SDX Interoperability

Global Research Platform
Calit2 UC San Diego
September 17, 2019

John Hess
Network Engineer - Research Engagement
CENIC - Pacific Wave

CENIC

PACIFIC WAVE

Agenda

- CENIC and Pacific Wave
- SDX Technologies and Services
- Routing Ecosystem
- Security for Routing and Infrastructure Services

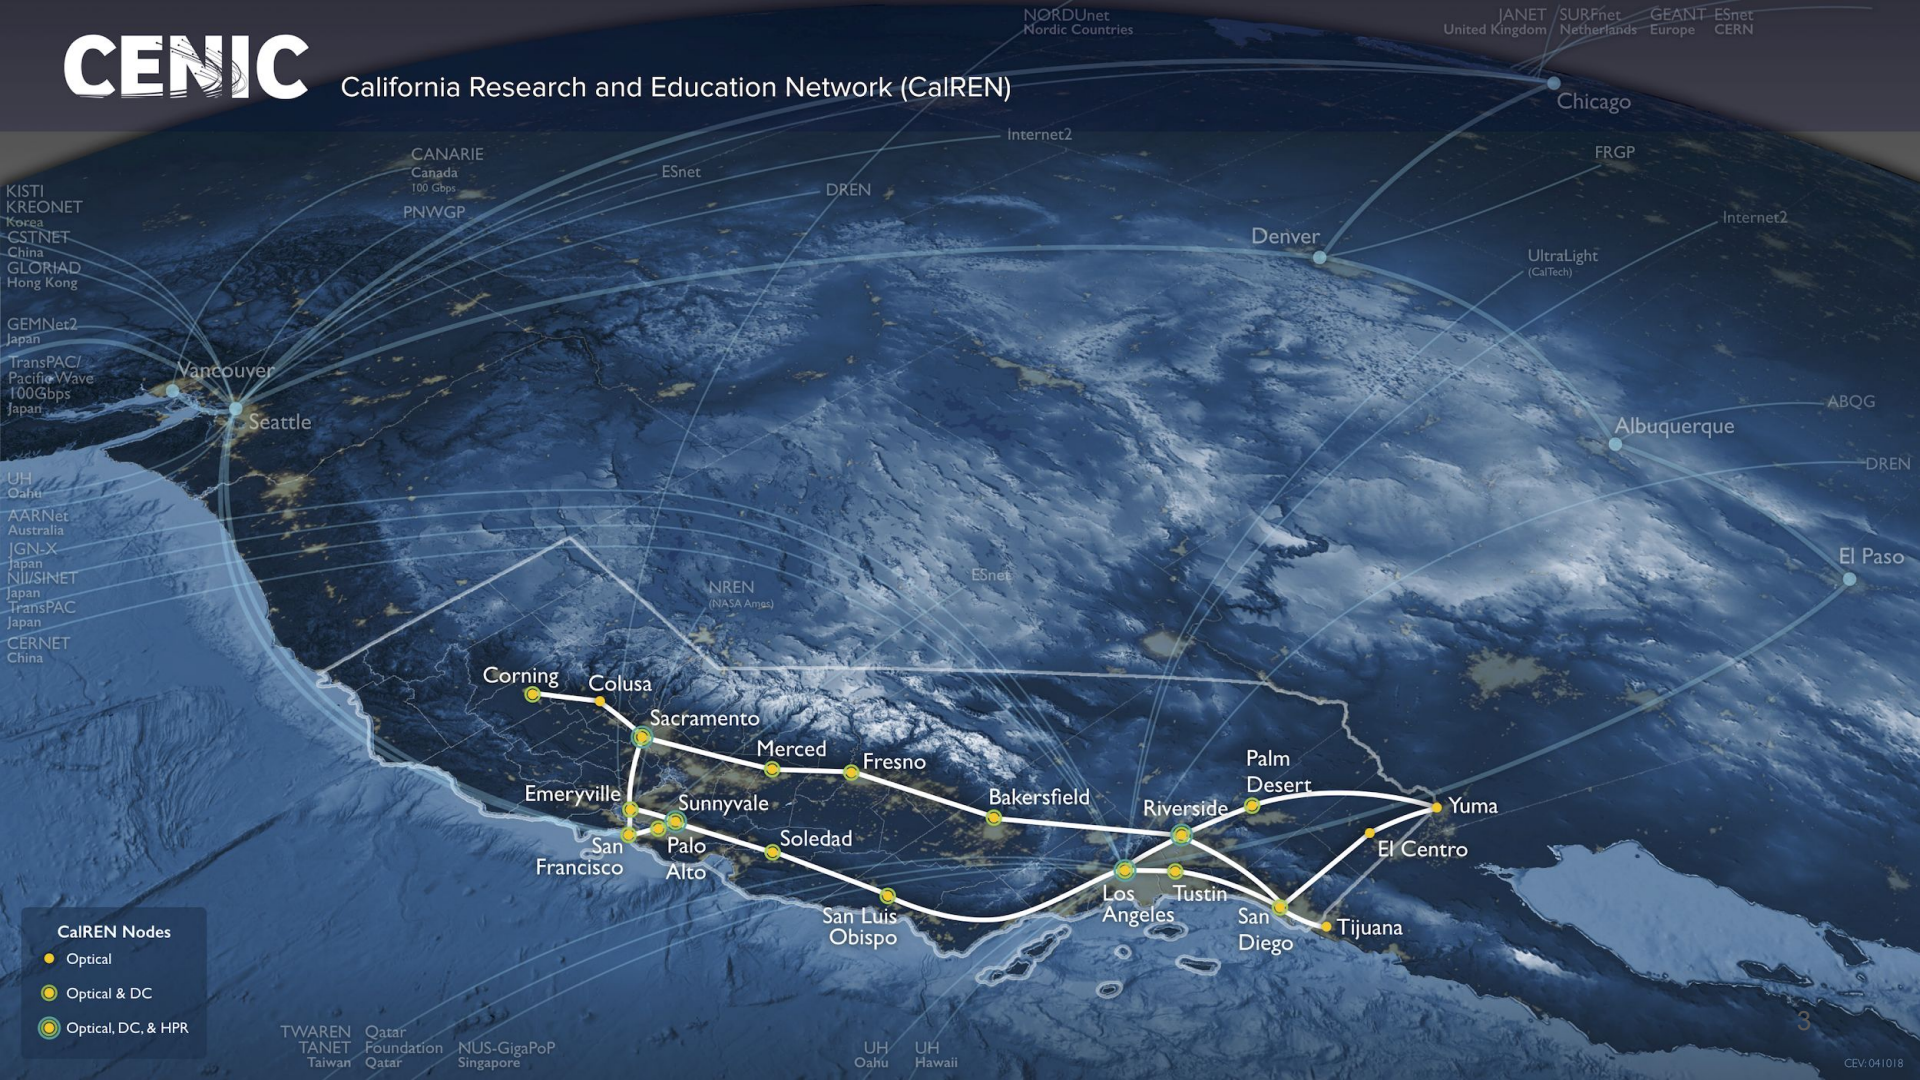
CENIC
California Research and Education Network (CalREN)
NORDUnet
Nordic Countries
JANET
United Kingdom
SURFnet
Netherlands
GEANT
Europe
ESnet
CERN
Chicago
Internet2
CANARIE
Canada
100 Gbps
ESnet
DREN
FRGP
KISTI
KREONET
Korea
CSTNET
China
GLORIAD
Hong Kong
PNWGP
Denver
Internet2
UltraLight
(CalTech)
GEMNet2
Japan
TransPAC/
Pacific Wave
100Gbps
Japan
Vancouver
Seattle
ABQG
Albuquerque
DREN
UH
Oahu
AARNet
Australia
JGN-X
Japan
NII/SINET
Japan
TransPAC
Japan
CERNET
China
El Paso
NREN
(NASA Ames)
ESnet
Corning
Colusa
Sacramento
Merced
Fresno
Palm
Desert
Emeryville
Sunnyvale
Bakersfield
Riverside
Yuma
San
Francisco
Palo
Alto
Soledad
El Centro
Los
Angeles
Tustin
San Luis
Obispo
San
Diego
Tijuana
CalREN Nodes
Optical
Optical & DC
Optical, DC, & HPR
TWAREN
TANET
Taiwan
Qatar
Foundation
Qatar
NUS-GigaPoP
Singapore
UH
Oahu
UH
Hawaii
3
CEV:041018

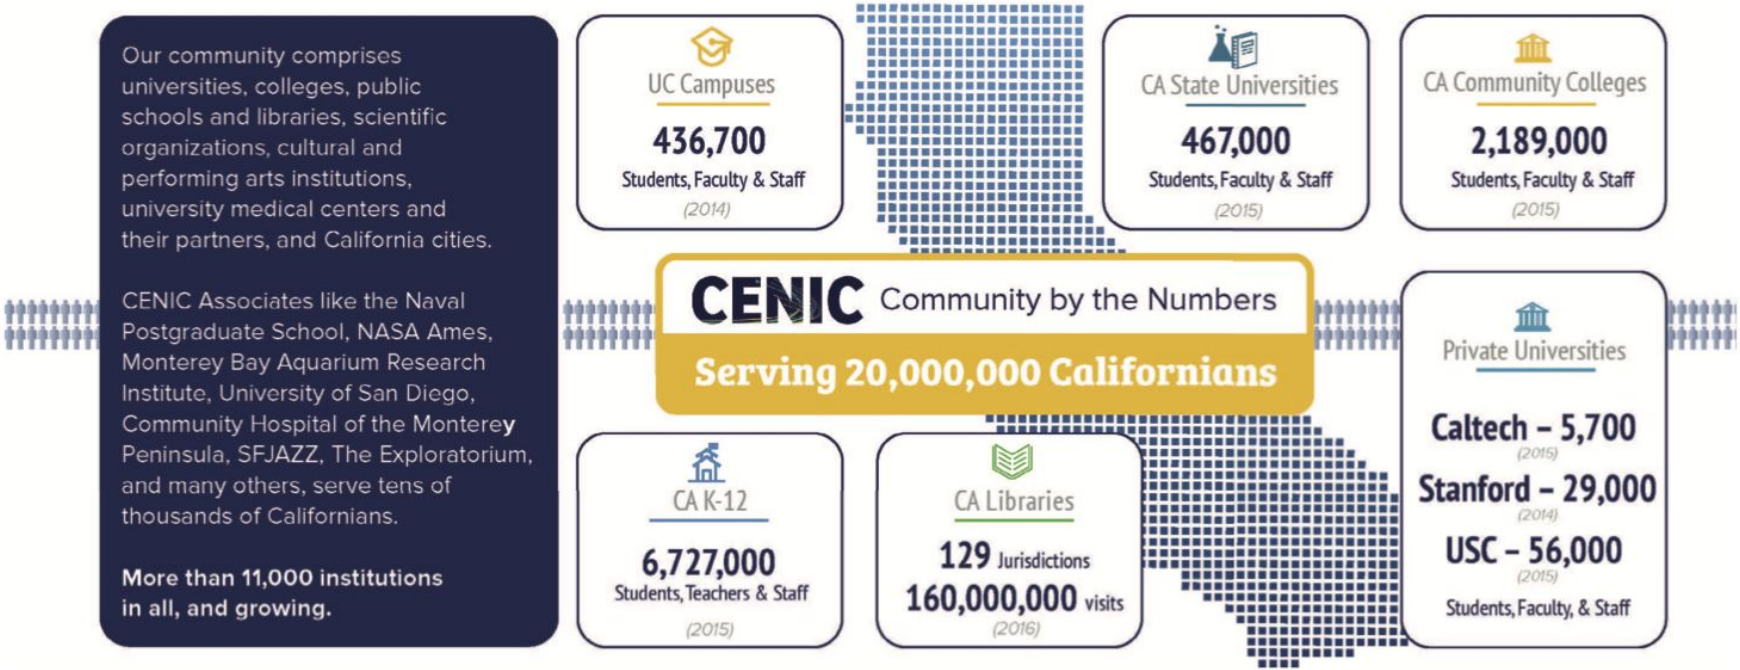
Our community comprises universities, colleges, public schools and libraries, scientific organizations, cultural and performing arts institutions, university medical centers and their partners, and California cities.
CENIC Associates like the Naval Postgraduate School, NASA Ames, Monterey Bay Aquarium Research Institute, University of San Diego, Community Hospital of the Monterey Peninsula, SFJAZZ, The Exploratorium, and many others, serve tens of thousands of Californians.
More than 11,000 institutions in all, and growing.
UC Campuses
436,700
Students, Faculty & Staff
(2014)
CA State Universities
467,000
Students, Faculty & Staff
(2015)
CA Community Colleges
2,189,000
Students, Faculty & Staff
(2015)
CENIC Community by the Numbers
Serving 20,000,000 Californians
CA K-12
6,727,000
Students, Teachers & Staff
(2015)
CA Libraries
129 Jurisdictions
160,000,000 visits
(2016)
Private Universities
Caltech – 5,700
(2015)
Stanford – 29,000
(2014)
USC – 56,000
(2015)
Students, Faculty, & Staff

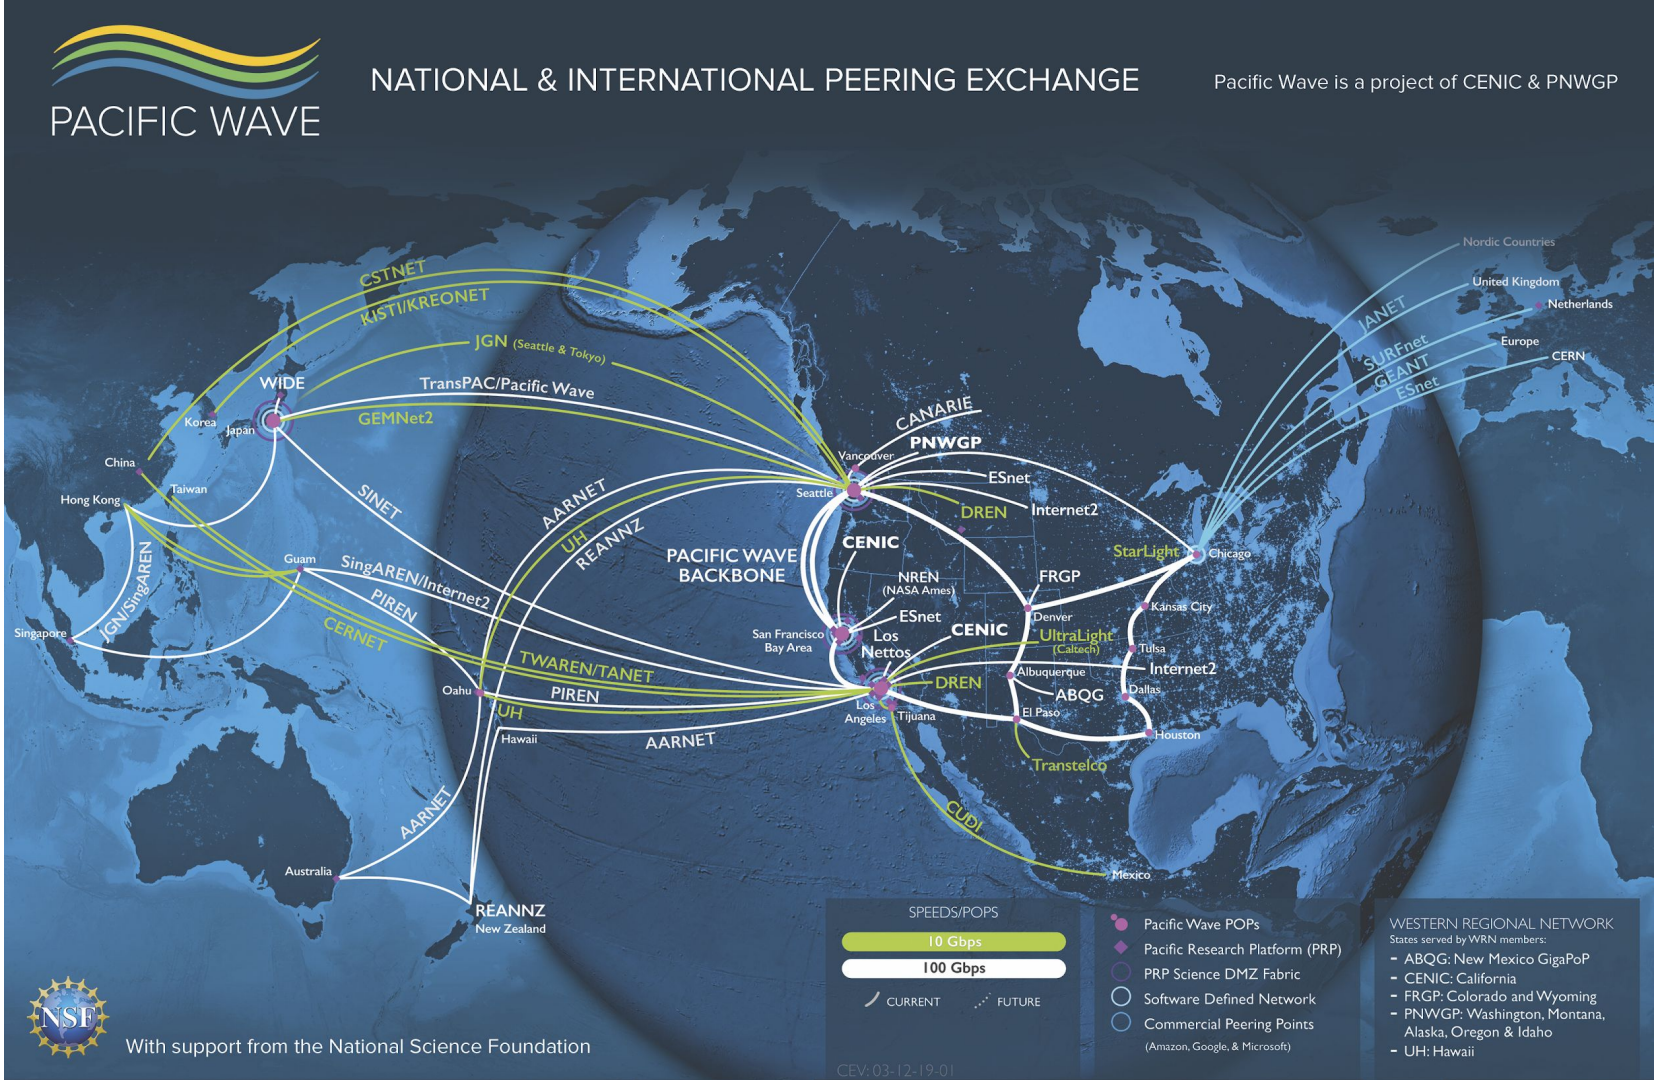
PACIFIC WAVE
NATIONAL & INTERNATIONAL PEERING EXCHANGE
Pacific Wave is a project of CENIC & PNWGP
CSTNET
KIST/KREONET
JGN (Seattle & Tokyo)
TransPAC/Pacific Wave
GEMNet2
WIDE
Korea
Japan
China
Taiwan
Hong Kong
SINET
Guam
SingAREN/Internet2
PIREN
CERNET
JGN/SingAREN
Singapore
AARNET
UH
REANNZ
TWAREN/TANET
PIREN
UH
Oahu
Hawaii
AARNET
AARNET
Australia
REANNZ
New Zealand
CANARIE
PNWGP
Vancouver
Seattle
ESnet
DREN
Internet2
CENIC
PACIFIC WAVE BACKBONE
NREN (NASA Ames)
ESnet
San Francisco Bay Area
Los Nettos
CENIC
DREN
Los Angeles
Tijuana
FRGP
Denver
UltraLight (Caltech)
Albuquerque
ABQG
El Paso
Transtelco
CUDI
Mexico
StarLight
Chicago
Kansas City
Tulsa
Internet2
Dallas
Houston
Nordic Countries
United Kingdom
Netherlands
Europe
CERN
JANET
SURFnet
GEANT
ESnet
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
With support from the National Science Foundation

Pacific Wave - International Peering Exchange

- A project CENIC and Pacific Northwest Gigapop (PNWGP), Pacific Wave was deployed as a geographically distributed peering facility in January 2004
- Open Exchange Points supporting both commercial and R&E peers
- Pacific Wave has been partially supported through three separate five-year National Science Foundation (NSF) grants supporting growth connectivity and innovation
- Currently serves 31 countries across the Pacific connecting to the Western USA
- Enables science-driven high-capacity data-centric projects, such as the Pacific Research Platform (PRP), enabling researchers to move data between collaborator sites, supercomputer centers, and campus Science DMZs without performance degradation

Agenda

- CENIC and Pacific Wave
- SDX Technologies and Services
- Routing Ecosystem
- Security for Routing and Infrastructure Services

AARNet - Pacific Wave - StarLight Inter-domain SDX Topology high-level v0.06.03A

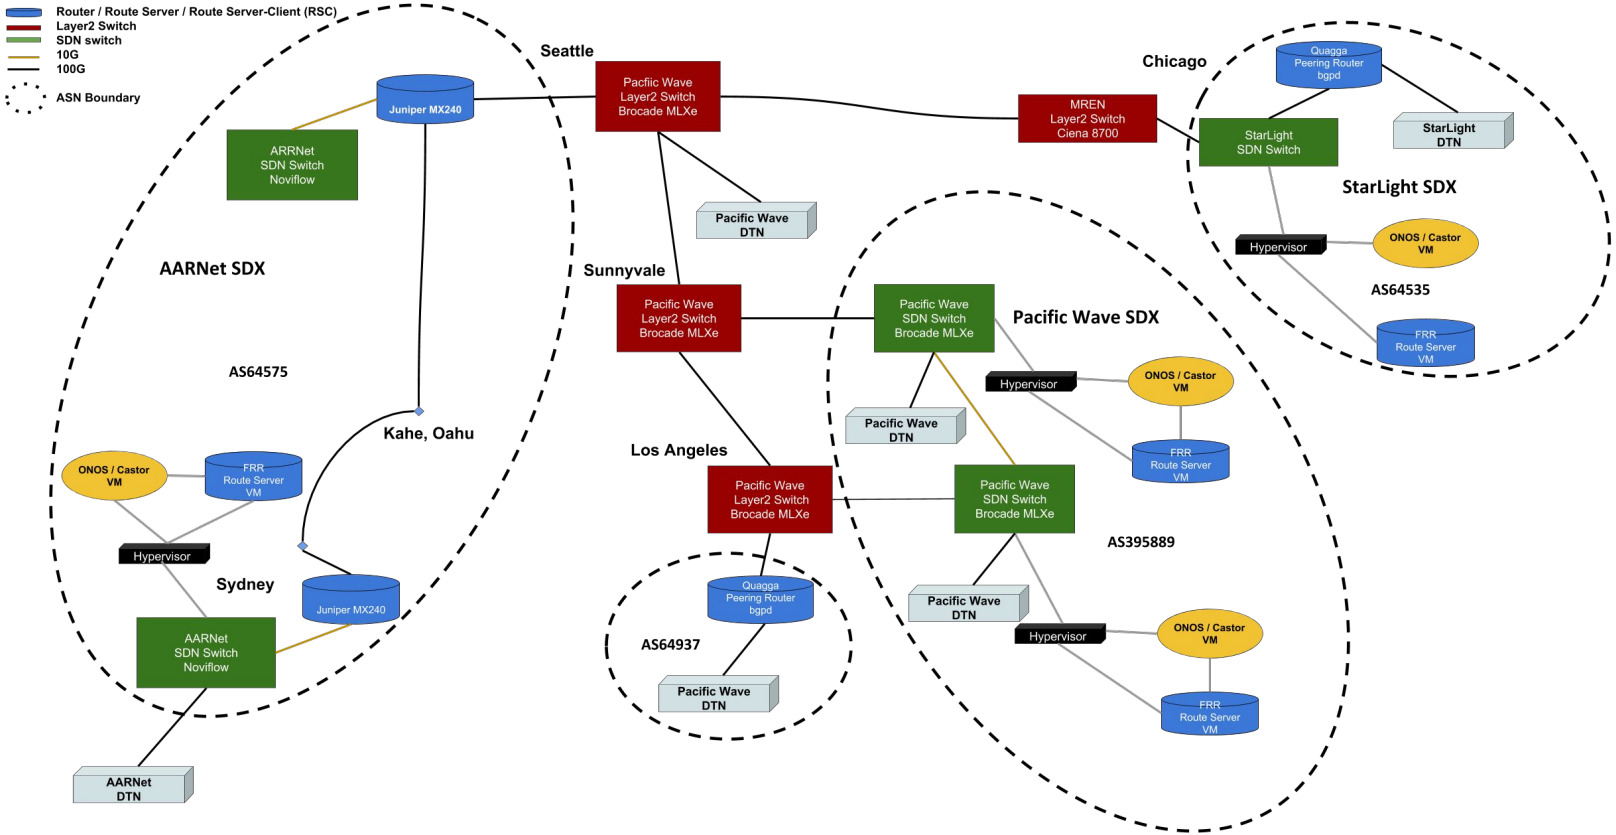

NOTE: this diagram represents a subset of sites, devices, and connections

V0.06.02A 20180306

SDX Technologies and Services - Forward Looking

- Capacity
  - Core moving from N x 100Gbps to 200Gbps and toward 400Gbps
- Segment Routing: SR-MPLS vs SRv6
  - Who remembers source-based routing? Some workflows desire or require an explicit path.
  - SR may be implemented as a control-plane architecture to provide traffic engineering (TE) at the ingress router. TE policy to steer a packet through a specific set of nodes and links in the network.
  - SR-MPLS uses a label stack to describe the desired path through the network. Examine the label; do a lookup; pop the label; forward the packet. SR-MPLS requires some specific control-plane software, but does not impact ASICs ability to forward packets.
  - SRv6 uses a Segment Routing Header (SRH) embedded as a new header within an IPv6 packet to describe the desired path through the network. SRv6-capable routers will want (or need) new ASICs and/or programmable silicon (e.g. Juniper Penta). non-SRv6 capable routers ignore the SRH and forward the packet normally.

SDX Technologies and Services

- Orchestration
  - Dynamic circuit and services provisioning
    - AutoGOLE / NSI + MEICAN (RNP!); and, SENSE & Big Data Express
  - NFV, vCPE, IaaS
  - Containers: Docker, Singularity, OCI. Federated across cloud-providers, within-network, at member institutions
- Federated access to Infrastructure-attached resources
  - Inter-domain and inter-cluster: access to storage, CPU / GPU / TPU, FPGA
  - ICN / NDN -- heiracheral, caching of research data, e.g. LHC/HEP, genomics, weather. SANDIE: https://www.nsf.gov/awardsearch/showAward?AWD_ID=1659403

Agenda

- CENIC and Pacific Wave
- SDX Technologies and Services
- Routing Ecosystem
- Security for Routing and Infrastructure Services

Routing Ecosystem

- IANA/ICANN
- Internet Routing Registries (IRRs) / Regional Internet Registries (RIRs)
  - http://www.irr.net/docs/list.html
  - American Registry for Internet Numbers (ARIN)
  - Merit's Routing Arbiter Database (RADb)
- Network Operators
- Internet Exchange Points (IXPs)
- Border Gateway Protocol
- Infrastructure Services and Tools

Agenda

- CENIC and Pacific Wave
- SDX Technologies and Services
- Routing Ecosystem
- Security for Routing and Infrastructure Services

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

Routing Security and Trust models

Mutually Agreed Norms for Routing Security (MANRS)

MANRS

- MANRS actions for Internet eXchange Points (IXPs)
  - Prevent propagation of incorrect routing information
  - Promote MANRS to the IXP membership
  - Protect the peering platform
  - Facilitate global operational communication and coordination between network operators
  - Provide monitoring and debugging tools to the members

Routing Security and Trust models -- Implementation

Resource Public Key Infrastructure (RPKI)

- Resource certificates digitally 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
- Defined in RFC 6480 (An Infrastructure to Support Secure Internet Routing)
- Currently five RIRs (AFRINIC, APNIC, ARIN, LACNIC & RIPE) provide a method for members to take an IP/ASN pair and sign a ROA
- These five RIRs act as Trust Anchors (TAs) -- similar to Certificate Authorities (CAs) -- to validate ROAs

Routing Security and Trust models -- Implementation

RPKI high-level implementation elements

- Origin validation uses X.509 certificates with extensions specified in RFC 3779
- RPKI Cache server (RPKI Validator) synchronizes its local database with TAs
- RPKI Validator exports a simplified version of ROAs as Route Validation (RV) records
- An RV record is a (prefix, maximum length, origin AS) triple
- Router interacts with cache server to query the RV status for a given route
- Router implements Border Gateway Protocol (BGP) policy based on validation status indicated in the RV record for the route

CENIC MANRS - RPKI regional pilot

- Pilot will focus on facilitating MANRS adoption and validation of routing information by implementing 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

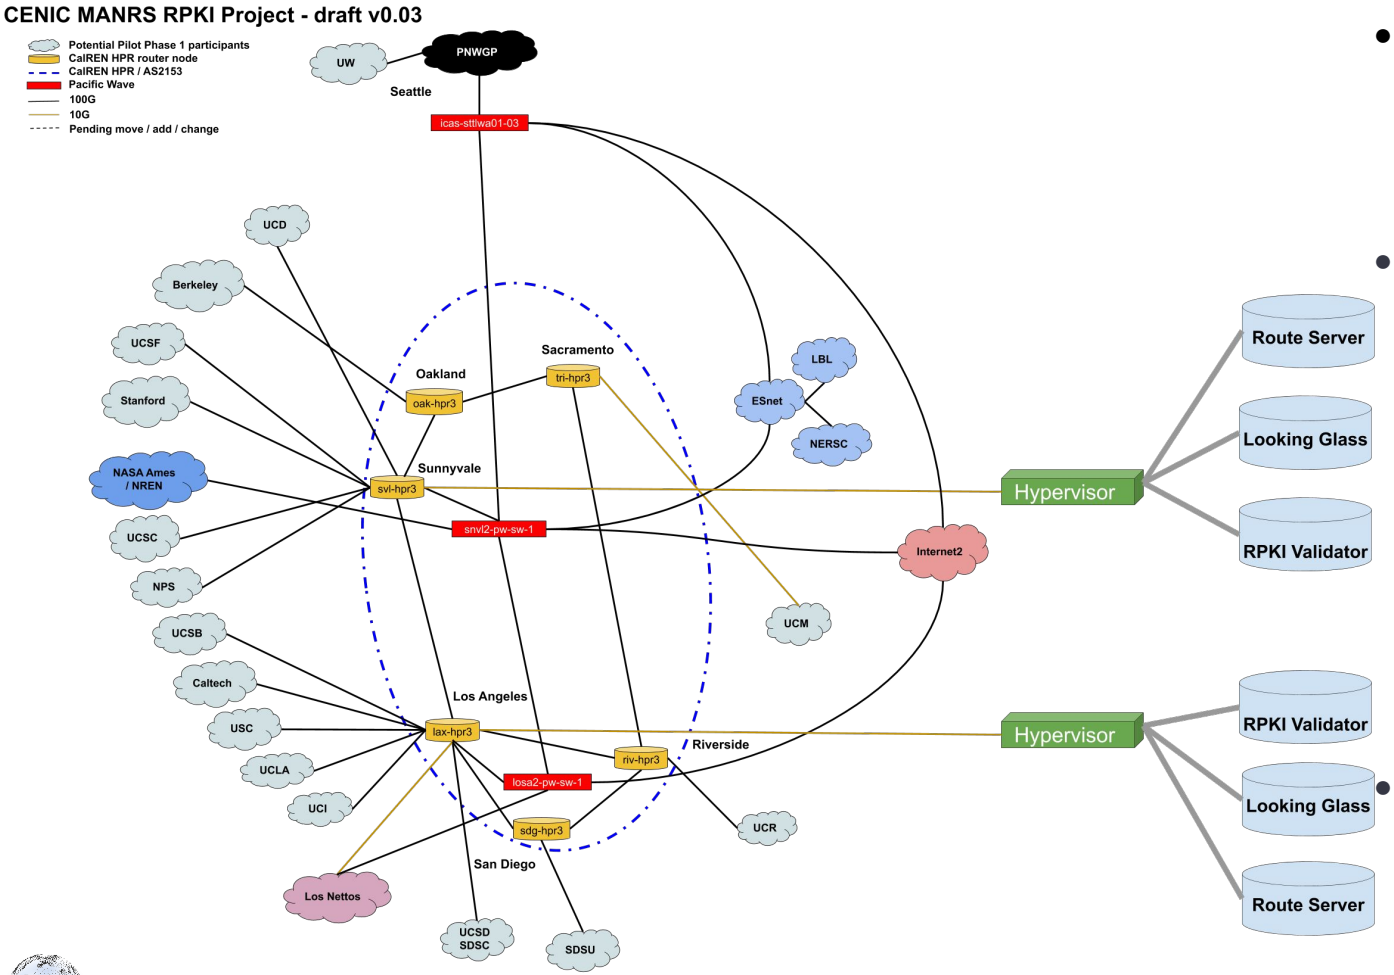

- RPKI Validator instances synchronize their local ROA database with the (RIRs) trust anchors
- Route Servers interact with RPKI Validators, using ROA validation status as a hook for determining BGP policy. Route Servers facilitate BGP policy for routing platforms which do not support for RPKI, and provide routing telemetry and other data to the Looking Glass instances
- Looking Glass instances provide monitoring and debugging tools to network operators and participants

CENIC MANRS - RPKI regional pilot

- Phase Two will seek to expand participation to include Western Regional Network members (University of Hawaii, Front Range Gigapop, New Mexico Gigapop), their respective research universities, as well as other key partners (Oregon Fiber Partnership, LEARN, OneNet, Great Plains Network, Nevada, and others) and their respective research universities. Phase Two will commence, based on the outcomes of Phase One, in the spring of 2020.
- Stretch Goals: Though this project will focus on implementing RPKI for route origin validation, we intend to investigate implementing complementary technologies for enhancing the security of core cyberinfrastructure functions. Examples include: filtering route announcements based on Internet Routing Registry (IRR) data; and, supporting Domain Name System Security Extension (DNSSEC).

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

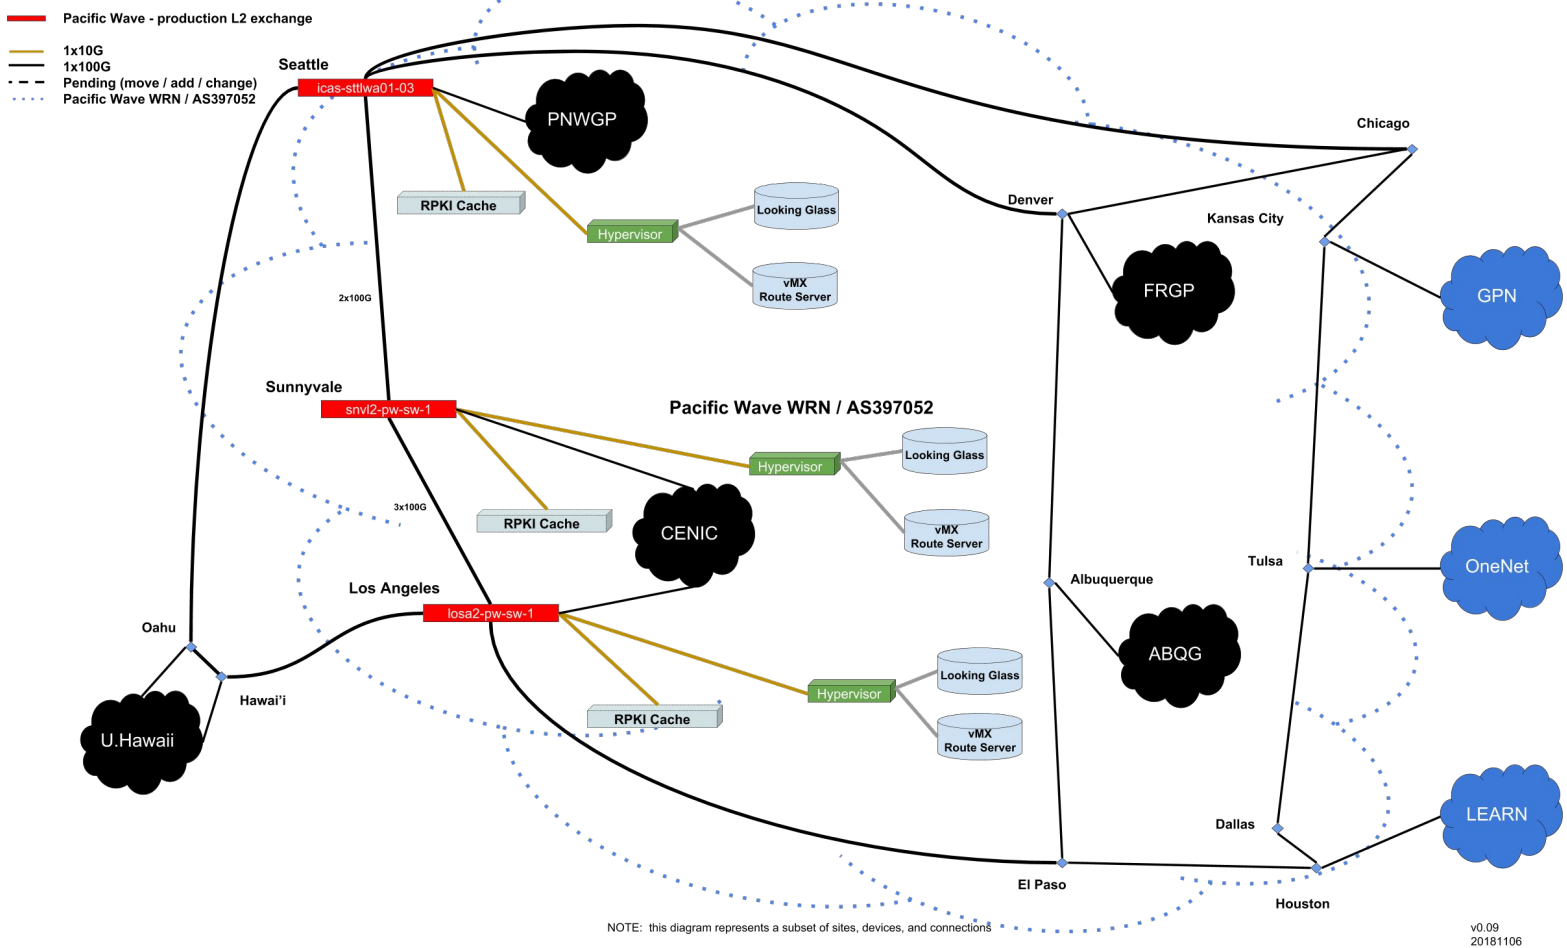

Securing Infrastructure Services, and Tools

Considerations for Network Operators and IXPs (not an exhaustive list)

- DNS (Domain Name System):
  - DNSSEC (DNS Security) https://www.internetsociety.org/deploy360/dnssec/basics/
- BGP: Implement RPKI, then add BGPsec?
  - https://www.internetsociety.org/blog/2017/10/bgpsec-reality-now/
  - BGPsec is implemented via an optional non-transitive BGP path attribute called BGPsec_Path, that carries digital signatures produced by each autonomous system propagating the update message.
- BGP monitoring tools:
  - Artemis: https://www.inspire.edu.gr/artemis/
  - Looking Glass

CENIC MANRS - RPKI regional pilot

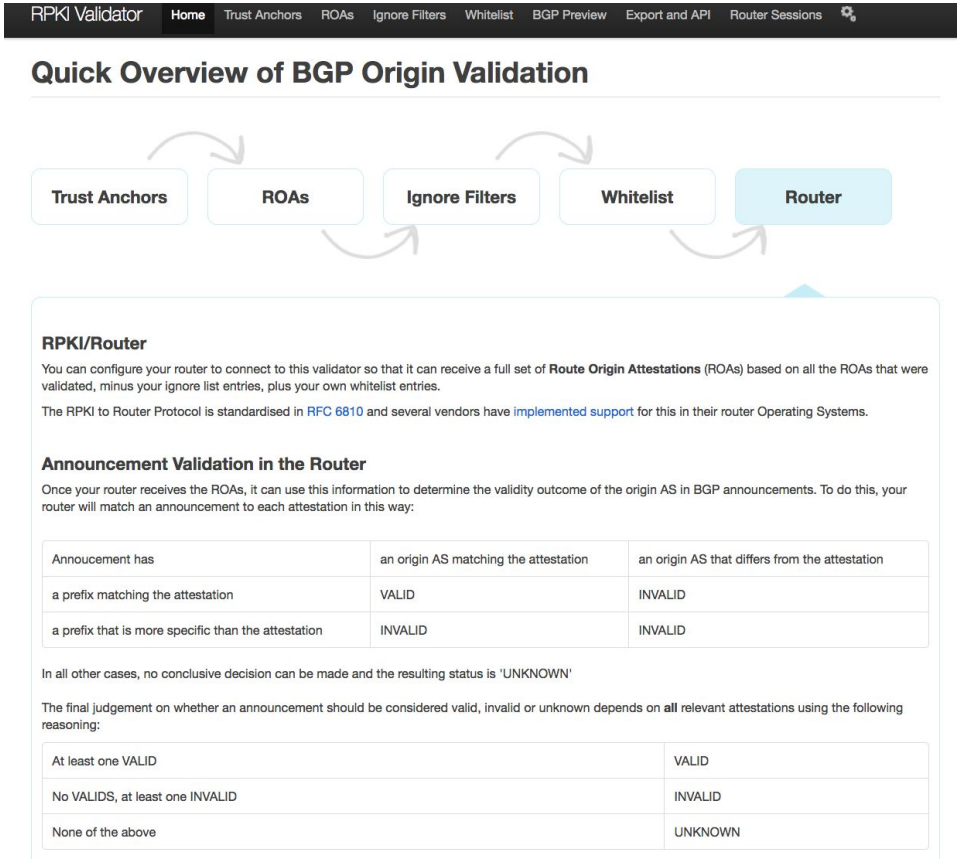


```
root@lax-vmx0> show validation session
Session                         State  Flaps    Uptime #IPv4/IPv6 records
2607:f380:2:8030::20              Up       0  17:21:23 56591/9933

root@lax-vmx0> show validation statistics
Total RV records: 66524
Total Replication RV records: 66524
  Prefix entries: 62340
  Origin-AS entries: 66524
Memory utilization: 12801844 bytes
....
root@lax-vmx0> show validation database
RV database for instance master

Prefix                Origin-AS Session                       State   Mismatch
1.0.0.0/24-24          13335 2607:f380:2:8030::20             valid
1.1.1.0/24-24          13335 2607:f380:2:8030::20             valid
1.9.0.0/16-24           4788 2607:f380:2:8030::20             valid
1.9.12.0/24-24         65037 2607:f380:2:8030::20             valid
1.9.21.0/24-24         24514 2607:f380:2:8030::20             valid
1.9.23.0/24-24         65120 2607:f380:2:8030::20             valid
1.9.31.0/24-24         65077 2607:f380:2:8030::20             valid
1.9.65.0/24-24         24514 2607:f380:2:8030::20             valid
….
```

Routing Security and Trust models -- Participating

- ARIN's Operational Test and Evaluation Environment (OT&E): https://www.arin.net/resources/ote.html
  - Allows experimentation and research within ARIN's main services, including RPKI, without affecting production data
  - Available to organizations which have signed a Registration Services Agreement (RSA) as well as those organizations which have not signed a RSA
  - Example: test Route Authorization Requests (ROAs) can be created and validated without impacting production RPKI data
- RPKI cache server (Validator) application:
  - RIPE:https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources
  - NLnet Labs Routinator: https://www.nlnetlabs.nl/projects/rpki/routinator/

Routing Security and trust models -- Participating

- Configuring Origin Validation for BGP:
  - Cisco: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s-book/irg-origin-as.pdf
  - Juniper: https://www.juniper.net/documentation/en_US/junos12.2/topics/topic-map/bgp-origin-as-validation.html
- Configure and Apply policies (implementation subjective to Network Operator)
  - Assign a local-preference to the RPKI validity attribute of the prefix
  - IETF standards discussions including setting local-preference to prefer valid over unknown, and valid and unknown over invalid
  - Apply extended community tags to signal ROA validity state to iBGP peers

Quick Overview of BGP Origin Validation

Trust Anchors

ROA

Trust anchors are the entry points used for val

This RPKI Validator is preconfigured with the t
repository, you will first have to accept their R
application.

RIPE NCC Copyright (

Configured Trust Anchors

Enabled	Trust anchor	Processed Items	Expires in	Last updated	Next update in	Update all
☑	APNIC RPKI Root	6715 0 0	4 years and 9 months	8 minutes ago	2 minutes	Update
☑	ARIN	3939 0 0	8 years and 9 months	5 minutes ago	5 minutes	Update
☑	AfriNIC RPKI Root	549 0 0	8 years and 9 months	10 minutes ago	43 seconds	Update
☑	LACNIC RPKI Root	5183 0 0	93 years and 9 months	1 minute ago	9 minutes	Update
☑	RIPE NCC RPKI Root	26517 0 0	98 years and 11 months	1 minute ago	9 minutes	Update

RIPE NCC Copyright ©2009-2018 the Réseaux IP Européens Network Coordination Centre RIPE NCC. All rights restricted. Version 2.25

RPKI Validator | Home | Trust Anchors | ROAs | Ignore Filters | Whitelist | BGP Preview | Export and API | Router Sessions

Validated ROAs

Validated ROAs from **APNIC RPKI Root**, **ARIN**, **AfriNIC RPKI Root**, **LACNIC RPKI Root**, **RIPE**

Show 10 entries

ASN	Prefix	Maximum Length
0	209.24.0.0/24	24
42	74.80.64.0/18	24
42	204.61.208.0/21	24
42	2001:500:15::/48	48
42	74.63.16.0/20	24
42	2001:500:14::/48	48
42	204.61.216.0/23	23
42	206.220.228.0/22	24
42	2620:0:870::/45	48
87	129.79.0.0/16	16

First | Previous | 1 | 2 | 3 | 4 | 5 | Next | Last

RIPE NCC Copyright ©2009-2018 the Réseaux IP Européens Network Coordinatic

RPKI Validator | Home | Trust Anchors | ROAs | Ignore Filters | Whitelist | BGP Preview | Export and API | Router Sessions

BGP Preview

This page provides a **preview** of the likely RPKI validity states your routers will associate with BGP announcements. This preview is based on:

- The RIPE NCC Route Collector information that was last updated 8 hours and 34 minutes ago.
- BGP announcements that are seen by 5 or more peers.
- The validation rules defined in RFC 6483.
- The validated ROAs found by this RPKI Validator after applying your filters and additional whitelist entries.

Please note that the BGP announcements your routers see may differ from the ones listed here.

Show 10 entries Search: 15169

ASN	Prefix	Validity
15169	89.207.231.0/24	VALID
15169	104.132.34.0/24	VALID
15169	185.25.28.0/23	VALID
15169	2620:0:1000::/40	VALID
15169	2620:15c::/36	VALID
15169	2a00:79e0::/32	VALID
15169	8.8.4.0/24	UNKNOWN
15169	8.8.8.0/24	UNKNOWN

MANRS RPKI References

MANRS:
https://www.manrs.org/wp-content/uploads/sites/14/2018/10/Routing-Security-for-Policymakers-EN.pdf

https://www.manrs.org/isps/

https://www.manrs.org/ixps/manrs-actions-for-ixps/

RPKI:

https://blog.cloudflare.com/rpki/

http://www.sanog.org/resources/sanog30/SANOG30-Tutorial_rpsl-rpki-sanog30.pdf

https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources

https://www.arin.net/resources/rpki

https://tools.ietf.org/html/rfc8210

Tools: https://www.inspire.edu.gr/artemis/

Questions?