

Global BCP for Routing Policy and Security

Philip Smith

Network Startup Resource Center

AfNOG 2021

31st May 2021



UNIVERSITY OF OREGON



Securing Routing BCPs

Industry initiative: MANRS (<https://www.manrs.org>)

- Filters on EBGP sessions
 - Prevent propagation of incorrect routing information
- Block traffic with spoofed source addresses
 - BCP 38 / unicast reverse path forwarding on access interfaces
- Communication between network operators
 - PeeringDB, route objects, AS objects, and NOC contact details up to date
- Validation of routing information
 - Route origination authorisation and validation

Focus on RPKI

Route Origin Authorisation

- Digital signature indicating the origin AS of an announced route
- Well established globally
- Widely deployed with more and more operators signing ROAs every week

Route Origin Authorisation

- A typical ROA would look like this:

Prefix	10.10.0.0/16
Max-Length	/18
Origin-AS	AS65534

- There can be more than one ROA per address block
 - Allows the operator to originate prefixes from more than one AS
 - Caters for changes in routing policy or prefix origin

Where to see ROAs?

RouteViews! (<http://www.routeviews.org>)

```
route-views>show ip bgp rpki table
184405 BGP sovc network entries using 29504800 bytes of memory
204243 BGP sovc record entries using 6535776 bytes of memory
```

Network	Maxlen	Origin-AS	Source	Neighbor
1.0.0.0/24	24	13335	0	184.171.101.188/3323
1.0.0.0/24	24	13335	0	184.171.101.187/3323
1.0.4.0/24	24	38803	0	184.171.101.188/3323
1.0.4.0/24	24	38803	0	184.171.101.187/3323
1.0.4.0/22	22	38803	0	184.171.101.188/3323
1.0.4.0/22	22	38803	0	184.171.101.187/3323
1.0.5.0/24	24	38803	0	184.171.101.188/3323
1.0.5.0/24	24	38803	0	184.171.101.187/3323
1.0.6.0/24	24	38803	0	184.171.101.188/3323
1.0.6.0/24	24	38803	0	184.171.101.187/3323
1.0.7.0/24	24	38803	0	184.171.101.188/3323
1.0.7.0/24	24	38803	0	184.171.101.187/3323
1.1.1.0/24	24	13335	0	184.171.101.188/3323
1.1.1.0/24	24	13335	0	184.171.101.187/3323
1.1.4.0/22	22	4134	0	184.171.101.188/3323
1.1.4.0/22	22	4134	0	184.171.101.187/3323
1.1.16.0/20	20	4134	0	184.171.101.188/3323
1.1.16.0/20	20	4134	0	184.171.101.187/3323
1.2.9.0/24	24	4134	0	184.171.101.188/3323
1.2.9.0/24	24	4134	0	184.171.101.187/3323

```
route-views>show bgp ipv6 unicast rpki table
31138 BGP sovc network entries using 5729392 bytes of memory
33612 BGP sovc record entries using 1075584 bytes of memory
```

Network	Maxlen	Origin-AS	Source	Neighbor
2001:200::/32	32	2500	0	184.171.101.188/3323
2001:200::/32	32	2500	0	184.171.101.187/3323
2001:200:136::/48	48	9367	0	184.171.101.188/3323
2001:200:136::/48	48	9367	0	184.171.101.187/3323
2001:200:1BA::/48	48	24047	0	184.171.101.188/3323
2001:200:1BA::/48	48	24047	0	184.171.101.187/3323
2001:200:900::/40	40	7660	0	184.171.101.188/3323
2001:200:900::/40	40	7660	0	184.171.101.187/3323
2001:200:8000::/35	35	4690	0	184.171.101.188/3323
2001:200:8000::/35	35	4690	0	184.171.101.187/3323
2001:200:C000::/35	35	23634	0	184.171.101.188/3323
2001:200:C000::/35	35	23634	0	184.171.101.187/3323
2001:200:E000::/35	35	7660	0	184.171.101.188/3323
2001:200:E000::/35	35	7660	0	184.171.101.187/3323
2001:218:3002::/48	48	1613	0	184.171.101.188/3323
2001:218:3002::/48	48	1613	0	184.171.101.187/3323
2001:240::/32	32	2497	0	184.171.101.188/3323
2001:240::/32	32	2497	0	184.171.101.187/3323
2001:260::/32	48	2518	0	184.171.101.188/3323
2001:260::/32	48	2518	0	184.171.101.187/3323



Where to see ROAs?



HURRICANE ELECTRIC
INTERNET SERVICES

 Search

(<https://bgp.he.net>)








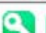






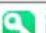














AS2497 Internet Initiative Japan Inc.

Quick Links

- [BGP Toolkit Home](#)
- [BGP Prefix Report](#)
- [BGP Peer Report](#)
- [Exchange Report](#)
- [Bogon Routes](#)
- [World Report](#)
- [Multi Origin Routes](#)
- [DNS Report](#)
- [Top Host Report](#)
- [Internet Statistics](#)
- [Looking Glass](#)
- [Network Tools App](#)
- [Free IPv6 Tunnel](#)
- [IPv6 Certification](#)
- [IPv6 Progress](#)
- [Going Native](#)
- [Contact Us](#)



AS Info Graph v4 Graph v6 Prefixes v4 Prefixes v6 Peers v4 Peers v6 Whois IRR IX

Prefix		Description	
49.239.64.0/18	 	Internet Initiative Japan Inc.	
58.138.0.0/17	 	Internet Initiative Japan Inc.	
58.138.128.0/18	 	Internet Initiative Japan Inc.	
61.211.96.0/19	 	Japan Network Information Center	
101.128.128.0/17	 	Internet Initiative Japan Inc.	
103.2.56.0/24		Rikei Corporation	
103.2.57.0/24	 	Internet Initiative Japan Inc.	
103.2.58.0/23	 	Internet Initiative Japan Inc.	
103.5.23.0/24		Videx Inc.	
113.197.128.0/17	 	Internet Initiative Japan Inc.	
116.118.192.0/20	 	Internet Initiative Japan Inc.	
118.151.0.0/17	 	Internet Initiative Japan Inc.	
118.151.128.0/18	 	Internet Initiative Japan Inc.	
119.10.192.0/18	 	Internet Initiative Japan Inc.	

Number of ROAs



Afrinic



APNIC



ARIN




LACNIC



RIPE NCC


This graph shows the total number of valid Route Origin Authorisation (ROA) objects created by the holders of a certificate



IPv4 address space in ROAs (/24s) 

 Afrinic  APNIC  ARIN  LACNIC  RIPE NCC


This graph shows the amount of IPv4 address space covered by ROAs, in /24 units



IPv6 address space in ROAs (/32s)



This graph shows the amount of IPv6 address space covered by ROAs, in /32 units



Focus on RPKI

- Route Origin Validation
 - Network operator checks if ROA exists for entry in the BGP table
 - If “valid”: **ALLOW**
 - If “invalid”: **DROP**
 - If “NotFound”: ALLOW but low preference

Major Operators deploying RPKI and ROV

- More and more operators are deploying RPKI and ROV
- Not just transit providers!
- But also:
 - Content providers
 - IXPs
 - R&E networks
 - Access providers
- Telia
- NTT
- Lumen (ex L3)
- HE
- GTT
- Workonline
- SEACOM
- Cloudflare
- AMS-IX
- LINX
- DE-CIX
- Terrehost
- Vocus
- Telstra
- REANNZ
- Cogent
- GR-IX
- Swisscom
- Netflix
- UAE-IX
- ...

Route Origin Validation – Deployment

Being cautious... 😊

1. Sign ROAs for **originated** address space
 - Has no negative operational impact
 - Add to standard operating procedure: **if it is originated, sign it!**
2. Observe:
 - Deploy a validator cache
 - EBGP speaking routers talk with the cache
 - Do NOT implement any policy (beware: Cisco IOS/IOS-XE implements policy by default)
 - What are the **invalids**?
3. Full deployment:
 - Deploy redundant validator caches
 - Implement ROV: **drop invalids on all EBGP routers**

Where to see Invalid Routes?

RouteViews! 😊

```
route-views3.routeviews.org> show ip bgp route-map invalid  
<snip>
```

Network	Next Hop	Metric	LocPrf	Weight	Path
* 1.6.161.0/24	202.150.221.33	0	38001	9583	64764 i
* 1.6.172.0/24	38.19.140.162	0	54574	6461	6453 9583 i
* 1.6.177.0/24	202.150.221.33	0	38001	9583	64764 i
* 1.6.219.0/24	38.19.140.162	0	54574	6461	3320 9583 137130 i
* 1.6.229.0/24	38.19.140.162	0	54574	6461	6453 4755 i
* 1.6.230.0/24	38.19.140.162	0	54574	6461	6453 4755 i
* 1.7.178.0/24	38.19.140.162	0	54574	6461	3320 9583 137130 i

```
route-views3.routeviews.org> sh bgp ipv6 route-map invalid  
<snip>
```

Network	Next Hop	Metric	LocPrf	Weight	Path
* 2001:200:e102::/48	2406:f400:8:35::1	0	38001	7660	i
* 2001:388:cf85::/48	2a0c:b641:7af:1::	0	207268	207268	58057 62240 6762 7575 i
* 2001:418:0:5000::3a4/127	2605:8200::5	0	19653	?	
* 2001:418:0:5000::a66/127	2607:ff18:1::a	0	40630	i	
* 2001:418:1401:1e::/64	2001:550:2:67::42:2	0	54574	6461	20940 i
* 2001:418:1401:63::/64	2001:550:2:67::42:2	0	54574	6461	20940 i
* 2001:470:1:6e1::/126	2602:fed2:fc0:e4::1	0	141237	17920	?



Creating ROAs: Best Practices

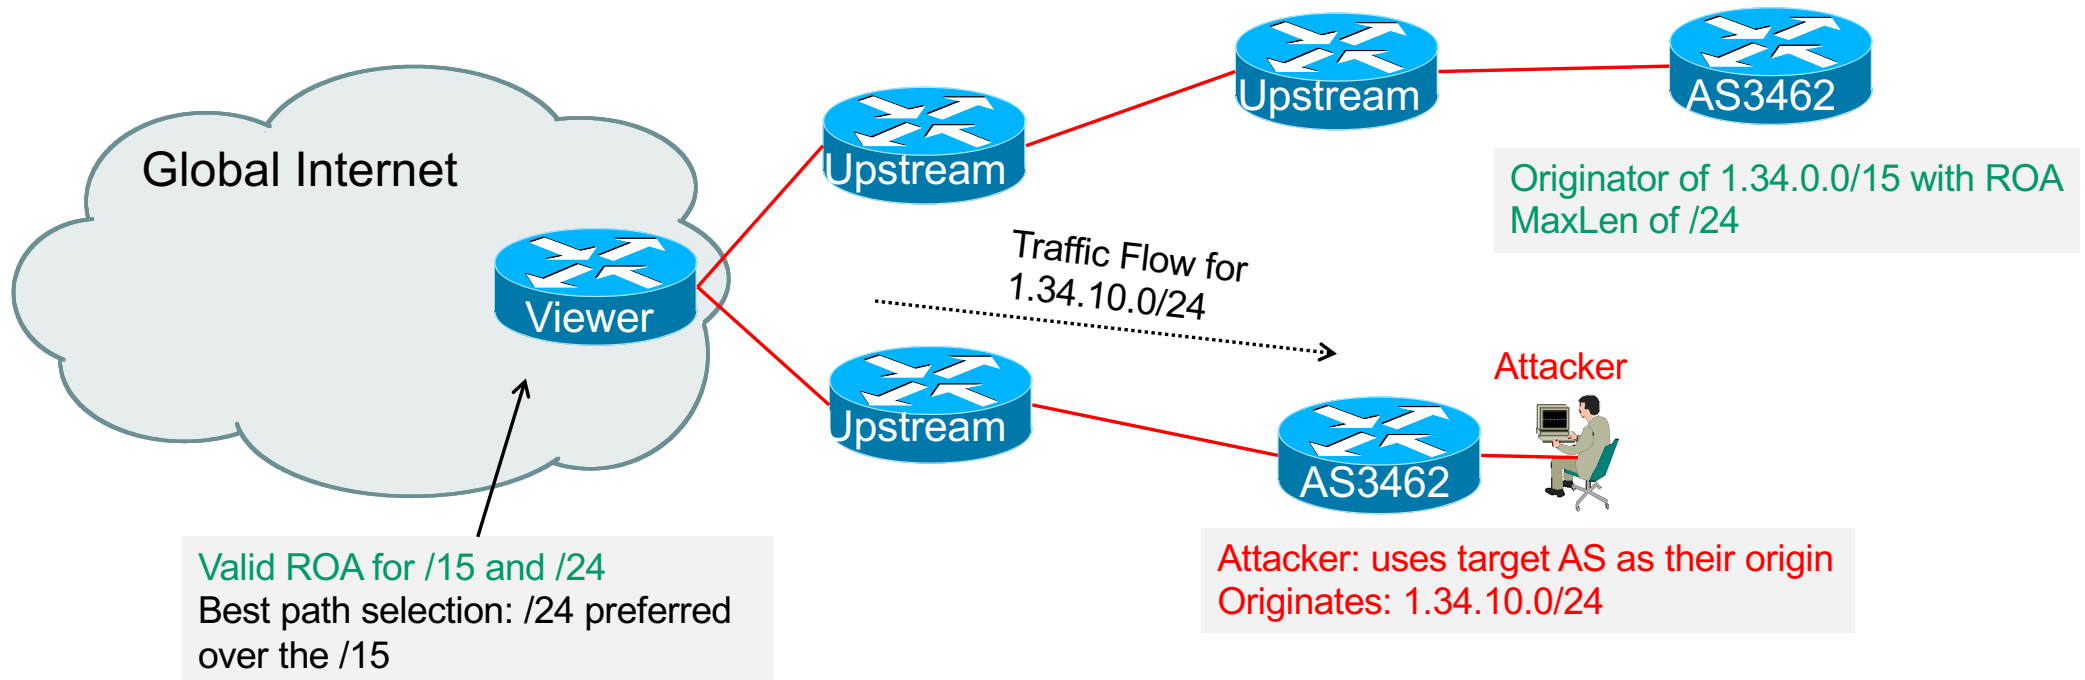
- Only create ROAs for aggregates and individual subnets being announced to the global Internet
- Do **NOT** create ROAs for unannounced subnets
 - This can lead to what is called “validated hijack”
 - Example of a (current) problematic ROA:

```
route-views3.routeviews.org> sh rpki prefix 1.34.0.0/15
Prefix                Prefix Length  Origin-AS
1.34.0.0              15 - 24       3462
```

- This means that any subnet of 1.34.0.0/15 down to a /24 as originated by AS3462 is valid
 - An attacker can use AS3462 as their origin AS to originate 1.34.10.0/24 to deny service to that address block



Creating ROAs: “Validated Hijack”



- If no ROA had been created for the 1.34.10.0/24 prefix, route origin validation would have dropped the invalid announcement at the upstream AS

Route Origin Validation – AS0

- RFC6483 also describes “Disavowal of Routing Origination”
 - AS 0 has been reserved for network operators and other entities to identify non-routed networks
 - Which means:
 - “A ROA with a subject of AS0 (AS0 ROA) is an attestation by the holder of a prefix that the prefix described in the ROA, and any more specific prefix, should not be used in a routing context”
- Any prefixes with ROA indicating AS0 as the origin AS need to be dropped
 - If these prefixes appear with any other origin, their ROAs will be invalid, achieving this goal

Route Origin Validation – AS0

- Possible use cases of AS0:
 - Internal use of a prefix that should not appear in the global BGP table
 - Internet Exchange Point LAN must never appear in the global BGP table
 - Private Address space (IPv4) and non-Global Unicast space (IPv6)
 - Unassigned address space
 - This is under discussion within the various RIR policy fora
 - IPv4 and IPv6 address resources which should not appear in the global BGP table
 - For example, the special use address space described in RFC6890

Route Origin Validation – AS0

- APNIC has now published its AS0 TAL
 - Operated separately from the regular TAL
 - <https://www.apnic.net/community/security/resource-certification/trust-anchor-locator/>
 - Simply add to the TAL folder in the validator cache
- Some examples of AS0 being used today:

RPKI/RTR prefix table

Prefix	Prefix Length	Origin-AS
2.57.180.0	22 - 24	0
5.57.80.0	22 - 22	0
23.4.85.0	24 - 24	0
23.173.176.0	24 - 24	0
23.211.114.0	23 - 24	0
45.12.44.0	22 - 22	0
58.181.75.0	24 - 24	0
109.122.244.0	22 - 22	0

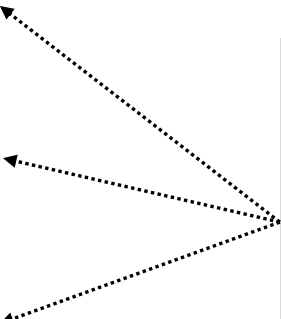


Route Origin Validation – Implementations

- Cisco IOS – available from release 15.2
- Cisco IOS/XR – available from release 4.3.2
- Juniper – available from release 12.2
- Nokia – available from release R12.0R4
- Huawei – available from release V800R009C10
- FRR – available from release 4.0
- BIRD – available from release 1.6
- OpenBGPD – available from OpenBSD release 6.4
- GoBGP – available since 2018
- VyOS – available from release 1.2.0-RC11
- Mikrotik ROS – available from release v7

RPKI Validator Caches

- NLnet Labs Routinator
 - <https://www.nlnetlabs.nl/projects/rpki/routinator/>
 - <https://github.com/NLnetLabs/routinator>
- LACNIC/NIC Mexico validator (FORT)
 - <https://fortproject.net/en/validator>
 - <https://nicmx.github.io/FORT-validator/>
- Cloudflare validator (OctoRPKI)
 - <https://github.com/cloudflare/cfrpki>
 - <https://blog.cloudflare.com/cloudflares-rpki-toolkit/>
- (RIPE NCC validator)
 - Will be discontinued as from 1st July 2021



Available as
Debian/Ubuntu
.deb packages
for easy install

Validator Cache Deployment

- Network Operator design advice:
 - Deploy at least two Validator Caches
 - Geographically diverse
 - Two different implementations
 - For software independence
 - Configure validator to listen on both IPv4 and IPv6
 - Configure routers with both IPv4 and IPv6 validator connections
 - Securing the validator: Only permit routers running EBGP to have access to the validators

Deploying RPKI within an AS

- For fully supported Route Origin Validation across the network:
 - All EBGP speaking routers need talk with a validator
 - Supporting ROV means dropping **invalids** as they arrive in the network
 - EBGP speaking routers are part of the operator IBGP mesh
 - IBGP speaking routers do not need to talk with a validator
 - Only **valid** and **NotFound** prefixes will be distributed from the EBGP speaking routers
 - The validation table is not distributed from router to router
- Even if network is not default free, there is still value in implementing ROV
 - Invalid routes are not in the internal table, relying instead on default to upstream (who hopefully will also run ROV)

RPKI Summary

- All AS operators must consider deploying:
 - Signing ROAs
 - Dropping Invalids (ROV)
- An important step to securing the routing system
- Doesn't secure the path, but that's the next important hurdle to cross
- With origin validation, the opportunities for malicious or accidental mis-origination are considerably reduced
- NIST RPKI Monitor
 - <https://rpki-monitor.antd.nist.gov/>
- NLnetLabs RPKI FAQ:
 - <https://nlnetlabs.nl/projects/rpki/faq/>

NIST RPKI Monitor

Version: 2.0 Last Update : 2021-05-30 06:00 [Feedback](#)

Announcements:

[Welcome to version 2 of the RPKI Monitor](#)

Reports:

[Spike in RPKI-ROV state changes seen on 2021-05-10 12:00](#)

RPKI-ROV Analysis: Global Analysis

Date

05/30/2021



Hour:

06



RIR:

All



Protocol:


IPv4



SUBMIT

RPKI-ROV Analysis: Global Analysis

RPKI-ROV Analysis of Unique Prefix-Origin Pairs (IPv4)



NIST RPKI Monitor: RPKI-ROV Analysis

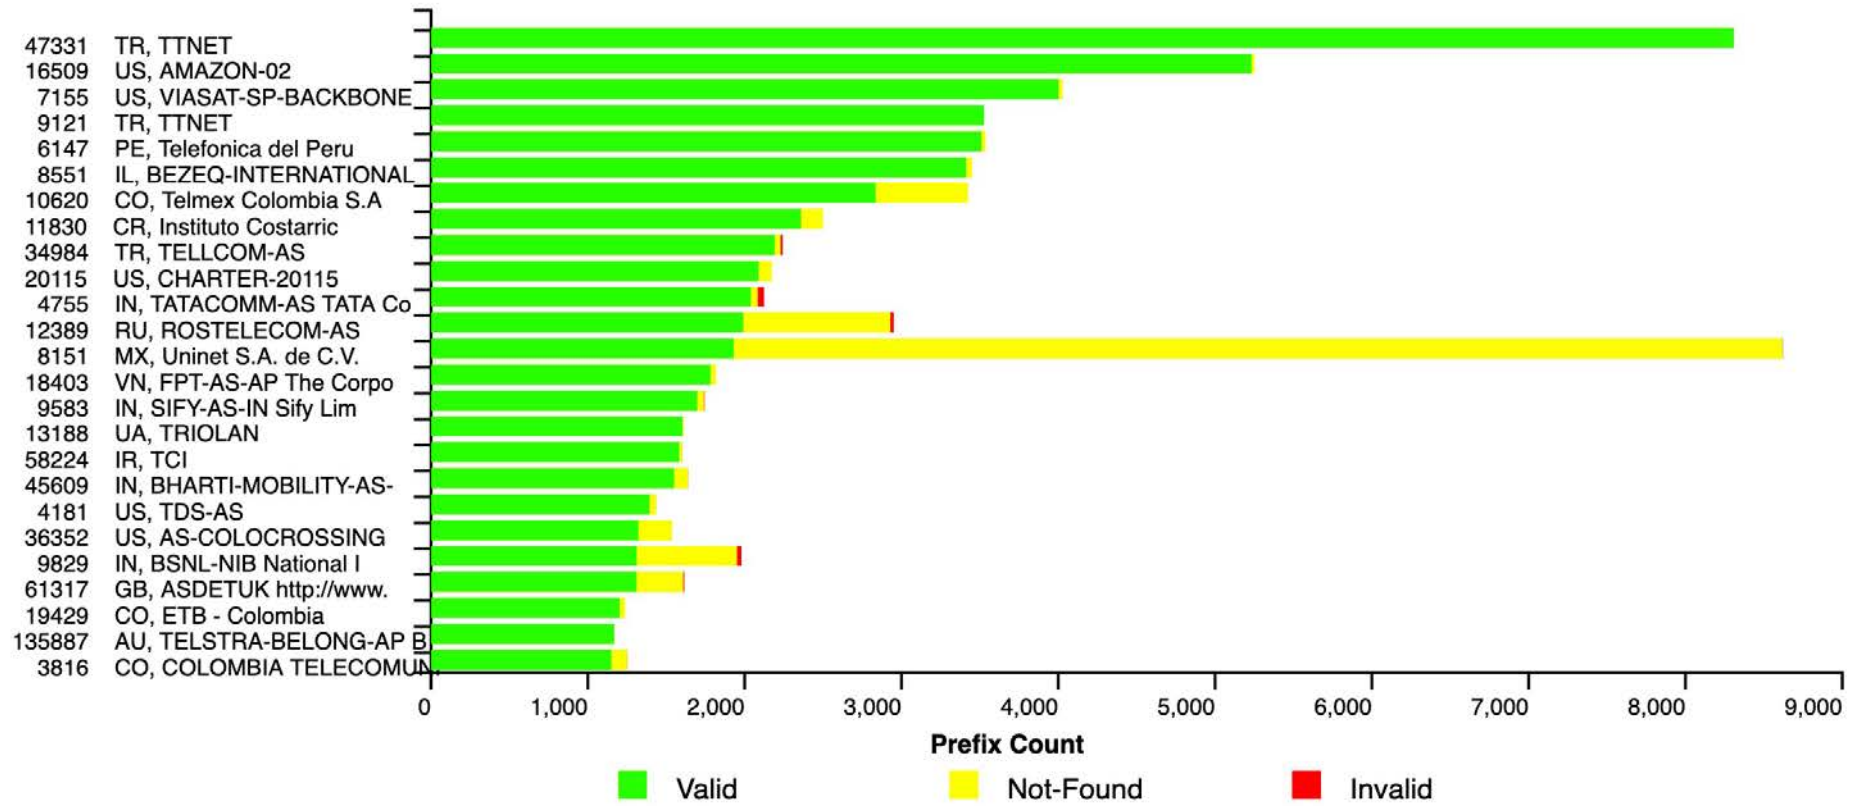
Protocol: IPv4

RIR: All

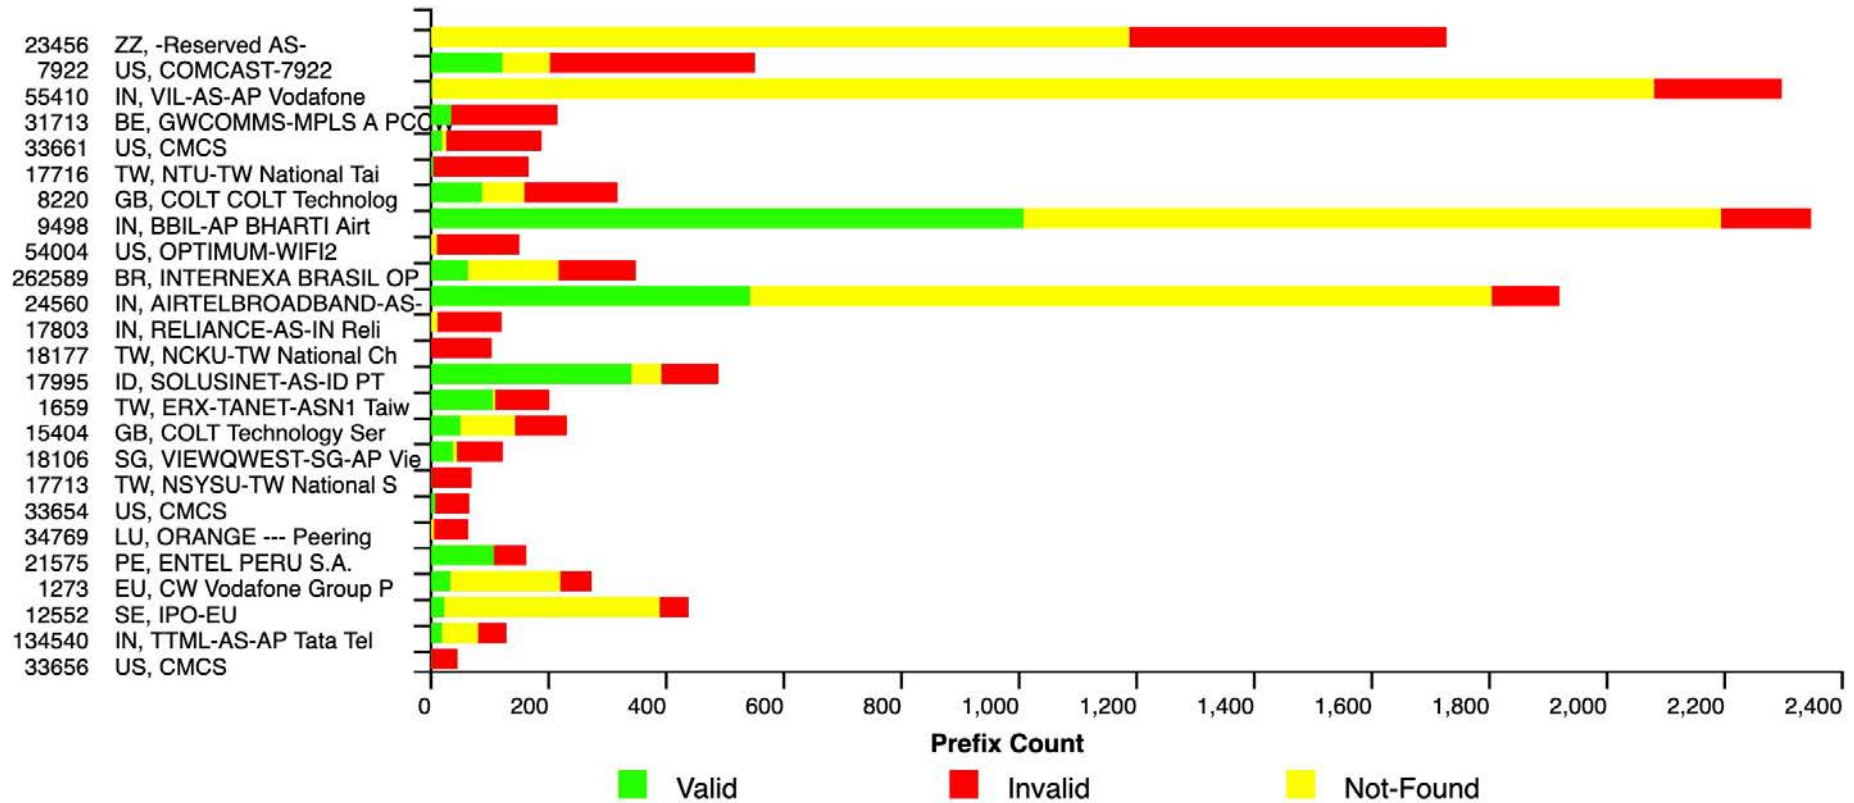
Date: 2021-05-30 06:00

Screenshot


25 Autonomous Systems with the most BGP Originated Prefixes VALID by RPKI-ROV (IPv4)



25 Autonomous Systems with the most BGP Originated Prefixes INVALID by RPKI-ROV (IPv4)



25 Autonomous Systems with the most Address Space (/24s) VALID by RPKI-ROV (IPv4)



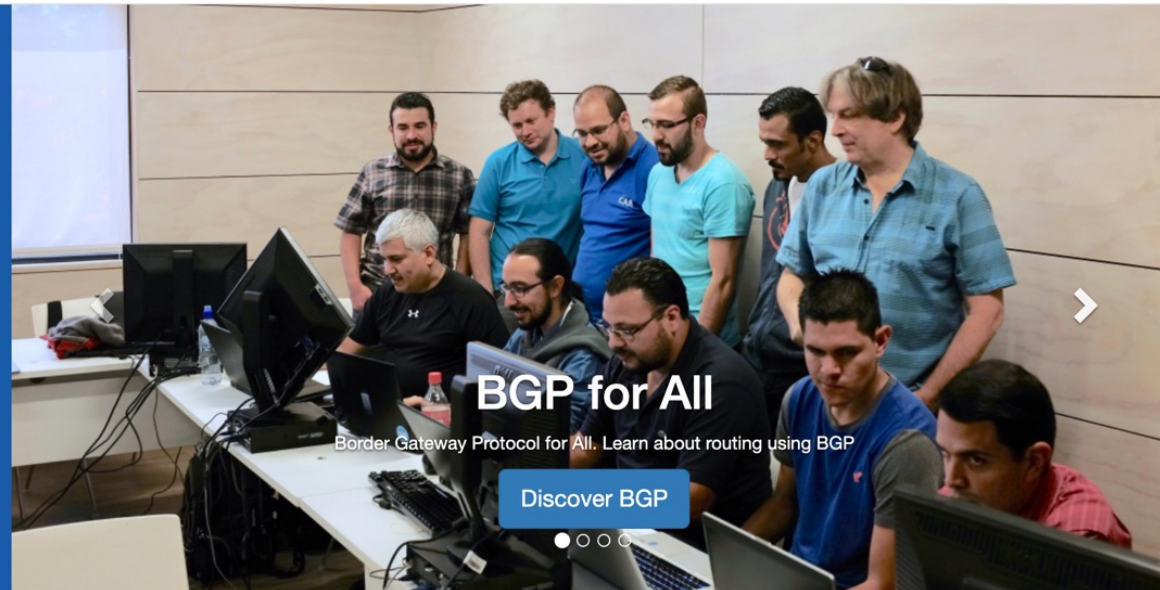
Video-based Education at <https://learn.nsrc.org/>

Welcome

The LEARN web site of the NSRC provides educational content about technical Internet topics with a mix of video clips, accompanying reference documents, command line examples and Exercises to reinforce learning.

BGP for All

Over 100 videos on the Border Gateway Protocol including Introduction to Routing and BGP, BGP attributes, policy, scaling techniques, best practices and BGP For NRENs as well as IXP design and implementation, peering, communities and detailed Multihoming scenarios.



CNDO

(Campus Network Design & Operations)

Learn how to improve campus network designs and how to implement best practices in switched and routed networks

91 Videos

BGP for All

(Border Gateway Protocol)

The primary routing protocol used to transfer data and information on the Internet or autonomous systems

109 Videos

perfSONAR

(performance Service-Oriented Network monitoring ARchitecture)

Networking tools for end-to-end monitoring and troubleshooting of multi-domain network performance

28 Videos

ScienceDMZ

Network design with equipment, configuration, and security policies optimized for high-performance scientific use

15 Videos

FedIdM

(Federated Identity Management)

Agreement between domains to allow users to access applications and services using the same digital identity

10 Videos

Acknowledgements

- National Science Foundation (NSF)
 - For promoting good routing security via the International Research Network Connections programme and the work of the Network Startup Resource Center
- Internet Society
 - Partnership in the global efforts towards improving routing security

Questions?



UNIVERSITY OF OREGON

