Observations during Testing of Route Origin Validation

IETF 103 Bangkok

O. Borchert, K. Lee, D. Montgomery, K. Sriram

National Institute of Standards and Technology

Testing Router Implementation of RPKI Origin Validation

- What is the impact on convergence time when the router performs RPKI ROV?
- Can routers be configured using multiple RPKI validation caches?
  - Stability of Connections.
  - Difference in received data sets.
- Does the router perform RPKI Route Origin Validation (ROV) as specified in RFC 6811?
- Do the routers support the signaling of ROV results using the extended community as specified in RFC 8097?
- Observations on iBGP traffic using RFC 8097 signaling.
  - Impact on iBGP traffic?
  - Prefix Packing?

Impact on BGP Convergence time

- Baseline convergence time
  - No RPKI ROV
  - No other route policies
- Independent of processor, memory, and other variable factors we encountered the following impact on route convergence:
  - Routing table (~700K), 100% ROA coverage, 1 peer
    - An average increase of 2% to 7% in convergence time
  - Routing table (~700K), 100% ROA coverage, 2 peers
    - An average increase of 2% to 7% in convergence time

Routers and Validation Caches

- Routers communicate using the RPKI to Router Protocol (RFC 6810 or RFC 8210) to receive the Validated ROA Payload (VRP)
- All routers tested were able to configure more than one validation cache
- All tested routers use the union of all VRPs from all validation caches
- Some routers respond almost immediately to loss of validation cache
    - Does increase churn if validation state changes.
- Some routers do not react immediately to the loss of a validation cache.
    - Prevents unnecessary churn for “flaky connections” but introduces temporary staleness
    - Here router configuration settings influence the detection of loss of connectivity

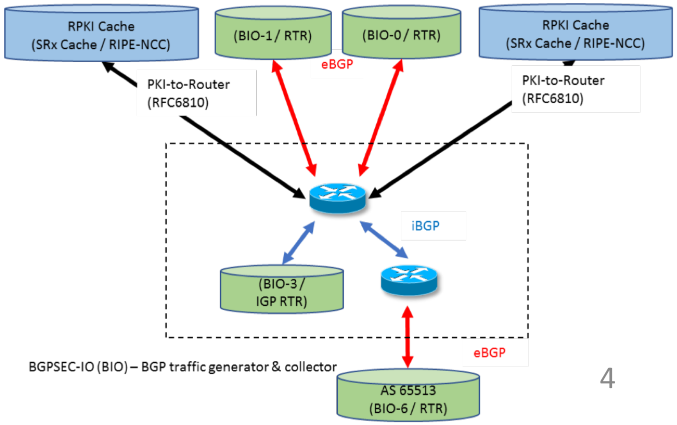

Signaling of validation state information using RFC 8097

- During our tests we noted that not all routers fully supported RFC 8097
  - During the course of our testing we received builds which fixed this issue*
- Some implementations labeled iBGP routes as “valid” if no validation value was signaled using RFC 8097
- Some implementations allowed custom configuration on how to process and label routes with the validation value signaled using RFC 8097

* Contact your vendor in case you observe any issues

RPKI ROV on eBGP learned routes

- Routes with AS_SET*:
  - Some routers had “issues” validating routes containing AS_SET and validated routes with AS_SET as NotFound regardless of ROA coverage.
    - It seems not all vendors chose to validate routes containing AS_SET and followed the suggestion of RFC 6811 to label non-validated routes as “NotFound”; otherwise they should be “Invalid” if covered by a ROA

- Routes without AS_SET:
  - All routers performed RPKI ROV validation correctly.
  - Changes learned through router to cache protocol were applied as soon as synchronization with the validation cache finished.

* we received some “fixes” but consult with your vendor if you experience this on your routers.

RPKI ROV on IGP / local scripted routes

- Not all routers perform origin validation on these routes
- In general these routes are labeled as “valid”
- In other cases these routes were labeled as ”unverified”
- Some routers can be configured how to evaluate these routes.

ROV on routes received via iBGP

- Some routers* allow to perform local validation on routes received via iBGP
- If extended community RFC 8097 was provided*:
  - Some routers use the value of the community string as validation value without performing local validation
  - Some routers allow configuration on how to use them. (apply the value, perform local validation, some hybrid of both)

- Without extended community*:
  - Some routers label these routes as “valid”
  - Some routers label these routes as “unverified”
  - Some routers allow configuration on how to label them.

consult with the router’s manual or vendor for capability questions

Impact on iBGP traffic using ROV signaling (RFC 8097)

- Routes contain the validation state encoded in the update
  - This does reduce the “packability” of prefixes due to different community value.
    - Increase of iBGP traffic due to partial loss of prefix packing (max. x4 – “valid”, “invalid”, “not-found”, no community string)

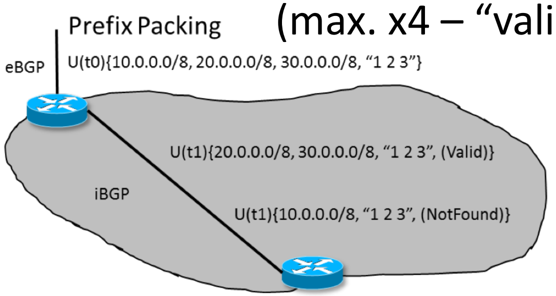
Prefix Packing
eBGP
U(t0){10.0.0.0/8, 20.0.0.0/8, 30.0.0.0/8, “1 2 3”}
U(t1){20.0.0.0/8, 30.0.0.0/8, “1 2 3”, (Valid)}
iBGP
U(t1){10.0.0.0/8, “1 2 3”, (NotFound)}

Example: One update becomes two updates

  - Even if route selection does not change, each change of validation state triggers the route being re-sent.
    - Increase of internal traffic – can become problematic with instable connections to validating cache

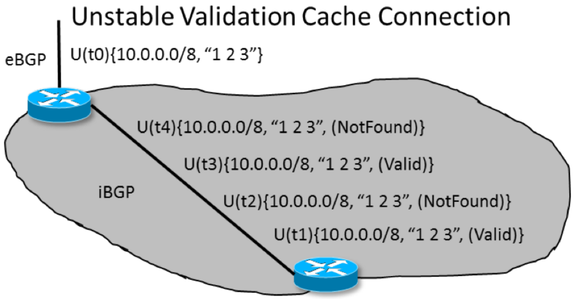
Unstable Validation Cache Connection
eBGP
U(t0){10.0.0.0/8, “1 2 3”}
U(t4){10.0.0.0/8, “1 2 3”, (NotFound)}
U(t3){10.0.0.0/8, “1 2 3”, (Valid)}
iBGP
U(t2){10.0.0.0/8, “1 2 3”, (NotFound)}
U(t1){10.0.0.0/8, “1 2 3”, (Valid)}

Example: Repetitive announcement of same update due to validation state changes in community string

“Conflicting” validation results in Route Origin Validation

- Using iBGP and eBGP with multiple iBGP peers we encountered “conflicting” validation results:
  - Each iBGP peer performs its own local origin validation
  - Each iBGP peer uses RFC 8097 to signal the local validation result
  - The router is configured to perform ROV on eBGP routes only and uses the signaled state for iBGP learned routes.

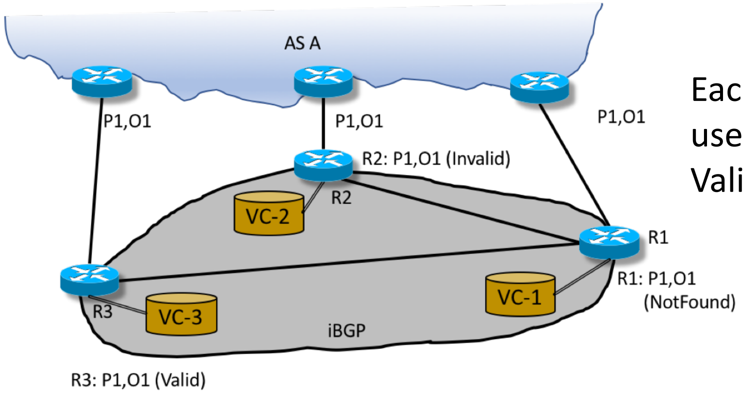

Each router R1, R2, R3 uses an independent Validation Cache (VC)

- It is possible to have different validation results for the same prefix origin in the same router at the same time
  - If origin validation is **not performed locally!**
  - If remotely received ROV results are applied!

Example1: Distributed Validation in iBGP with “conflicting” results

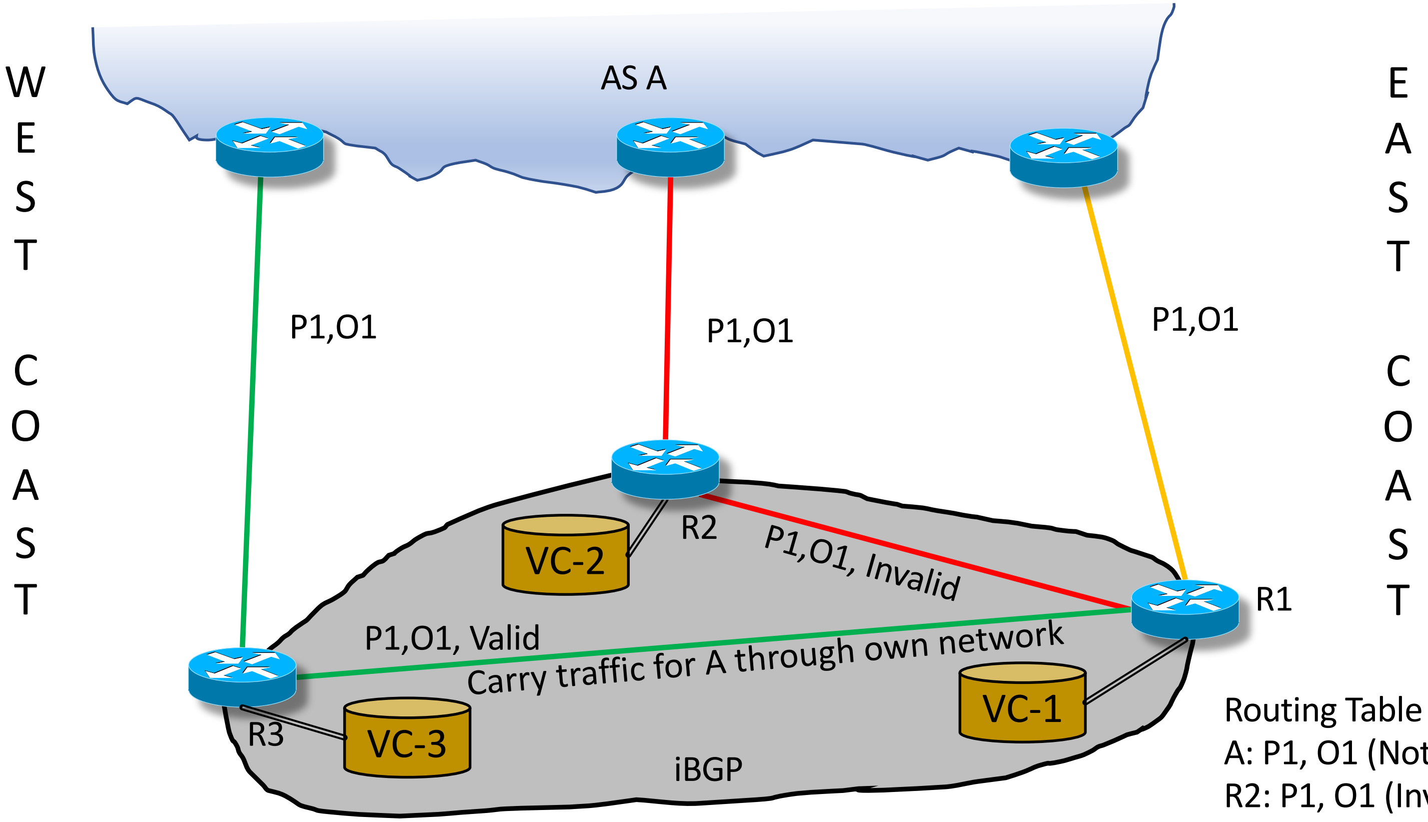

Routing Table in R1
A: P1, O1 (Not-Found)
R2: P1, O1 (Invalid)
R3: P1, O1 (Valid)

Router R1 chooses the VALID route through R3 and transports the traffic through its own network rather than directly to **A.** This could be prevented if R1 performs local validation on iBGP routes or all validation caches would be synchronized.

Example2: Default "valid" for iBGP routes without ROV state (No RFC 8097)

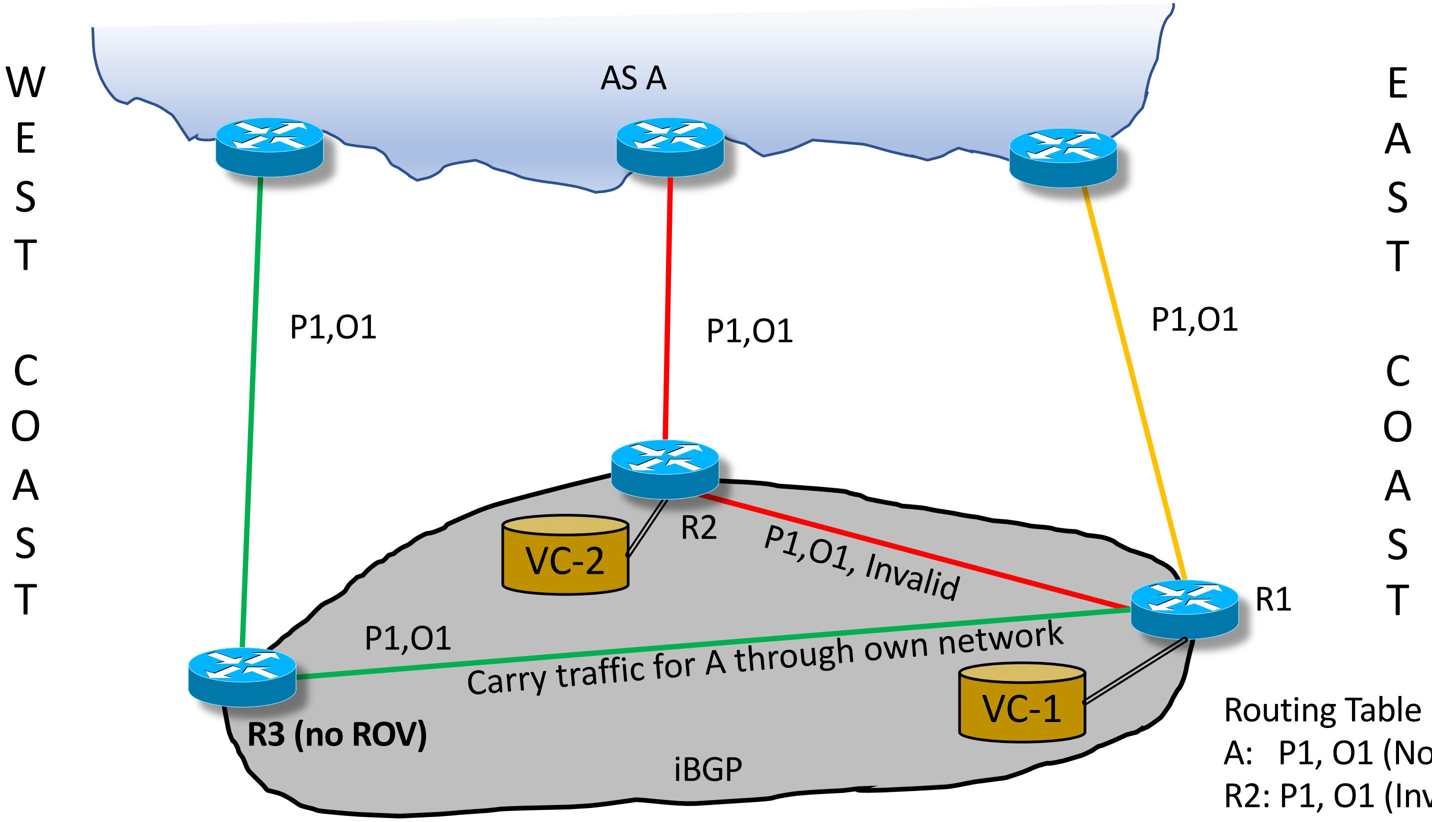

Routing Table in R1
A: P1, O1 (Not-Found)
R2: P1, O1 (Invalid)
R3: P1, O1 (Valid - Default)

Router R1 chooses the VALID route through R3 and transports the traffic through its own network rather than directly to **A.** If the state Unverified would exist with the rule **prefer Not-Found to Unverified**, R1 would route directly to A.

Summary

- In general, router documentation could be better!
  - Sometimes vague / incomplete – questions remain!
  - We observed some configuration labels from within the console but did not find any documentation on them!
  - Not enough examples!
- VRP tables from multiple caches within a router are combined as a union.
- Detection of connection loss to validation caches differs between router implementations.
  - To reduce general churn be more tolerant to connection loss
    - Increase the timeout time before declaring a connection as dead.
    - Don't clear all tables immediately if connection is lost (Configuration ??)
- Signaling validation state in iBGP - RFC 8097
  - Reduction of packing due to signaling validation state – x4
    - Max Factor 4: 3 Validation states + 1 no signaling for selected routes
  - Increased churn if connection to Validation Cache is flaky due to changes in validation
  - Conflicts of route prefix validation states
    - Might cause carrying other carriers traffic though own network
- Verify with your vendor if you experience / have questions regarding any of the observations discussed here!

Questions?

- Community interest in an informational RFC?
- Other Community Input?

oliver.borchert@nist.gov