RPKI Case Study
in
Tashi InfoComm Limited

PRESENTER: CHOKI WANGDA

Disclaimer

- We have tested on

Sl No.	Model	Software	Version
1.	Cisco ASR1001-X	Cisco IOS-XE	16.03.08
2.	Cisco ASR1002	Cisco IOS-XE	03.16.08.S
3.	Cisco ASR1002-HX	Cisco IOS-XE	16.03.06
4.	mx240	Junos	17.3R3-S2.2
5.	mx104	Juno tos	15.1R6.7

NETWORK DIAGRAM

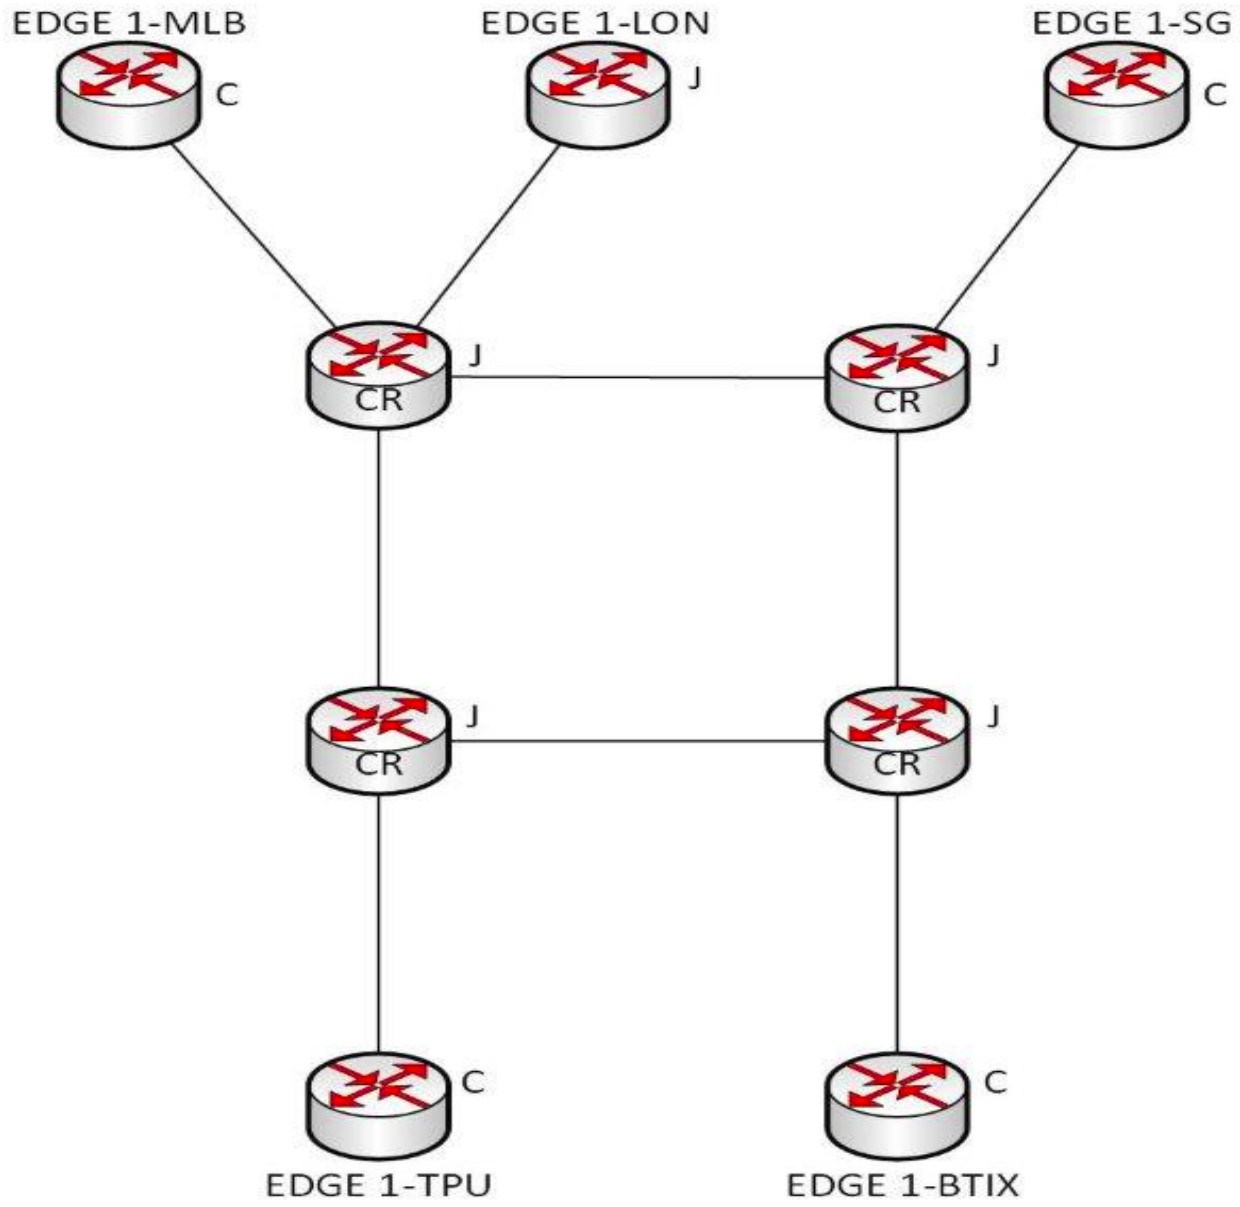

IOS-XE DEFAULT BEHAVIOR

- Created RTR session between RPKI Validator & all the edge routers.
- No rpki-related configuration on Core Routers.
- No “announce rpki-state” configured on edge(cisco) for validation state propagation in ibgp.

RESULT:

- All the prefixes learned on edge(cisco) via ibgp is marked as valid(**inconsistent RPKI states**)
- Valid is preferred over unknown/invalid (overwrites the default bgp path calculation), hence **altering the path.**

IOS-XE DEFAULT BEHAVIOR

OUTPUT:

- **Inconsistent RPKI State**

edge1-btix.as38740#sh ip bgp **43.229.124.0/22**
43.241.139.238 from 43.241.139.238 (103.80.108.5)
Origin IGP, localpref 100, valid, **external, best**
Community: 38740:300
path 44CC9B68 **RPKI State not found**

edge1-tpu.as38740#sh ip bgp **43.229.124.0/22**
43.241.139.1 (metric 20) from 43.241.139.1 (43.241.139.1)
Origin IGP, metric 0, localpref 100, valid, **internal, best**
Community: 38740:300
Originator: 43.241.139.12, Cluster list: 1.1.1.5
path 5E68530C **RPKI State valid**

IOS-XE DEFAULT BEHAVIOR

OUTPUT: Altering a path

- **BEFORE RPKI**

edge1-sg#sh ip bgp **1.0.128.0/19**
38040 23969
43.241.139.39 (metric 20) from 43.241.139.39 (43.241.139.39)
Origin IGP, metric 0, localpref 100, valid, **internal**
Community: 8714:65010 8714:65011 38740:501
Originator: 43.241.139.122, Cluster list: 2.2.2.2, 43.241.139.45

38040 23969
27.111.228.150 from 27.111.228.123 (27.111.228.123)
Origin IGP, localpref 100, valid, **external, best**
Community: 38740:500

IOS-XE DEFAULT BEHAVIOR

OUTPUT: Altering a path

- **AFTER RPKI**

edge1-sg#sh ip bgp **1.0.128.0/19**
38040 23969
43.241.139.39 (metric 20) from 43.241.139.39 (43.241.139.39)
Origin IGP, metric 0, localpref 100, valid, **internal, best**
Community: 8714:65010 8714:65011 38740:501
Originator: 43.241.139.122, Cluster list: 2.2.2.2, 43.241.139.45
path 7F9F9AE8D368 **RPKI State valid**

38040 23969
27.111.228.150 from 27.111.228.123 (27.111.228.123)
Origin IGP, localpref 100, valid, **external**
Community: 38740:500
path 7F9F394DCA88 **RPKI State not found**

With “announce rpki-state” on IOS-XE

Why announce validation state in ibgp ?

- It saves that core routers/access routers from having RTR session with RPKI server.
- To get rid of path alteration.
- Consistent rpki validation state throughout the network.
- I don’t always have to go to edge to check the rpki state

With “announce rpki-state” on IOS-XE

- With “**announce rpki-state**” configured on edge(cisco) under iBGP policy
- Still haven't configured anything on core(Junos)

RESULT:

- Prefixes learned via ibgp the validation state is marked correctly between the cisco edge

  Problem solved???????

  Not yet...because all the prefixes learned via ibgp from edge1-lon(juons) is marked as valid in Cisco edge.

With “announce rpki-state” on IOS-XE

- **but on core(Junos):**

  - unknown iana opaque 0x4300:0x0:0x0

  - unknown iana opaque 0x4300:0x0:0x1

  Junos unable to recognize the ext-comm send by Cisco.

With “announce rpki-state” on IOS-XE

Output:

Validation state is marked correctly on cisco edge

edge1-btix.as38740#sh ip bgp **43.229.124.0/22**
BGP routing table entry for 43.229.124.0/22, version 59632
43.241.139.238 from 43.241.139.238 (103.80.108.5)
Origin IGP, localpref 100, **valid, external, best**
Community: 38740:300
path 44CC9B68 **RPKI State not found**

edge1-tpu.as38740#sh ip bgp **43.229.124.0/22**
43.241.139.1 (metric 20) from 43.241.139.1 (43.241.139.1)
Origin IGP, metric 0, localpref 100, **valid, internal, best**
Community: 38740:300
Extended Community: **0x4300:0:1**
Originator: 43.241.139.12, Cluster list: 1.1.1.5
path 5E68530C **RPKI State not found**

Output on Cores(junos)

choki@bb1.thimphu# run show route **43.229.124.0/22** logical-system cr1-thimphu detail

43.229.124.0/22 (1 entry, 1 announced)
State: <FlashAll>
*BGP Preference: 200/-101
Source: 43.241.139.12
Next hop: 43.241.139.129 via xe-2/0/1.0, selected
Protocol next hop: 43.241.139.12
Local AS: 38740 Peer AS: 38740
Validation State: unverified
AS path: 136039 I
Communities: 38740:300 **unknown iana opaque 0x4300:0x0:0x1**
Localpref: 100
Router ID: 43.241.139.12

Findings

As per JTAC :

- “Basically, Junos is expecting the community **0x43:AS:value,** AS in the second column is the additional check.”
- Since Cisco is sending “0x4300:0:1” and juniper is considering 0 as AS, hence there’s error.

Junos complains when manually configuring the ext-community

```
choki@bb1.thimphu# set community test1 members 0x4300:0:1

{master}[edit policy-options]
choki@bb1.thimphu# commit check
re0:
[edit policy-options community test1 members]
  '0x4300:0:1'
    invalid autonomous system value at '0' not in range 1 to 65535. Use '0L' Long
format to specify 4 byte AS
error: configuration check-out failed
```

but AS PER RFC

RFC8097

The origin validation state extended community is an opaque extended community with the following encoding:

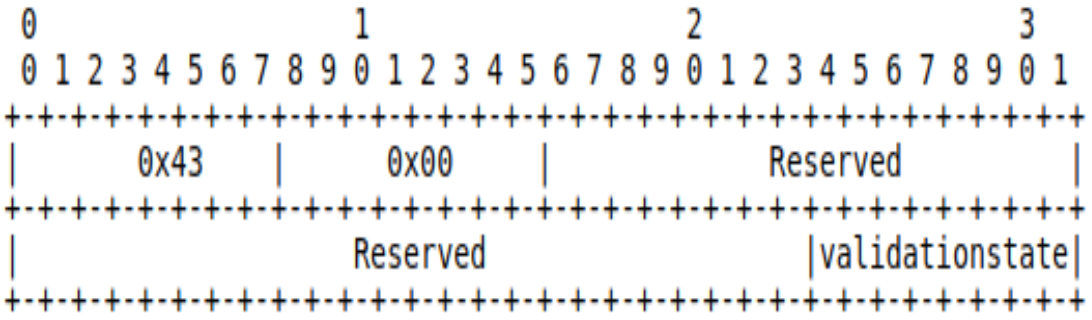

Value	Meaning
0	Lookup result = "valid"
1	Lookup result = "not found"
2	Lookup result = "invalid"

1. **Cisco sends "0x4300:0:X" in ibgp, but as per RFC8097 cisco should be sending "0x43:0:x"**
2. **Juniper is expecting the second field as ASN but as per RFC8097 juniper should consider sub-type(0x00) field as "BGP Origin Validation State" not ASN.**

Refer rfc4360 & rfc7153 for more details.

HIT-AND-TRIAL

METHOD 1:

CONFIGURED COMMUNITY AS SUGGESTED BY JUNIPER ON JUNIPER ROUTER.

community origin-validation-state-invalid members 0x43:38740:2;
community origin-validation-state-unknown members 0x43:38740:1;
community origin-validation-state-valid members 0x43:38740:0;

BUT WE COULDN'T CONFIGURE SAME EXT-COMMUNITY ON CISCO ASR ROUTERS.

Result:

- Validation state is correct on Junos
- But error in community output “unknown iana 43”
- All the prefixes learned via ibgp from edge1-lon(juons) is marked as valid in Cisco edge (**inconsistent rpki states**)

HIT-AND-TRIAL

METHOD 1 OUTPUT:

choki@edge1-lon# run show route **1.1.8.0/24** detail

*BGP **Validation State: unknown**

AS path: 1299 4134 I

Communities: 1299:20000 38740:1299 **unknown iana 43**

choki@bb1.malbase# run show route 1.1.8.0/24 logical-system cr1-malbase detail

*BGP Protocol next hop: 43.241.139.122

Validation State: unknown

AS path: 1299 4134 I

Communities: 1299:20000 38740:1299 **unknown iana 43**

Router ID: 43.241.139.122

HIT-AND-TRIAL

METHOD 1 OUTPUT:

edge1-sg#sh ip bgp **1.1.8.0/24**
1299 4134
43.241.139.39 (metric 20) from 43.241.139.39 (43.241.139.39)
Origin IGP, metric 0, localpref 100, valid, **internal, best**
Community: 1299:20000 38740:1299
Extended Community: 0x43:38740:1
Originator: 43.241.139.122, Cluster list: 2.2.2.2, 43.241.139.45
path 7F9F9510F5B8 **RPKI State valid**

2914 4134
116.51.31.77 from 116.51.31.77 (129.250.0.184)
Origin IGP, metric 12235, localpref 100, valid, **external**
Community: 38740:2914
path 7F9FBE261FA8 **RPKI State not found**

HIT-AND-TRIAL

METHOD 2: CUSTOM COMMUNITY

- We tried with our own community

  38740:0 – Valid

  38740:1 – Unknown

  38740:2 - Invalid

RESULT:

- All the prefixes learning via iBGP is marked as valid when it reaches cisco edge (**inconsistent RPKI states**)

HIT-AND-TRIAL

METHOD 2: OUTPUT:

edge1-mal#sh ip bgp **1.1.20.0/24**
180.87.38.156 from 180.87.38.156 (66.110.10.202)
Origin IGP, localpref 100, valid, **external, best**
Community: 38740:6453
path 7EFFB06C1188 **RPKI State not found**

edge1-btix.as38740#sh ip bgp **1.1.20.0/24**
43.241.139.2 (metric 20) from 43.241.139.2 (43.241.139.2)
Origin IGP, metric 0, localpref 100, valid, **internal, best**
Community: 38740:1 38740:6453
Originator: 43.241.139.44, Cluster list: 1.1.1.5, 43.241.139.45
path 51F58CEC **RPKI State valid**

JUNIPER NEW RELEASE

- **Juniper release to fix this issues**

  1. junos:17.3R3-S5 **14-JUN-2019**
  2. junos:17.4R3 **13-SEP-2019**
  3. junos:18.2R3 **19-JUN-2019**
  4. junos:18.3R3 **no date has been released yet**
  5. junos:18.4R2 **07-JUN-2019**

Maybe we can manually configure “0x4300:0:x” on JUNIPER to recognise what CISCO is sending and make it work!!!!!!

CONCLUSION

- Validation state is not consistent in cisco edge when you don’t configure “announce rpki-state” in ibgp.
- Juniper router doesn't recognize “announce rpki-state” sent by cisco.
- Custom Community doesn't work if you want to propagate rpki states in ibgp
- Can’t configure ext-community in juniper as sent by cisco.

[edit policy-options community test1 members]
'0x4300:0:1'
invalid autonomous system value at '0' not in range 1 to 65535. Use '0L' Long format to specify
4 byte AS
error: configuration check-out failed

Both cisco & juniper doesn't follow rfc 8097

Thank you