collaboration with Andrew Howard and in support of a SingAREN-initiated request for assistance the Los Angeles DTN was used for transferring genomic data-sets between the US and researchers at Singapore's NSCC. This configuration also was useful in comparing throughput to Singapore of data-sets on the Pacific Wave DTN over R&E paths versus moving the same data-sets from Amazon's AWS.

Exploration of an Openflow 1.0 environment

While Openflow 1.0 was initially considered a testbed for development and experimentation, after further examination and discussion with other participants, it was decided that utilizing Openflow wasn't a path that most entities wanted to pursue. As such, we decided against exploring this testbed environment any further.

Training and professional development

Two members of CENIC's core engineering team, John Hess and Will Black, have rotated into the project for an opportunity to learn more about perfSONAR, Data Transfer Nodes, SDN/SDX, AutoGOLE, and NSI. Both have been contributing their knowledge to the project while learning about the advanced technologies used by Pacific Wave and the participating National Research and Education Networks (NRENs). John is also active in the Pacific Research Platform (PRP) project, and provides a strong technical bridge between the two projects.

Result dissemination to communities of interest

Pacific Wave staff regularly attend meetings related to R&E Exchange Points where they participate in discussions regarding the technologies deployed and under consideration within Pacific Wave. Presentations include updates about the projects and, as appropriate, typically contain a "lessons learned" section.

Pacific Wave also hosts the bi-weekly PIREN calls with participants from CENIC, PNWGP, UH, AARNet, REANNZ, and Guam. These calls are used to discuss items of mutual interest, share project plans, and review outages that may have impacted participants.

Year 3 Goals:

Software Defined Networking/Software Defined Exchange (SDN/SDX)

For the second year of this project, we spent most of our time getting the needed pieces in-place for testing. For the third year, we will be in a position to start utilizing those pieces quite rapidly. Below is a summary of the integration plans for this third year.

Initially, we will configure an eBGP session between the FRR route server and peers at AARNet and Starlight using general purpose routers at both entities. Any Layer2 connectivity that may be needed between the FRR route server and these entities will utilize either the AutoGOLE/GLIF NSI mechanism via the MEICAN tool or, if the NSI mechanism is unavailable, traditional Layer2 VLAN's. Once these BGP sessions are established, the routes for the 100 Gbps connected DTN hosts will be advertised via these BGP peerings. Assuming there are no other issues, once these routes are advertised, the various DTN hosts will be able to communicate with either over the SDX switches using the software-defined data-plane forwarding path that will be programmed by the ONOS SDN software. Using this model, some amount of the forwarding-plane will be SDN-controlled, and some will utilize traditional networking methods. [Will be completed by September 2017]

After we have confirmed this basic functionality is working, we will move to setup BGP sessions between the Castor route-servers at the three different entities. This modified configuration should expand the size and scope of the SDX networks to include more software-controlled pieces. [Will be completed by November 2017]

Following all of these tests, we will be in a position to start peering with other general Pacific Wave participants and leverage the SDX network even for entities that aren't using software-controlled networks. [Will be completed by February 2018]

In additional to expanding the size of the SDX network, we will continue improving the reliability of the SDX network by also installing the FRR software onto the Seattle hypervisor server, along with the SDN controllers already present in Seattle. [Will be completed by November 2017]

Virtual Customer Equipment (VCE) Software

First and foremost, our primary goal is to get the VCE software working well on the Brocade in the CENIC HQ lab. Following this, in late August, 2017, we plan to start a production trial using the Los Angeles PacWave switch. This trial will be aimed at ensuring the VCE functions properly in a production environment and doesn't cause any unintended issues, while also ensuring that customer configurations are kept logically separate from one another.

After these initial production trials are completed, we will move forward in providing access to PacWave participants, initially on a limited basis but eventually with full deployment. [Initial limited access completed by September 2017, fully deployed by December 2017]

Add additional capacity between Sunnyvale and Los Angeles

Year 3 of the project calls for adding a second 100G circuit between Sunnyvale and Los Angeles to the Pacific Wave (and WRN) backbone.

Pacific Wave traffic peaked at around 85 Gbps during SuperComputing 2016. It is anticipated that traffic peak for this SuperComputing will increase further in 2017, with a requirement for an additional 100 Gbps of capacity to be added before SC 2018.

This upgrade calls for CENIC to acquire another pair of 100 Gbps transponders for the DWDM system on this route. CENIC will procure the Ciena transponders for this route and manage the deployment to be completed and operational by the end of December 2017.

PACIFIC WAVE

Work Breakdown Schedule for new 100G circuit, Los Angeles to Sunnyvale

9/5/17	Work with Ciena and Internet2 on optical system design
9/6/17	Confirm 100 GE port availability on Brocades in Sunnyvale and Los Angeles
9/18/17	Request quotes for Ciena transponders from at least 3 resellers
10/2/17	Issue purchase order for transponders
11/20/17	Coordinate with Level3/Internet2 for installation of transponders
12/11/17	Complete circuit testing and place circuit into production

Data Transfer Nodes (DTN)
As part of increasing our understanding of data-movement workflows and the performance of various toolsets, we intend to instrument the Pacific Wave DTNs to measure and characterize traffic patterns. Tstat (TCP STatistic and Analysis Tool):

http://tstat.polito.it/

Tstat is a passive 'sniffer' which can gather statistical data of TCP (and UDP) traffic using standard libraries on commodity server hardware. As we expect to able to run Tstat directly on the Pacific Wave DTNs, the utility should provide valuable insight into the bi-directional traffic flows associated with data-movement activities. To address privacy concerns, Tstat provides a utility to anonymize the host portion of the IP address (IPv4 and IPv6) in the collected flow data.

Tstat-collected flow data is available in several of output formats, including an integration for RRDtool, allowing for visualization of performance indexes over time while the measurement archive consumes a (configurable) fixed amount of storage. We intend to make summary traffic statistics available through a web interface and to explore other reporting and visualization mechanisms.

Tstat has been used successfully by ESnet to characterize traffic patterns associated with their four public anonymous DTNs, as well as by the NetSage project:
http://internationalnetworks.iu.edu/projects/netsage/index.html

PRODUCTS

Presentations produced on the project

Sep 22, 2016 | NorduNet | David Reese | "Pacific Wave and PRP; Big News for Big Data"
Sep 30, 2016 | TechExchange/GLIF | John Hess | *"Pacific Wave SDN/SDX Update"*
Oct 13, 2016 | e-Research Australasia | Celeste Anderson | *"Overcoming the Tyranny of Distance in 21st Century Research"*
Oct 18, 2016 | CANS | John Hess | *"PRP measurement, operations, engagement"*
Feb 21, 2017 | PRP| John Hess | *"Pacific Wave SDN/SDX Update"*
Apr 26, 2017 | Internet2 | David Reese | *"Western Region Network"*
Jun 01, 2017 | CUDI2017 | John Hess | *"Pacific Research Platform"*

PARTICIPATING ORGANIZATIONS

Organizations involved as partners or collaborators.

Asia Pacific Advanced Network (APAN)

At the recent APAN meeting in New Delhi, India, Pacific and APAN executed an MOU, renewing and extending their long-standing relationship. The MOU calls for the two organizations to work together in promoting their respective and mutual objectives through collaborations between customers/members. It also calls for joint work reporting on the R&E use of interconnections between networks and encouraging use of their networks in supporting the development of next-generation infrastructure and applications.

Pacific Islands Research and Education Network (PIREN)

Pacific Wave continues to host be-weekly conference calls with participants from PIREN (University of Hawaii), AARNet, REANNZ, and, new this year as a regular on the calls, Rommel Hidalgo, representing the University of Guam. The calls are used to discuss joint activities of interest (such as NetSage and perfSONAR), activities related to the AARNet links, and upcoming events. The calls are lead by Celeste Anderson, and notes summarizing the calls are distributed to the PIREN mailing list.

Pacific Research Platform (PRP)

Pacific Wave continues to provide the PRP with connectivity to regional, national, and international participants. Pacific Wave staff continue to participate in weekly PRP engineering calls as well as bi-weekly PRP science engagement calls. Pacific Wave is supporting PRP efforts to explore traffic engineering and data-movement technologies by providing participants with more options and tuning knobs for affecting the announcement profile of the routes for their Science DMZ resources as well as to develop the potential to steer related traffic across specific high-speed, inter-domain paths (aka SuperChannels).

Pacific Wave successfully facilitated a request through ARIN for the assignment of an Autonomous System Number (ASN) specifically for the use of the PRP — Pacific Research Platform / AS395889. Pacific Wave is providing hypervisor VM resources to support the PRPv2 BGP pilot, which is examining the potential of route-servers operating a BGP ‘view’ of PRP/AS395889 to provide a BGP-signaled control-plane among participants’ networks. Pacific Wave Layer2 infrastructure providesthe separate, high-speed data-plane for data-movement.

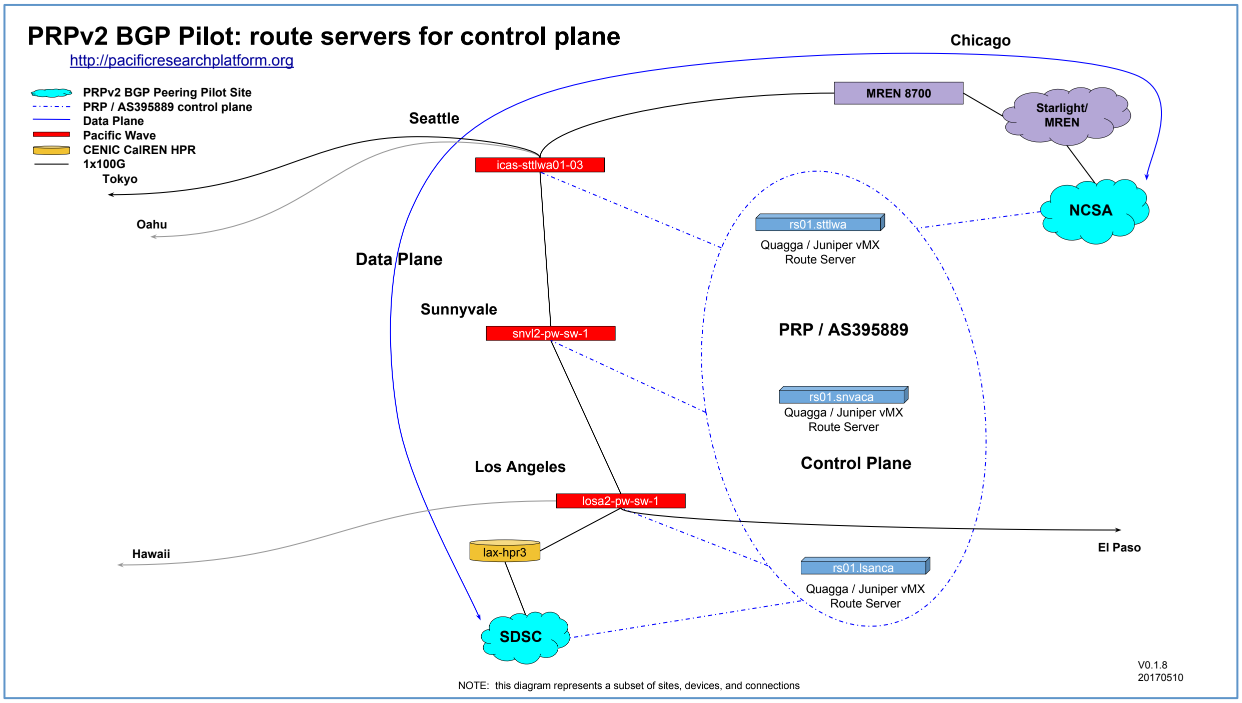

Among the goals of the PRPv2 BGP pilot is to inform the process of assessing various models for scaling from a regional a national and/or global model, from a PRP to an NRP, or GRP:

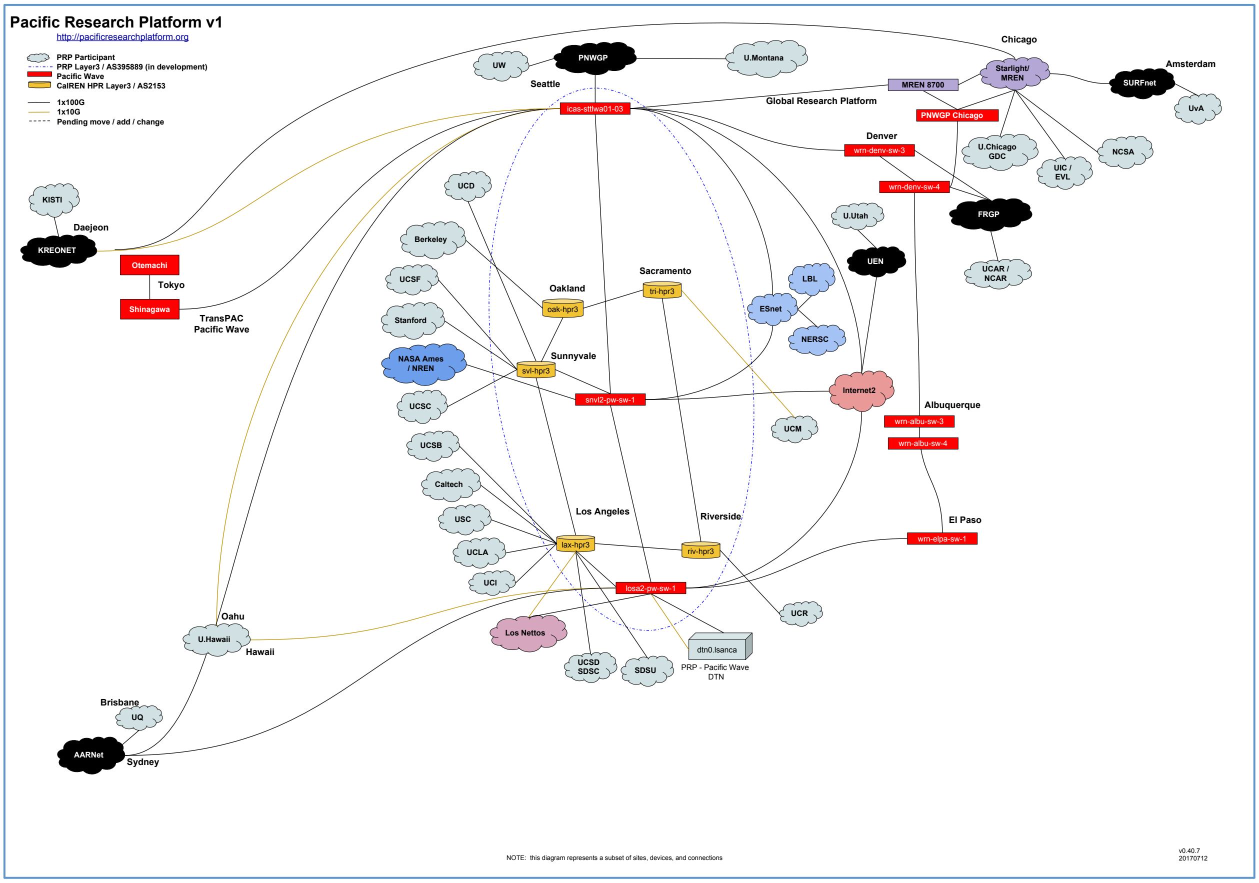

Network Startup Resource Center (IRNC:NSRC)

In May, Pacific Wave staff participated in a 5-day NSRC Direct Engineering Assistance (DEA) engagement with the University of Guam (UoG). The goals of the DEA include providing on-site training and engineering toward developing local expertise over time to not only benefit the specific institution but provide a resource for the surrounding region as well. The DEA team consisted of network engineers and system administrators from NSRC, University of Oregon, Pacific Wave (CENIC), UoG, and Northern Marianas College (Saipan). The team collaborated on several activities, including:

- Enhancing the UoG campus core and border networks in preparation for high-speed R&E transpacific connectivity through PIREN-coordinated activities;
- Reprovisioning connectivity of the UoG perfSONAR node from 100 Mbps to 1 Gbps. The perfSONAR had been donated by the NSRC as part of the initial UoG DEA engagement in July, 2016. The perfSONAR is currently participating in regular testing with the University of Hawaii, and is intended to be included in a perfSONAR MaDDash hosted by CENIC.
- Installation of a PRP-donated DTN. The DTN is a 'FIONette' (a low-cost version of the Flash I/O Network Appliances developed for the PRP), and is configured with 1 Gbps connectivity and 2 TB

of storage. The DTN is participating as a node in the PRP GridFTP disk-to-disk throughput testing regime, and is intended to inform the data-movement requirements for supporting science workflows of UoG-based researchers, including:
  - Marine Laboratory (Guam EPSCoR - Ecosystems Collaboratorium database for coral specimens and oceanographic data)
  - WERI (Water and Environmental Research Institute) - Professor Yuming Wen is involved, and we are moving their GIS lab elsewhere on campus, where other GIS researchers can share the equipment.
  - CLASS (College of Liberal Arts and Social Sciences) Geography department GIS Lab, run by Dr. Romina King.
- Instantiation of core IT infrastructure elements such as an NTP server, a jump/bastion host, as well as reorganizing the servers and systems supporting the campus DNS.

IRNC funded Network Operations Center (IRNC:NOC)

Pacific Wave staff are partnering with the GlobalNOC to expand the mapping capability for Exchange Points, circuits, and peers beyond that of the NSF-funded IRNC activities. The work will allow Exchange Points to export measurement data to the NOC that can be specifically classified based upon the policies of the Exchange Point and its Participant (see IRNC:AMI for more information on export tags).

The mapping capabilities were also proposed to the Global Network Architecture (GNA) Technical Team, and interest was expressed by the GNA participating exchanges to join the mapping effort to create a global map of R&E exchange points and related inter-connections.

IRNC funded Advanced Measurement Infrastructure (IRNC:AMI)

To increase the visibility and monitoring of Pacific Wave's resources, we developed an interface tagging scheme that, combined with SNMP polling, will allow us to export interface data to the IRNC:AMI collector. It will also allow us granular control over which interfaces will export utilization and link-state data to public-facing websites and which still stay completely private. Work began in March, 2016, and should be completed by September, 2017.

The tag definitions are:

- <no tag> The interface will be ignored by NOC and AMI as if it did not exist.
- [ns] The data will be collected and held private, and the link will not show on maps.
- [ns-stat] The data will be collected and held private, and the link will show on maps with an "up-down" status indicator.
- [ns-exp] The data will be collected and be visible publicly, and the link will show on maps with traffic data.

Global Lambda Integrated Facility (GLIF)

The Global Lambda Integrated Facility (GLIF) has been working to identify its future in the Global/International Exchange landscape given the recent introduction of the Global Network Architecture (GNA). Pacific Wave continues to participate and to support the GLIF, as we believe this is where experimentation on and future development of exchange point technologies can be done on a global scale. Activities such as Pacific Wave's own SDX experimentation require multiple participants to explore and test technologies that can later be deployed in a production environment. Pacific Wave will be partnering with StarLight and AARNet on some demonstrations of SDX capabilities at the upcoming GLIF meeting in Sydney, Australia, to be held September 25-27, 2017.

Global Network Architecture (GNA)

Pacific Wave participation in GNA continues through attendance at all GNA technical meetings. Pacific Wave is currently working on the final edits of a privacy policy that will be published on the Pacific Wave website per GNA recommendations. As mentioned under IRNC:NOC, Pacific Wave is working with the GlobalNOC on a mapping tool that will display GNA exchange points, circuits, and peers.

IMPACTS

Impacts on the development of the principal division(s) of the project

Findings, results, techniques or other products that were developed or extended from the project, and made an impact, or are likely to make an impact on the base of knowledge, theory, and research and/or pedagogical methods in the principal disciplinary field(s) of the project.

Pacific Wave is leading the way in the deployment of Science DMZ resources within an Exchange Point. During this project year, Pacific Wave will add Data Transfer Nodes (DTNs) in Seattle, San Jose. and Los Angeles. The outcome of this experiment could significantly alter how Exchange Points are built and operated in the future. In addition, the Pacific Research Platform (PRP) was initially built on the Pacific Wave and Western Region Network (WRN) infrastructure. This has already resulted in several interested organizations reaching out to Pacific Wave and WRN with an expressed interest in expanding the infrastructure and in a "National Research Platform" meeting scheduled in Bozeman, MT in August 2017.

The real impact is creating and sustaining the infrastructure to support the work carried out by others. In some ways, predicting when and if high capacity networks are needed is difficult, but when they are needed, capacity needs to be available. We are poised to support multiple large-scale projects with very large data requirements, such as the Large Synoptic Sky Telescope (LSST) and the Square Kilometer Array (SKA). There are increasing needs to support large data transfers for collaborations in genomics, physics, astronomy, and earth sciences, to name a few.

What is the impact on other disciplines?

Findings, results, or techniques that were developed or improved, or other products from the project that made an impact or are likely to make an impact on other disciplines.

Pacific Wave is working with other regional and national research and education networks to deploy components of the network. Examples of this are the various software components in the SDN/SDX deployment (ONOS, etc), and the MEICAN provisioning tool developed by RNP and used by Pacific Wave.

Through experimental deployments of these software stacks and/or tools, the products themselves are improved through feedback regarding problems reported and/or features requested. This process results in more stable solutions for future production-quality demands and wider deployment.

How the project made an impact or is likely to make an impact on human resource development in science, engineering, and technology.

The deployment of Software Defined Networks/Software Defined Exchanges (SDN/SDX) is likely to have an impact on the training and skills required of future network engineers. Engineers will no longer just be running experiments on a particular product, but will also require significant software development skills in order to design, build, operate, and debug future advanced networks.

What is the impact on physical resources that form infrastructure?

Ways the project made an impact, or is likely to make an impact, on physical resources that form infrastructure, including facilities, laboratories, or instruments.

Data-intensive science applications have always been major drivers of technology because they encounter technical barriers long before these limitations are experienced within other application domains. Multiple science domains are motivating the creation of new advanced communication services to meet the requirements of the most data-intensive applications in the world, including those in high-energy physics, astronomy, computational astrophysics, genomics, and computational chemistry.

New services are required to support large-volume, high-quality digital media, 3D digital media, and ultra-high-definition digital media. Digital media must be supported as a common, ubiquitous service, not as a highly specialized capability.

The Pacific Wave infrastructure is continuously evolving to meeting these growing demands. The addition of AutoGOLE/NSI and SDN/SDX capabilities were specifically selected to support the demands of research laboratories and distributed network-connected instruments.

What is the impact on institutional resources that form infrastructure?

Ways, in which the project made an impact, or is likely to make an impact, on institutional resources that form infrastructure.

Pacific Wave, working with the Pacific Research Platform (PRP), is refining the definition of the institutional Science DMZ and how it interconnects with, and integrates into, the regional, national and international infrastructure. Bridging the gap between the scientific resources within the institutional Science DMZ and those of their collaborators has led to the concept of 'the network as an instrument,' first proposed by Greg Bell of ESNet.

Creating an easily accessible and configurable method for constructing point-to-point and multi-point virtual networks across many boundaries will greatly enhance the ability of education and research institutions to serve the needs of their researchers and scientists.

Ways that the project made an impact, or is likely to make an impact, on information resources that form infrastructure.

Pacific Wave is engaged with the other IRNC funded projects (PIREN, AmLight, Starlight, NOC, and AMI) as part of both the operational and experimental aspects of operating and maintaining an exchange point. Pacific Wave is also engaged on an international level with organizations like the Asia Pacific Advanced Network (APAN), and is now a regular participant and contributor to the Global Network Infrastructure (GNA).

Through these efforts, Pacific Wave shares our experiences and uses that knowledge to help grow and expand the research and education infrastructure available to the research and education community.

The impact on technology transfer

Pacific Wave, as an exchange point, is essentially a modern-day technical "meeting place" of networks from around the world. Networks connect to exchange points for the explicit purpose of connecting to other networks connected to the exchange point for the interchange of research and education traffic. Because technology development moves at such a rapid pace, the networks connected at an exchange point are often at different stages of experimentation and development. This leads to a common "lifting" of networks to the more advanced uses of technology while providing a proving ground for the leading-edge ideas of network researchers.

The impact on society beyond science and technology

The ability of researchers and scientists to flexibly fashion international collaborations creates new avenues to further research in many areas. As the architecture of the exchange evolves to provide easily constructed pathways through multiple networks and organizations between participants around the globe, the speed with which *ad hoc* internet collaborations can be formed to solve real-time problems should greatly shorten the timeline of new discoveries.

CHANGES/PROBLEMS

Changes in approach and reasons for change

Actual or Anticipated problems or delays and actions or plans to resolve them

CENIC and Pacific Wave are continuing to monitor the recent acquisition of Brocade by Extreme Networks. As the acquisition is not set to close until the end of August, 2017, there have been no announcements regarding the future of the Brocade MLX-E product line deployed throughout Pacific Wave. Until these plans are announced, we will be closely watching our available pool of ports on the Pacific Wave switches and will only purchase additional capacity (hardware) on an as-needed basis.