PACIFIC WAVE

IRNC PIREN and Pacific Wave Updates

Americas' Research Platform WG
Calit2 UC San Diego
September 17, 2019

John Hess
Network Engineer - Research Engagement
CENIC - Pacific Wave

CENIC

PACIFIC NORTHWEST GIGAPOP

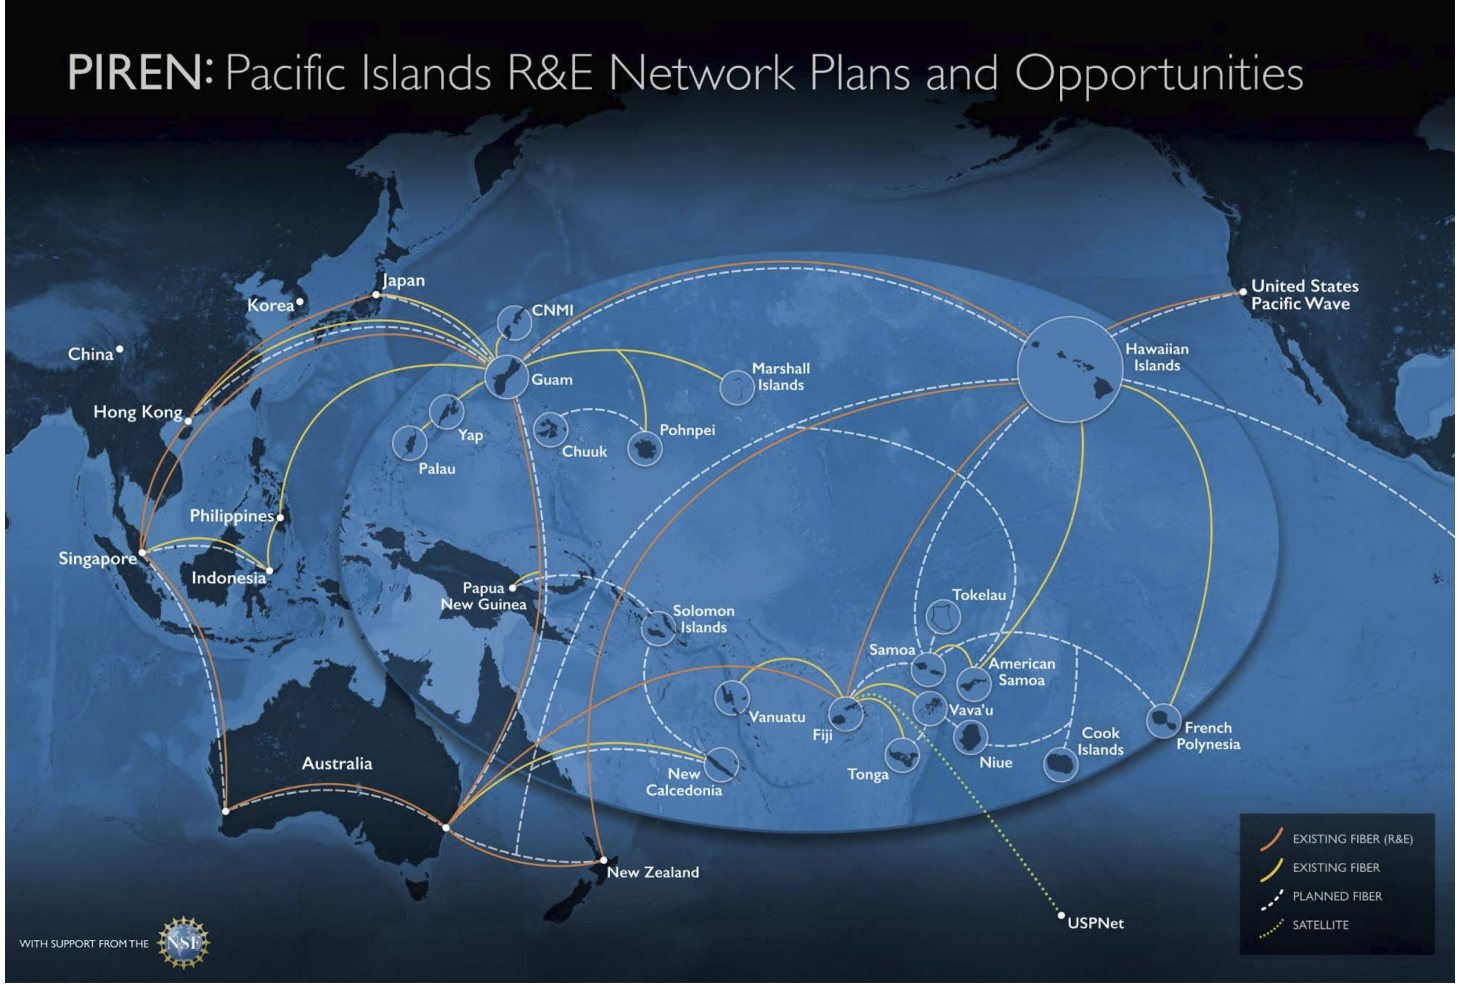
PIREN: Pacific Islands R&E Network Plans and Opportunities
Japan
Korea
China
CNMI
Guam
Marshall Islands
United States Pacific Wave
Hawaiian Islands
Hong Kong
Yap
Palau
Chuuk
Pohnpei
Philippines
Singapore
Indonesia
Papua New Guinea
Solomon Islands
Tokelau
Samoa
American Samoa
Vanuatu
Fiji
Vava'u
Cook Islands
French Polynesia
Niue
Tonga
New Calcedonia
Australia
New Zealand
USPNet
EXISTING FIBER (R&E)
EXISTING FIBER
PLANNED FIBER
SATELLITE
WITH SUPPORT FROM THE NSF

PIREN: Pacific Islands Research and Education Network

- Provide full domestic support for AARNET's current 2x100Gbps R&E circuits from Australia and New Zealand to the U.S., via Hawaii (including Mauna Kea)
- Continue to foster research and education (R&E) network capacity to interconnect Pacific Islands with each other and the global R&E network fabric by building on previous projects and relationships
- Opportunistically connect Mauna Kea and Haleakala, sites of major international astronomy observatories
- Collaborate and cooperate with IRNC measurement, NOC, Engagement, and Open Exchange awardeesCollaboration with IRNC awardees on SDX development

PIREN: Pacific Islands Research and Education Network

- AARNet: https://www.aarnet.edu.au/
- CENIC: http://cenic.org
- Pacific Wave: http://pacificwave.net
- REANNZ: https://reannz.co.nz/
- Pacific Northwest Gigapop – Seattle: http://pnwgp.net
- Global NOC – Indiana University: https://globalnoc.iu.edu/
- Network Startup Research Center: http://nrsc.org

PIREN: Pacific Islands Research and Education Network

- ## American Samoa launch HoloCampus 3D learning platform

  (August 22, 2019) American Samoa Community College (ASCC) has switched on a ground-breaking digital platform that will stream 3D holograms of University of Hawai'i faculty members to deliver classes and engage with ASCCstudents in real-time. Running on Hawaiki transpacific submarine cable, which directly connects American Samoa and Hawai'i. https://youtu.be/7_F3oqA3Ev8

- ## AstroFlows - U.Hawai'i Institute for Astronomy (IfA)

  Supporting high-speed data movement, traffic measurement, and flow characterization from observatories on Maunakea, Haleakalā, and Mauna Loa to researchers on connecting through PIREN to US mainland, Asia, and beyond. https://github.com/edubergeek/astroflows

- ## Japan-Guam-Australia North (JGA-N) expected Q1 2020

  (September 13, 2019) RTI Connectivity (RTI) has announced that the installation on the Japan-Guam-Australia North (JGA-N) cable system between Minamiboso (Japan) and Piti (Guam) has begun. The system will have an initial design capacity of 24Tbps, which is expected to be expanded significantly by adopting a higher optical signal modulation format using the latest available NEC transponders. JGA-N is expected to be commercially available in Q1 2020.

  JGA-N will interconnect at Guam with Japan-Guam-Australia South (JGA-S) system. The JGA-S will consist of a main trunk between Piti (Guam) and Sydney (Australia), which will have a total length of approximately 7,081km and consist of two fibre pairs with a design capacity of a minimum of 18Tbps per fibre pair. (source: https://www.telegeography.com)

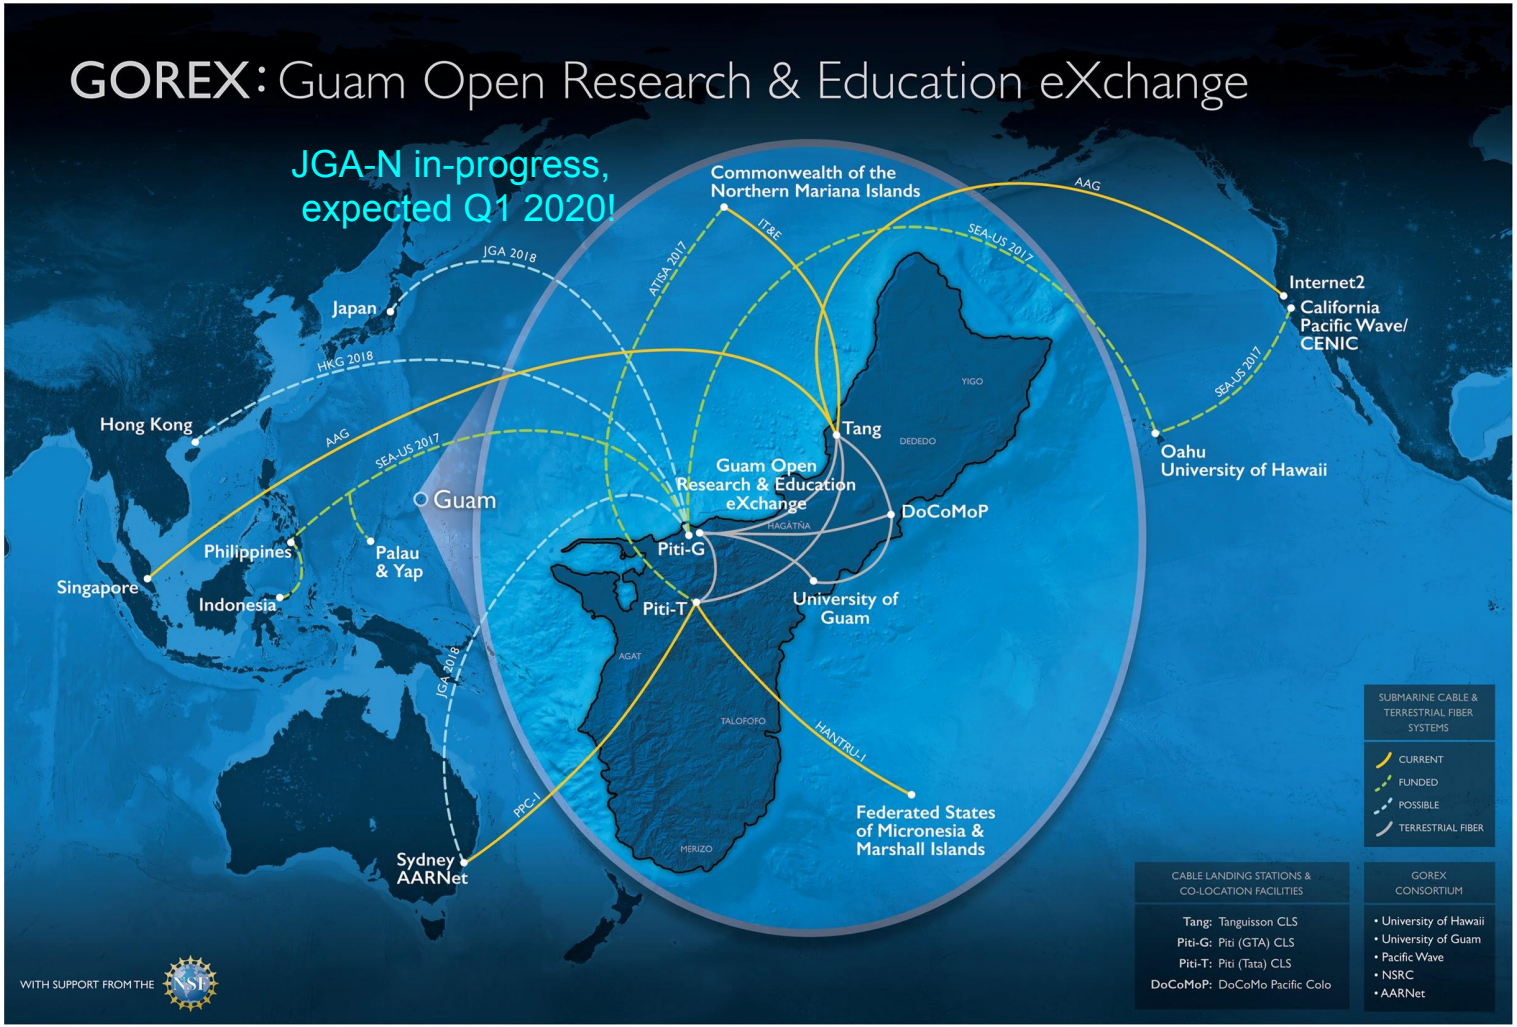

PACIFIC WAVE

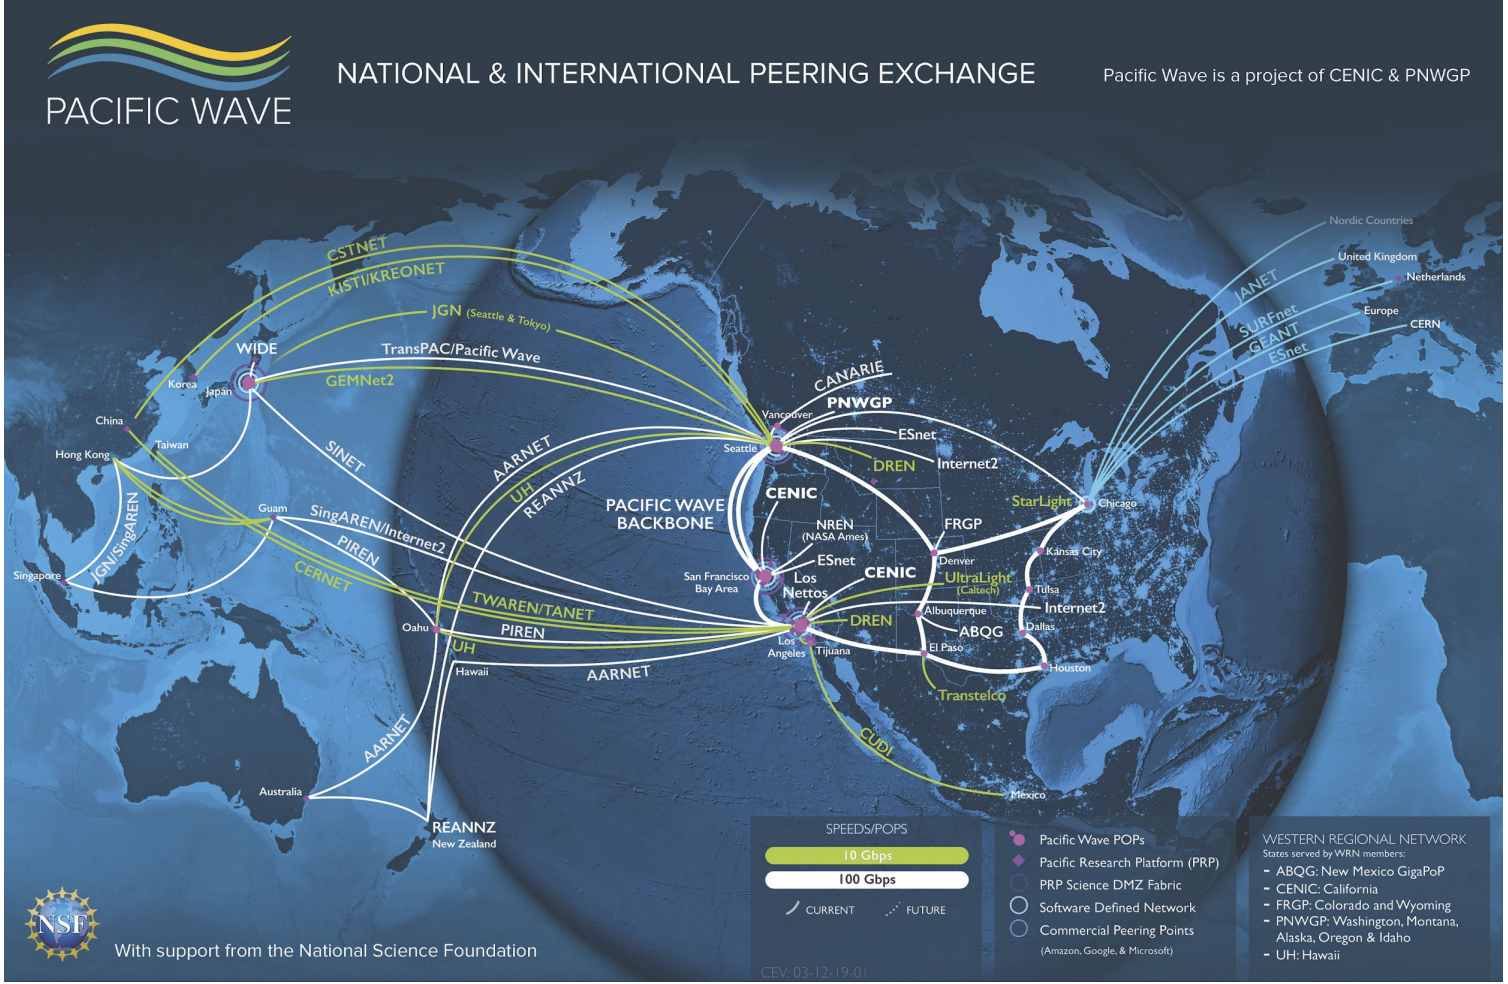
PACIFIC WAVE
NATIONAL & INTERNATIONAL PEERING EXCHANGE
Pacific Wave is a project of CENIC & PNWGP
CSTNET
KISTI/KREONET
JGN (Seattle & Tokyo)
TransPAC/Pacific Wave
GEMNet2
WIDE
Korea
Japan
China
Taiwan
Hong Kong
SINET
IGN/SingAREN
Singapore
Guam
SingAREN/Internet2
PIREN
CERNET
TWAREN/TANET
Oahu
UH
Hawaii
AARNET
AARNET
Australia
REANNZ
New Zealand
AARNET
UH
REANNZ
PACIFIC WAVE BACKBONE
CANARIE
PNWGP
Vancouver
Seattle
ESnet
DREN
Internet2
CENIC
NREN (NASA Ames)
ESnet
San Francisco Bay Area
Los Nettos
CENIC
DREN
Los Angeles
Tijuana
StarLight
Chicago
FRGP
Denver
Kansas City
UltraLight (Caltech)
Tulsa
Albuquerque
Internet2
ABQG
Dallas
El Paso
Houston
Transtelco
CUDI
Mexico
Nordic Countries
United Kingdom
Netherlands
Europe
CERN
IANET
SURFnet
GEANT
ESnet
With support from the National Science Foundation
SPEEDS/POPS
10 Gbps
100 Gbps
CURRENT
FUTURE
Pacific Wave POPs
Pacific Research Platform (PRP)
PRP Science DMZ Fabric
Software Defined Network
Commercial Peering Points
(Amazon, Google, & Microsoft)
WESTERN REGIONAL NETWORK
States served by WRN members:
- ABQG: New Mexico GigaPoP
- CENIC: California
- FRGP: Colorado and Wyoming
- PNWGP: Washington, Montana, Alaska, Oregon & Idaho
- UH: Hawaii
CEV: 03-12-19-01

Pacific Wave: NSF IRNC RXP Award

Pacific Wave Expansion Supporting SDX & Experimentation (ACI-1451050, September 2015)

- Continued enhancement, upgrade and evolution of Pacific Wave to support more 100G connections
- Additional 100G capacity between exchanges points along West Coast
- SDN/SDX testbed deployment on parallel infrastructure to enable experimentation while maintaining production use of the Pacific Wave exchange
- Instrumentation for measurement, monitoring, and analysis: including perfSONAR, measurement visualization with MaDDash
- Collaboration with IRNC awardees on SDX development

Pacific Wave SDX

Testbed Resources Other Uses (DTNs, pS nodes, ...)	Science Driver Applications
Network Testbeds Circuit Building (NSI/OpenNSA)	SDX middleware OpenFlow Controllers (ONOS, Ryu, ...)
SDX Switches On-ramp Locations (Ethernet / virtual circuits)	

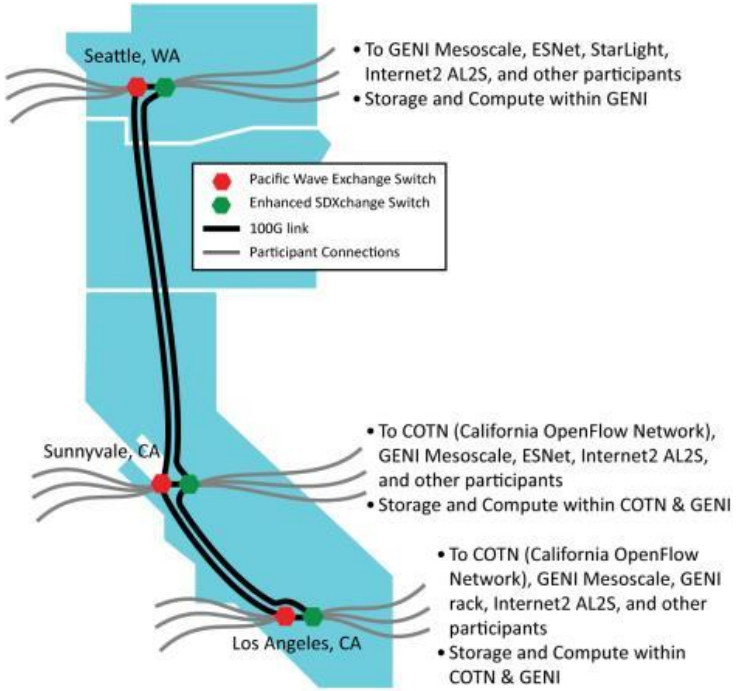

Traditional Peerings between all Pacific Wave participants can remain on the Pacific Wave Exchange Switches.

SDX Services can be accessed via direct connection to Enhanced SDXchange Switches *OR* via connections to the Pacific Wave Exchange Switches.

Mixed connections can also take place between participants where one is connected to the Pacific Wave Exchange Switch and the other to an Enhanced SDXchange Switch.

Pacific Wave: Platform and infrastructure-attached resources

Platform:

Existing -- Extreme (formerly Brocade) MLXe, both for production and SDX testbed

RFP / tender pending -- Juniper MX10008:

- Initially for Seattle, Sunnyvale, and Los Angeles
- Will explore ‘node slicing’ for collapsing production, testbed, __? functions in same chassis

Resources:

(3) 100G-connected DTNs / K8s Storage Nodes

- Dual-socket – (2) E5-2667v4 8-cores @ 3.2GHz w/256GB DRAM
- 6.4TB raw on NVMe -- (2) Kingston (Liqid) 3.2TB NVMe PCIe 3.0 x8
- 80TB raw on SAS3 – (8) 10TB HGST off of an LSI 9300-8i HBA
- Mellanox ConnectX-5 – interconnect using QSFP28 passive copper DACs in Sunnyvale and Los Angeles, and QSFP28 LR4-lite (2km) optic in Seattle

(3) 100G-connected pS nodes -- moving to replace original (QLogic) NICs with Mellanox ConnectX-5

(3) 10G-connected pS nodes

(3) 10G-connected ‘dynamic’ pS / K8s master nodes

(3) 10G-connected x86 Hypervisor VM servers

Pacific Wave: SDX/SDN testbed control plane

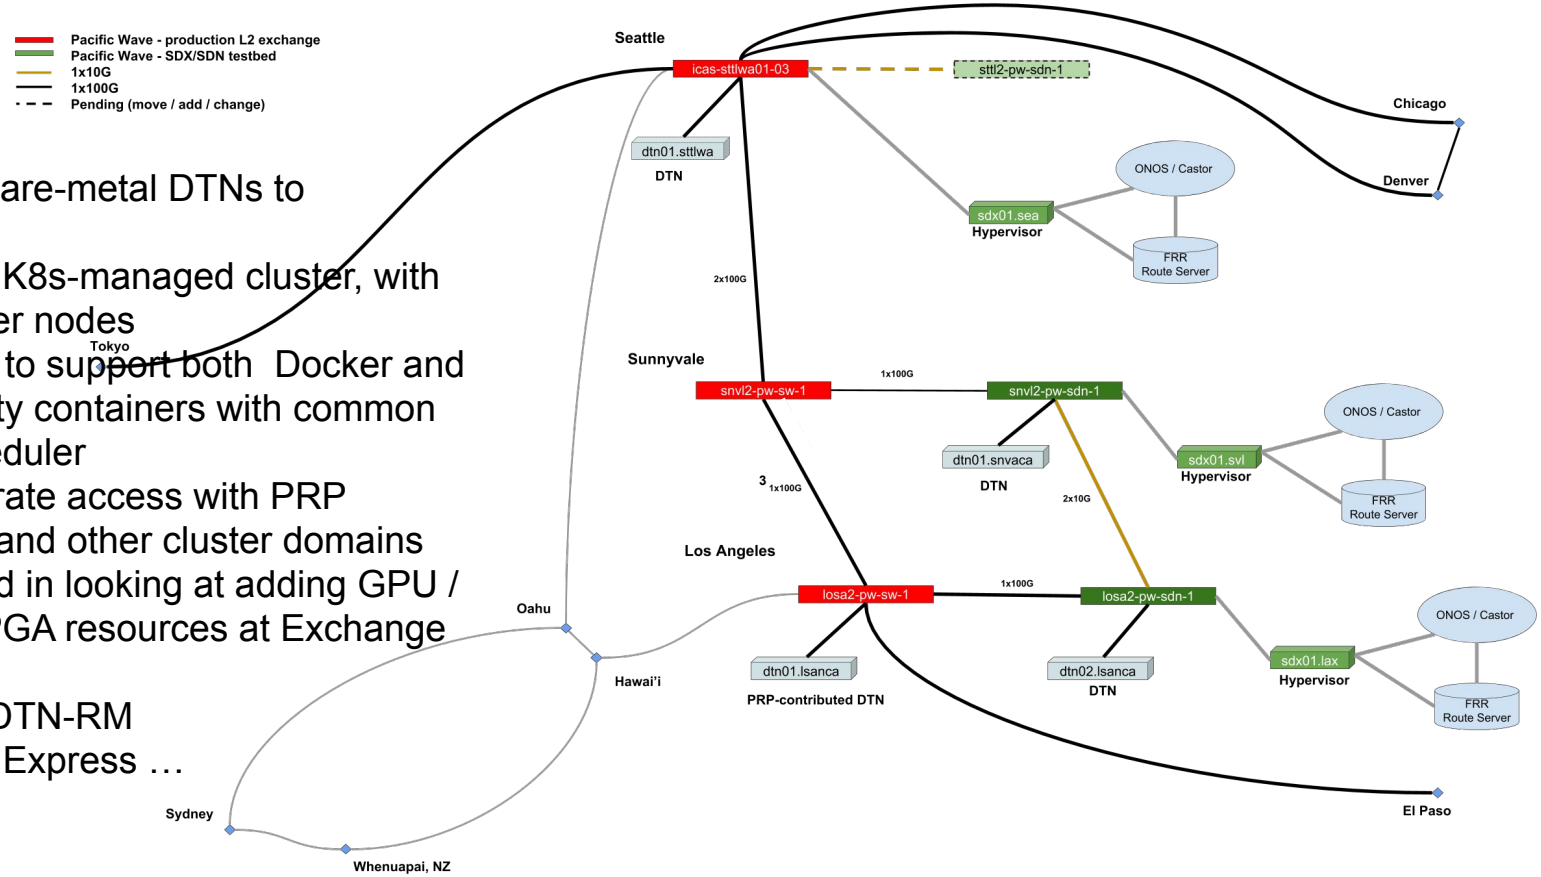

Migrating the bare-metal DTNs to containers:

- Will be a K8s-managed cluster, with HA master nodes
- Planning to support both Docker and Singularity containers with common K8s scheduler
- Will federate access with PRP Nautilus and other cluster domains
- Interested in looking at adding GPU / TPU / FPGA resources at Exchange Points
- SENSE DTN-RM
- Big Data Express …
- __?

NOTE: this diagram represents a subset of sites, devices, and connections

Automated GOLE

- **AutoGOLE fabric delivers dynamic network services among GOLEs and participating networks**
  - GOLE - GLIF Open Lightpath Exchange
- Based on NSI Connection Service v2.0
  - Architectural standard developed by Open Grid Forum (OGF)
    - https://www.ogf.org/documents/GFD.212.pdf
  - Redundant Aggregator backbone with a leaf ultimate provider agent (uPA) per network (Hub and Spoke architecture)
  - 29 Network Service Agents (6 Aggregators, 23 uPA) advertising 30 networks
- Using Document Distribution Service (DDS) for NSA discovery and document propagation between aggregators
- Introduction of monitoring, troubleshooting, and provisioning tools
- Dashboard, MEICAN, DDS Portal, etc
- Advanced capabilities
  - Experimenting with new path finding and signaling algorithms
  - Additional network modeling for optimizations

Pacific Wave Dynamic Circuit Services: AutoGOLE / NSI

- NSI-orchestrated circuit services available to participants traversing each of the Seattle, Sunnyvale, Los Angeles, and Tokyo GOLEs
- In our current implementation each Pacific Wave GOLE in the pilot operates as its own NSI domain, e.g. Los Angeles as lsanca.pacificwave.net:2016
- NSA -- OpenNSA, separate OpenNSA instances for each GOLE; each instance managing a single device; each instances configured for c-plane & d-plane peering with adjacent Pacific Wave GOLE
- C-plane peering (with NSI Aggregators) -- ESnet, NetherLight, StarLight, SINET
- D-plane peering -- ESnet, StarLight, SINET, JGN-X, Caltech, and with Calit2 - UCSD (pending)
- Provisioning -- we are participating in the pilot of RNP's MEICAN webUI
  - https://wiki.rnp.br/display/secipo/AutoGOLE+MEICAN+Pilot
- Participation -- contributors from engineering, operations, and some found cycles toward DevOps

Atlantic Pacific Research and Education Exchange (AP-REX)

A Pilot Project of Internet2 and Pacific Wave (CENIC and PNWGP)

PRPv2 BGP Pilot: route servers for control plane

http://pacificresearchplatform.org

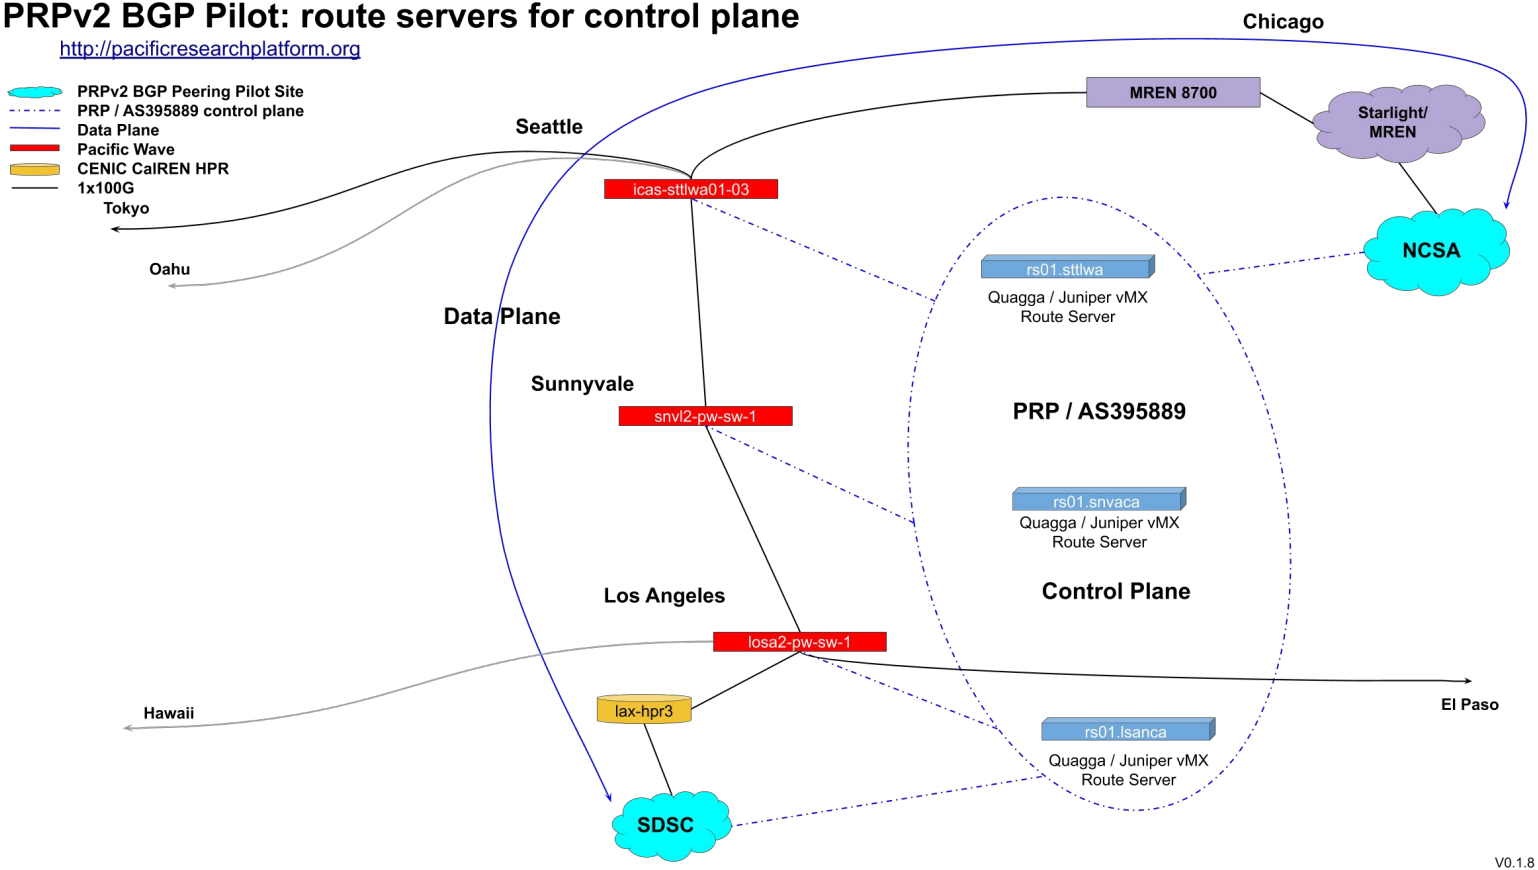

V0.1.8
20170510

NOTE: this diagram represents a subset of sites, devices, and connections

Routing Security and Trust models

- Mutually Agreed Norms for Routing Security (MANRS) https://www.manrs.org
- A global initiative, supported by the Internet Society, toward reducing the most common threats to the routing ecosystem
- MANRS actions for Network Operators / Internet Service Providers (ISPs)
  - Filtering -- Prevent propagation of incorrect routing information
  - IP source validation -- Prevent traffic with spoofed source IP addresses
  - Coordination -- Facilitate global operational communication and coordination between network operators
  - Global validation -- Facilitate validation of routing information on a global scale
- Resource Key Public Infrastructure (RPKI)
  - Certificates verify that a resource has been assigned to a specific entity
  - Route Origin Authorization (ROA) - a cryptographically-signed record that associate a BGP route announcement with the correct originating AS number

MANRS RPKI regional pilot

- Pilot will focus on facilitating MANRS adoption and validation of routing information by implementing Resource Public Key Infrastructure (RPKI) on a regional scale among CENIC and Pacific Wave research universities.
- The pilot is a collaborative effort involving contributors from CENIC, NSRC, ESnet, ARIN, as well as from the CENIC research university community.
- Phase One will focus on participation from California-based research institutions. We will also seek participation from Pacific Wave collaborators outside of California, including the University of Washington and its regional network, the PNWGP (Pacific NorthWest GigaPop).
- Results will be part of a public discussion at the 2020 CENIC Annual Conference

CENIC MANRS RPKI Project Phase 1- v0.04

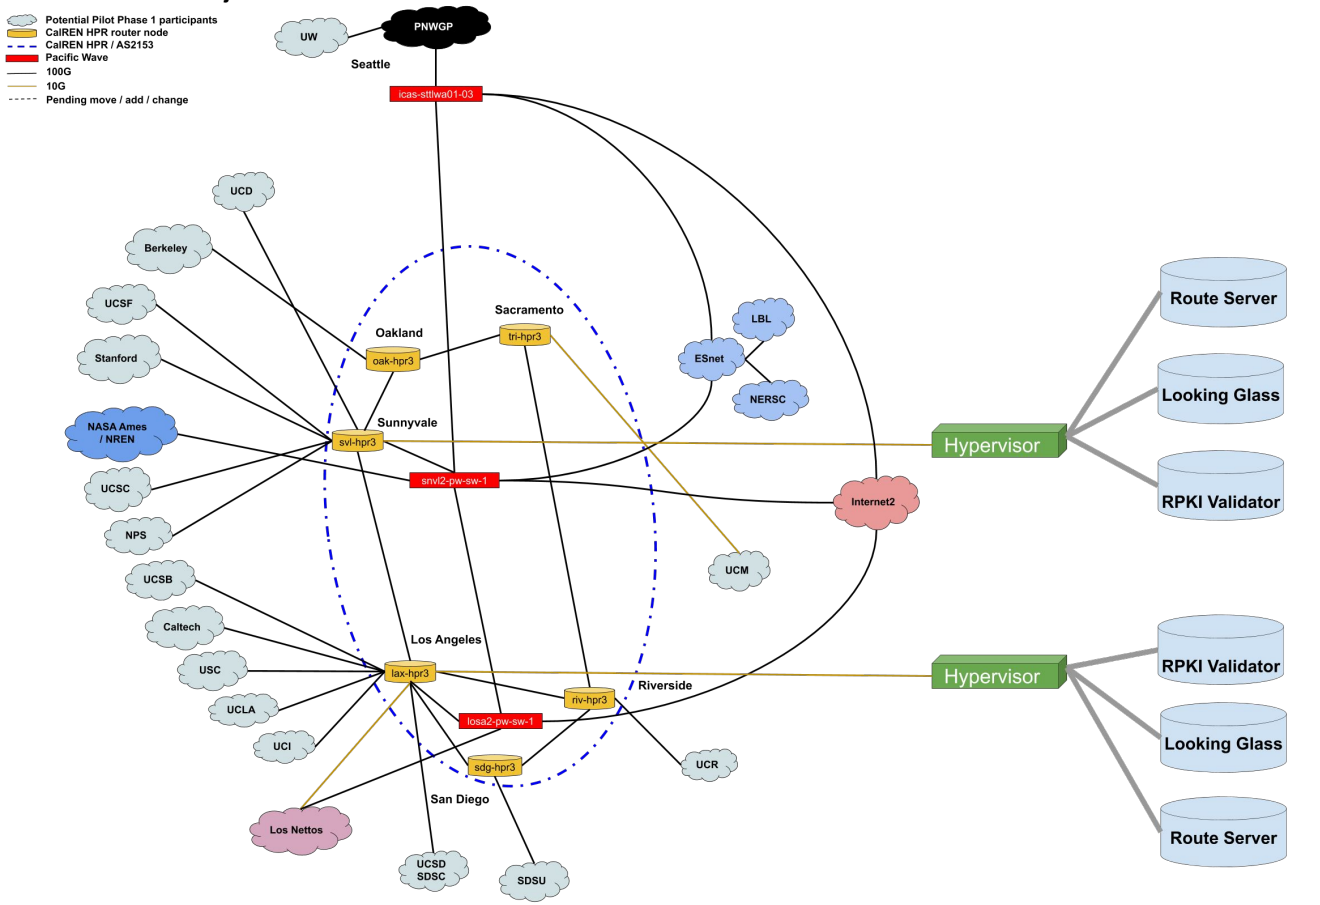

RPKI Validator instances synchronize their local ROA database with the (RIRs) trust anchors

Route Servers interact with RPKI Validators, using ROA validation status as a hook for determining BGP policy. Route Servers facilitate BGP policy for routing platforms which do not support for RPKI, and provide routing telemetry and other data to the Looking Glass instances

Looking Glass instances provide monitoring and debugging tools to network operators and participants

Pacific Wave - Western Region Network: GXP Route-Servers with RPKI pilot

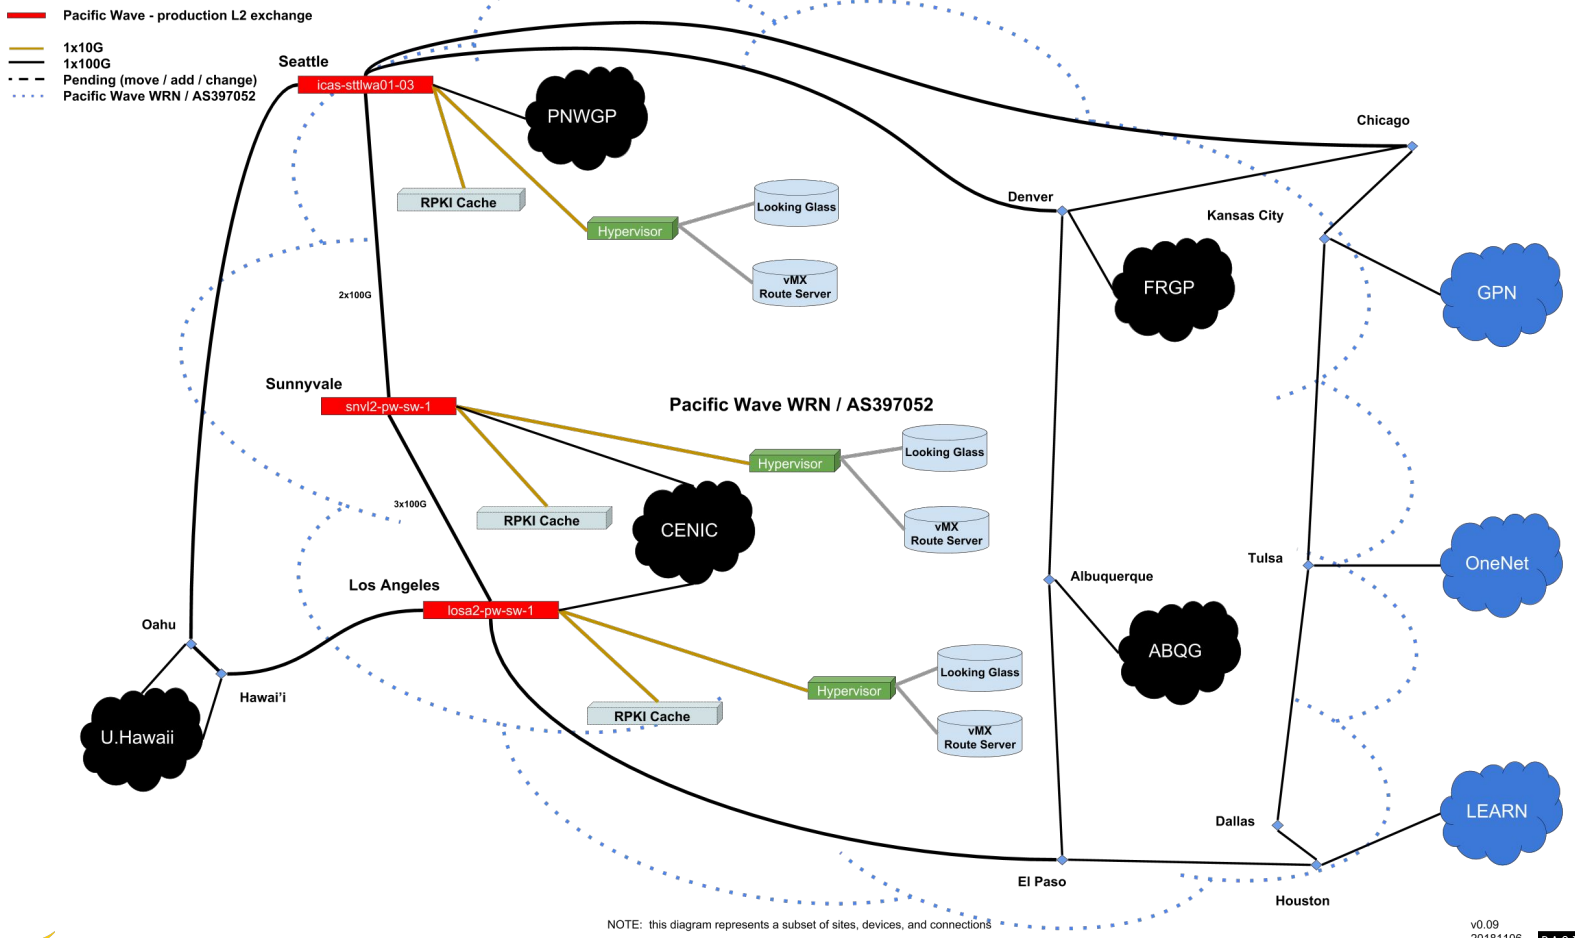

v0.09
20181106

Questions?

Thanks!