Resource Public Key Infrastructure for Secure Border Gateway Protocol

George Chang, Majid Arianezhad, and Ljiljana Trajković
gkchang@sfu.ca, arianezhad@live.com, ljilja@sfu.ca

Communication Networks Laboratory
http://www.ensc.sfu.ca/~ljilja/cnl/
Simon Fraser University, Vancouver
British Columbia, Canada

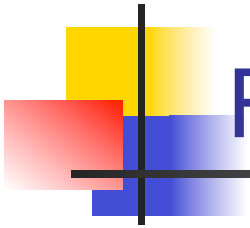

Roadmap

- Introduction
- Securing the Internet
- Testbed: configuration of a router between SFU and BCNET
- Simulation scenario and results
- Conclusion and references

Border Gateway Protocol (BGP)

- Security issues:
  - message insertion, message deletion, and modification to the routes or packets
- Man-in-the-middle attack
- Denial of Service (DoS)
- Distributed Denial of Service (DDoS)
- BGP lacks protection and verification mechanisms for invalid route advertisements

2008 YouTube Incident

- Cause:
  - Pakistan Telecom (AS 17557) re-routed most of YouTube's traffic to itself due to unauthorized advertisement of a more specific route
- Consequence:
  - YouTube network was brought down globally for more than two hours on Feb. 24th 2008

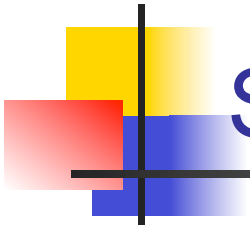

Securing BGP

- Resource Public Key Infrastructure (RPKI):
  - utilizes the Public Key Infrastructure (PKI) to secure resources (routes) for advertisements
  - uses public and private keys to encrypt the certificate that proves route validity
  - implements guards against unauthorized advertisement of routes and resources to neighbouring peers
  - ensures accurate inter-Autonomous System (AS) route advertisement

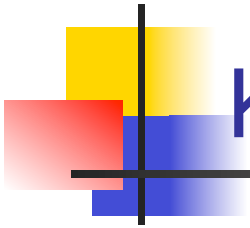

Keys and Certificates

- RPKI uses the well developed public key cryptographic technology
- The public and private keys are generated from the Regional Internet Registry (RIRs) for individual resource holders
- RPKI uses X.509 v3 standard and format specification that is adopted for PKI

RPKI Participants

- Certificate Authorities (CA)
- Authentication built in a hierarchical system:
  - IANA → RIR → ISP → Customers
  - IANA: Internet Assigned Numbers Authority
  - RIR: Regional Internet Registry
  - ISP: Internet Service Provider

RPKI Hierarchy Structure

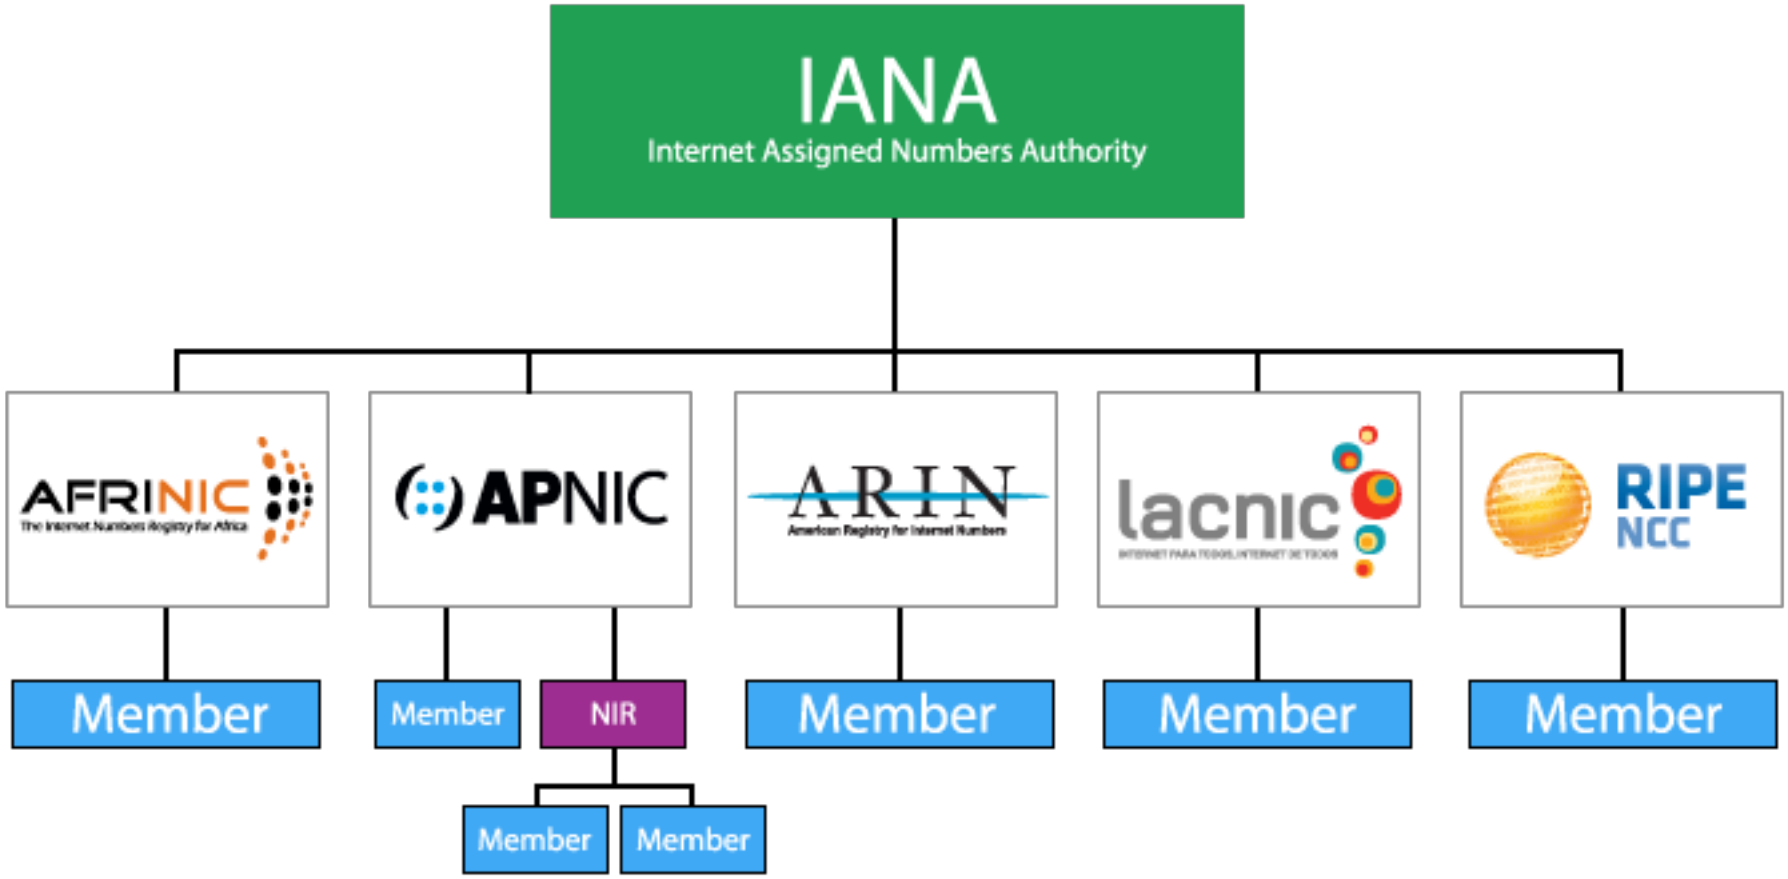

https://www.ripe.net/participate/internet-governance/internet-technical-community/the-rir-system

RPKI Tools

- RIPE and ARIN provide validation tools to the RPKI data repository:
  - web interface
  - cache validator
  - verified routes data
  - automatic queuing of validated ROAs or resources

RIPE: Réseaux IP Européens

ARIN: American Registry for Internet Numbers

ROA: Route Origin Authorization

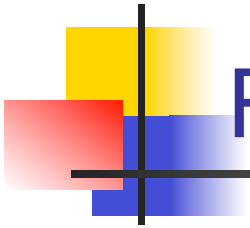

Routing Rules

- Routing decisions are made by the network administrator based on RPKI validity states
- Each route is assigned one of the three validity states:
  - valid: authorized announcement
  - invalid: unauthorized announcement
  - not found: not assigned or not backed by ROA

Testbed Architecture

- Two routers were connected via secure tunneling between two ASes:
  - BCNET (AS 271)
  - SFU (AS 11105)
- Both routers/ASes were connected to the RPKI cache validator obtained from RIPE
- Default RIR was selected as a trust anchor to validate BGP announcements (ARIN)

Testbed Specifications

- Two logical routers were instantiated between SFU and BCNET using Juniper JunOS
- Ubuntu virtual machine was used as the local cache validator hosted on a PC
  - UNIX based system running Oracle JDK 7, rsync, and RIPE's validator package
  - 1 GB of memory allocated
- SFU and BCNET obtained IP resources from ARIN used for route validation

Testbed Topology

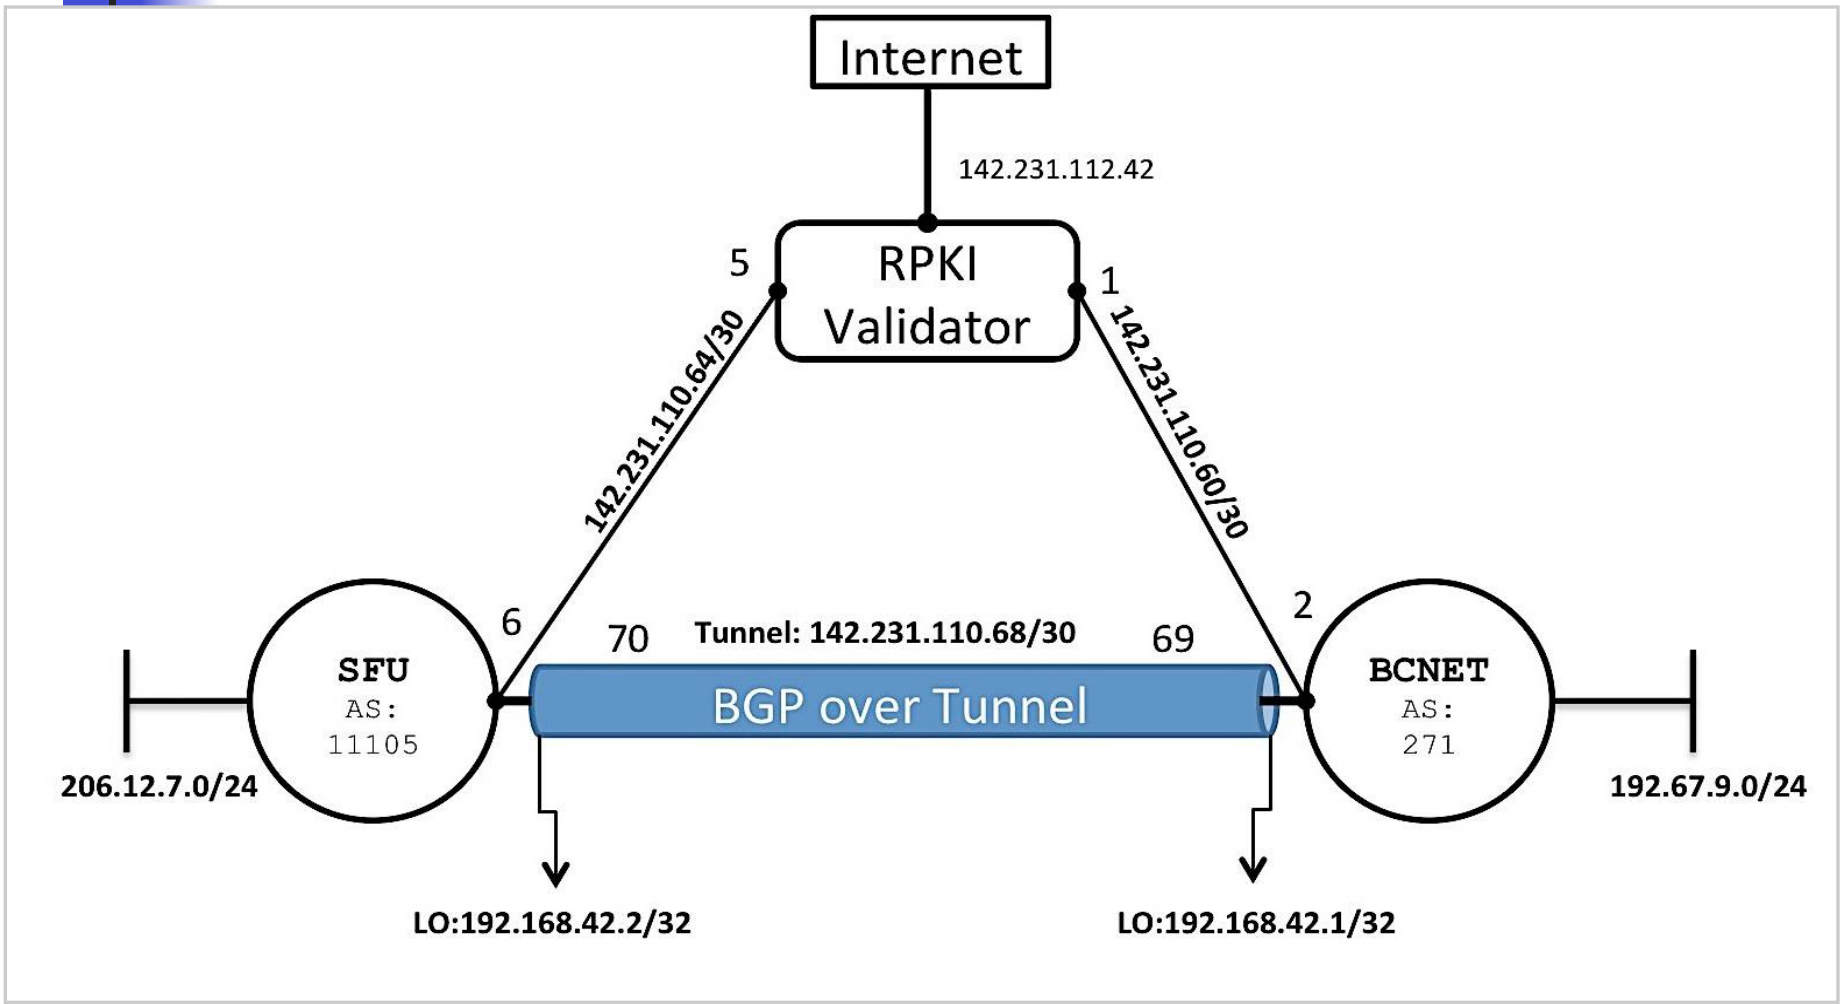

Decision Making via Route Validation

- Verification of the applied routing policy:
  - valid, invalid, and not found statements were set to 110, 90, and 100, respectively
  - decisions are made based on these values chosen by the administrator during router setup
- A rouge test router was introduced to deliberately advertise false information
  - advertising false route to BCNET, if accepted, would reroute traffic from SFU

Results: Valid States

arianezhad@tr1.vncv1> show route protocol bgp validation-state valid

inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

206.12.7.0/24 *[BGP/170] 3w6d 05:23:33, localpref 110
AS path: 11105 I, validation-state: valid
> to 142.231.110.70 via lt-0/2/10.69

Results: Invalid States

arianezhad@tr1.vncv1> show route protocol bgp validation-state invalid

inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

206.12.7.0/24 [BGP/170] 3d 08:00:09, localpref 90

AS path: 4476 I, validation-state: invalid

> to 142.231.110.66 via lt-0/3/10.65

Route 206.12.7.0 Validity

arianezhad@tr1.vncv1> show route 206.12.7.0

inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

206.12.7.0/24 *[BGP/170] 3w6d 05:27:15, localpref 110
AS path: 11105 I, validation-state: valid
> to 142.231.110.70 via lt-0/2/10.69
[BGP/170] 3d 08:03:15, localpref 90
AS path: 4476 I, validation-state: invalid
> to 142.231.110.66 via lt-0/3/10.65

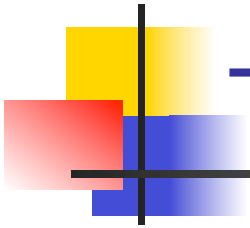

Testbed Summary

- We implemented the testbed using physical routers and the RPKI local cache server
- Validation states were received for the advertised routes
- A falsified route was injected and verified that the route is identified as invalid by the validator

Simulation Goals

- Implement the RIPE RPKI Validator as a network administrator:
  - use the TAL received from the local RIR to fetch route data
  - verify that the validator is reliably stable over long periods and remains online during simulation
- Fetch validated production routes from the validator implemented in the simulator

TAL: Trust Anchor Locator

RAR: Regional Internet Registry

Simulation: Network Configuration

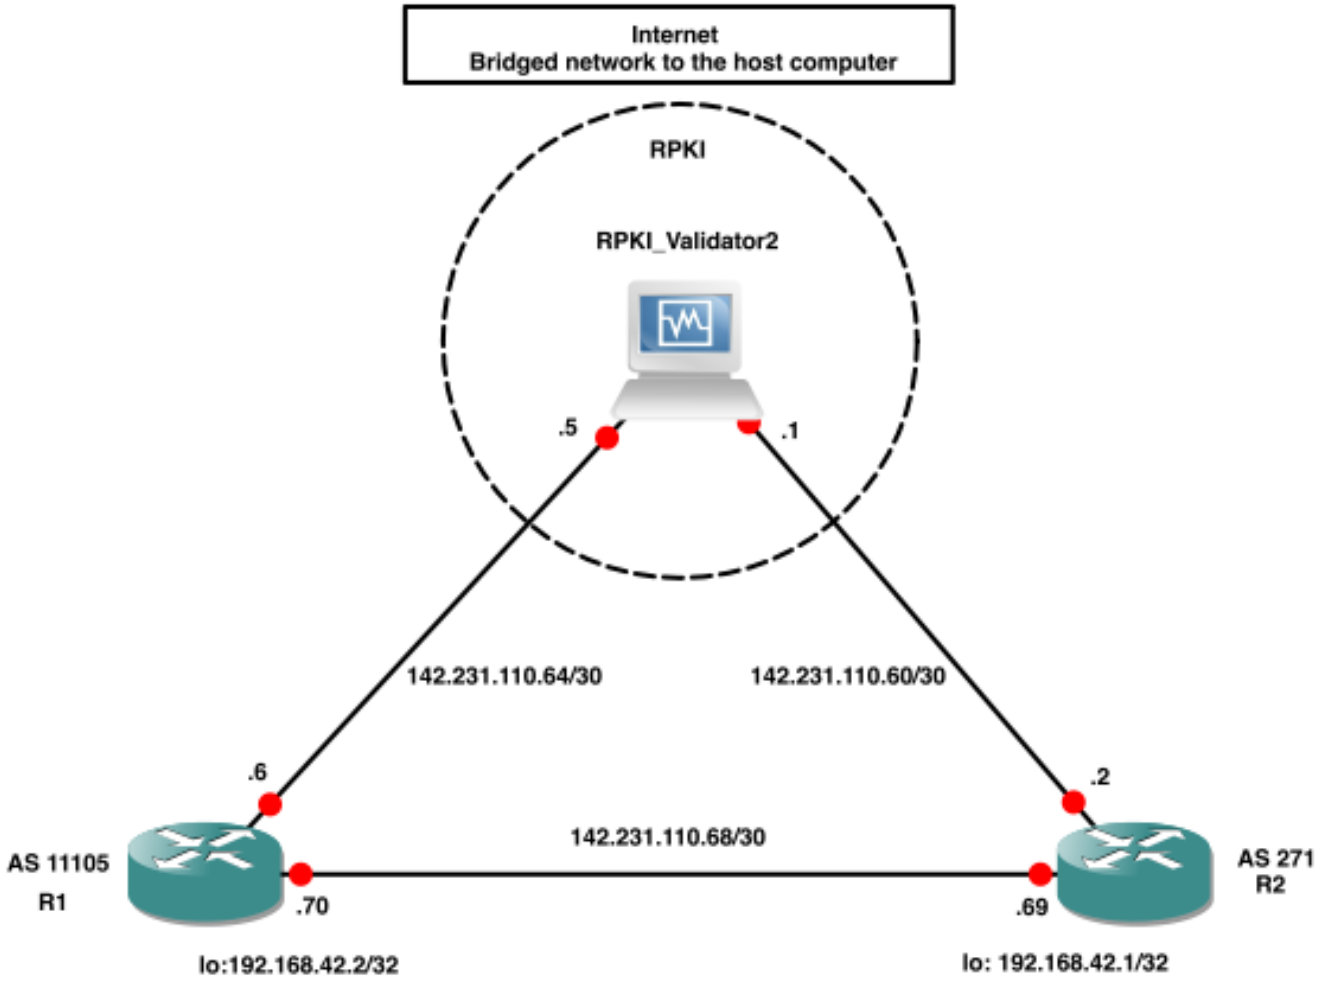

Trust Anchor Locator: TAL

- All the verified routes from ARIN for North America were downloaded by adding the TAL file for ARIN
- ARIN routes as of Aug. 17, 2015:
  - 950 valid routes
  - 1 not found route
  - 0 invalid routes
- In total, 17,432 verified routes were downloaded to the RPKI validator

RPKI Validator Web UI: Trust Anchors Page

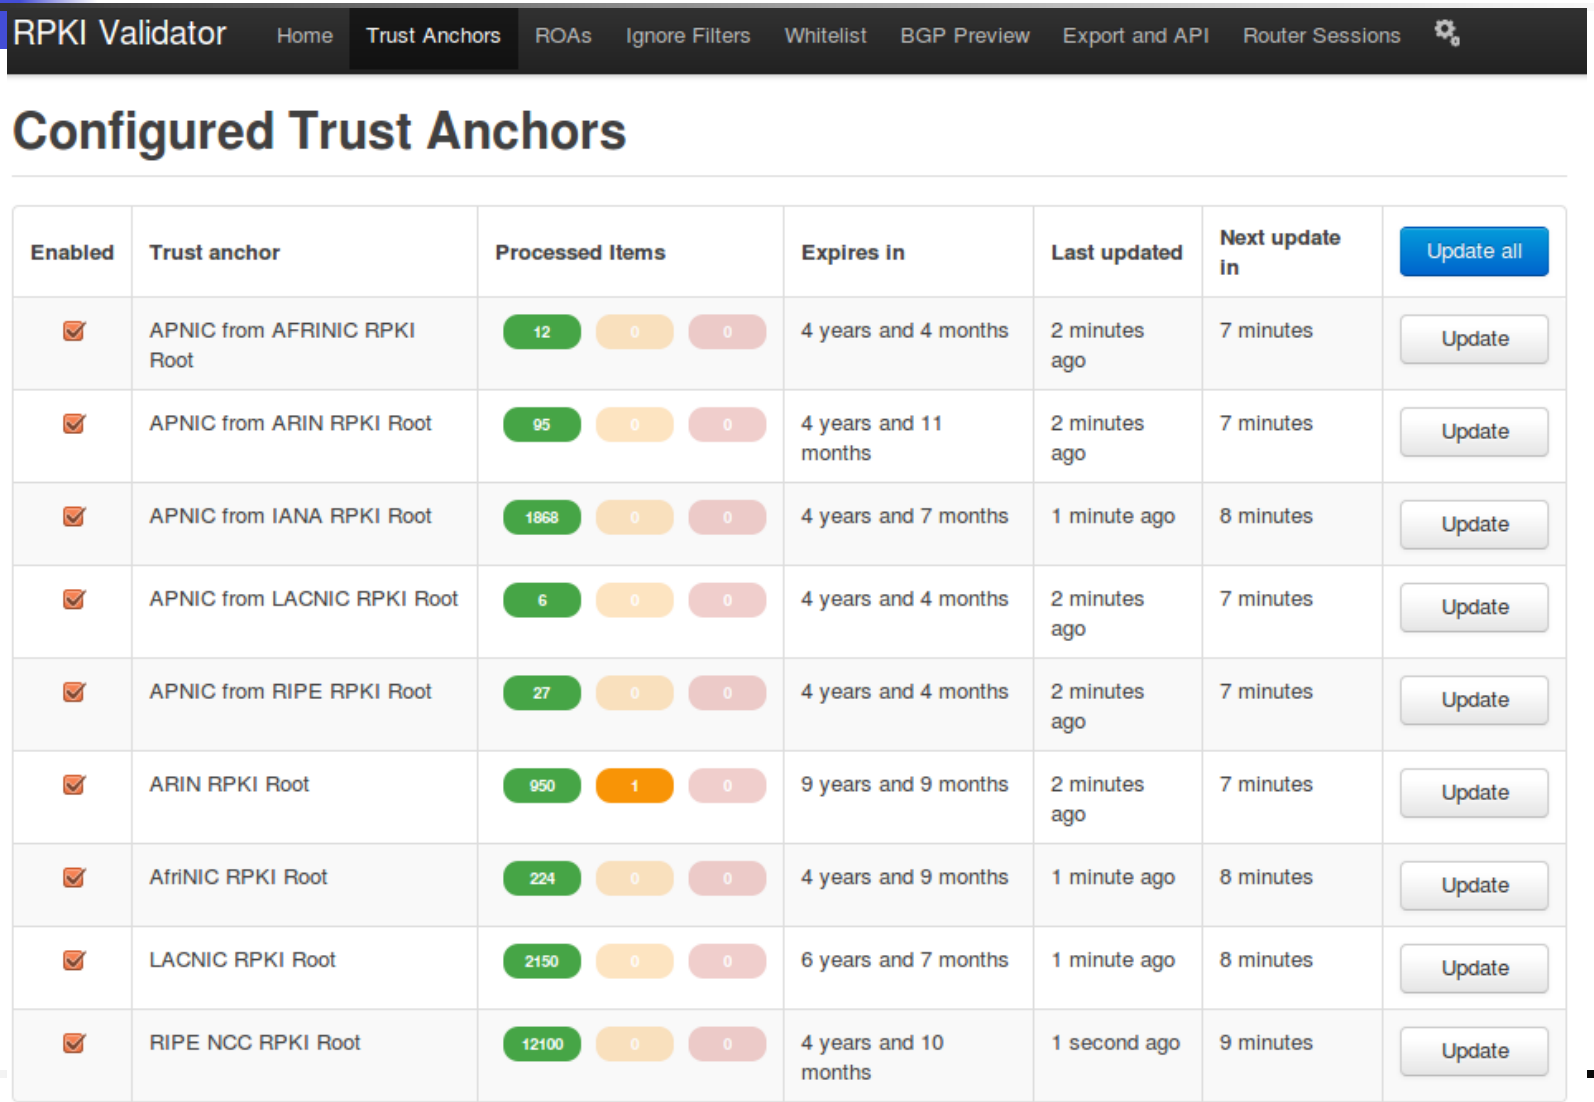

Validated Production Routes Downloaded to the Router

Network	Maxlen	Origin-AS	Source	Neighbor
2.0.0.0/16	16	3215	0	142.231.110.65/8282
2.0.0.0/12	16	3215	0	142.231.110.65/8282
2.1.0.0/16	16	3215	0	142.231.110.65/8282
2.2.0.0/16	16	3215	0	142.231.110.65/8282
2.3.0.0/16	16	3215	0	142.231.110.65/8282
2.4.0.0/16	16	3215	0	142.231.110.65/8282
2.5.0.0/16	16	3215	0	142.231.110.65/8282
2.6.0.0/16	16	3215	0	142.231.110.65/8282
2.8.0.0/16	16	3215	0	142.231.110.65/8282
2.9.0.0/16	16	3215	0	142.231.110.65/8282
2.10.0.0/16	16	3215	0	142.231.110.65/8282
2.11.0.0/16	16	3215	0	142.231.110.65/8282
2.12.0.0/16	16	3215	0	142.231.110.65/8282
2.13.0.0/16	16	3215	0	142.231.110.65/8282
2.14.0.0/16	16	3215	0	142.231.110.65/8282
2.80.0.0/14	14	3243	0	142.231.110.65/8282
2.148.0.0/14	14	2119	0	142.231.110.65/8282
2.248.0.0/13	13	3301	0	142.231.110.65/8282
5.1.0.0/19	19	21219	0	142.231.110.65/8282
--More--				

Advertisement Results

- Using the rpki-loc-pref, each individual state was set and a preference number was assigned to each advertised route:
  - route-map rpki-loc-pref permit 10
    match rpki invalid
    set local-preference 90
    !
    route-map rpki-loc-pref permit 20
    match rpki not-found
    set local-preference 100
    !
    route-map rpki-loc-pref permit 30
    match rpki valid
    set local-preference 110

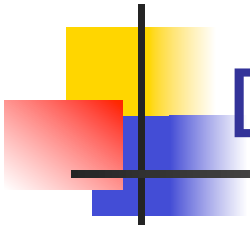

Decision Making

- Network administrators may:
  - use the local-preferences value to help make routing decisions
  - accept routes that are unknown or not found
  - design rules to handle the validity information via assigned local preferences

Advertisement Results: valid

- Route 206.12.7.0 was advertised to router R2 (AS 271)
- This original route was advertised by R1 (AS 11105)
- Router R2 identified that the route was valid and a localpref of 110 was set:

```
R2#show ip bgp 206.12.7.0
BGP routing table entry for 206.12.7.0/24, version 3
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 2
  11105
    142.231.110.70 from 142.231.110.70 (142.231.110.70)
      Origin IGP, metric 0, localpref 110, valid, external, best
      path 68DB44CC RPKI State valid
      rx pathid: 0, tx pathid: 0x0
R2#
```

Advertisement Results: invalid

- An invalid route was advertised to R1 (AS 11105) from R2 (AS 271)
- Router R1 identified that the route was invalid and a localpref of 90 was set:

```
R1#sh ip bgp 193.175.146.0
BGP routing table entry for 193.175.146.0/24, version 0
Paths: (1 available, no best path)
  Not advertised to any peer
  Refresh Epoch 9
  271
    142.231.110.69 from 142.231.110.69 (142.231.110.69)
      Origin IGP, metric 0, localpref 90, valid, external
      path 682CAF34 RPKI State invalid
      rx pathid: 0, tx pathid: 0
```

Advertisement Results: not found

- A not found route was advertised to R2 (AS 271) from R1 (AS 11105)
- Router R2 identified that the route was not found and a localpref of 100 was set:

```
R2#sh ip bgp 6.0.0.0
BGP routing table entry for 6.0.0.0/8, version 5
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 10
  11105
    142.231.110.70 from 142.231.110.70 (142.231.110.70)
      Origin IGP, metric 0, localpref 100, valid, external, best
      path 68DB4424 RPKI State not found
      rx pathid: 0, tx pathid: 0x0
R2#
```

Simulation Summary

- Two stand-alone virtual production routers were connected to a Virtualbox Ubuntu “router” running the RPKI Validator tool
- The validator was connected to the Internet to download the latest route information from RIRs
- The route validity states were downloaded to the router and verified with the advertised route
- Routing decisions may be made based on the state and its localpref value

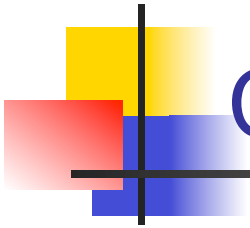

Conclusion

- RPKI is becoming a widely accepted technology
- It calls for additional participants to validate their routes
- The validation tool is user friendly:
  - easy to implement
  - easily maintained
  - limited resources are required to monitor the system, which automatically updates local data
- The experimental results indicate that RPKI may provide protection against route origin hijacks

References

- Y. Rekhter and T. Li, “A Border Gateway Protocol 4 (BGP-4),” *IETF RFC 1771*, Mar. 1995.
- S. Murphy, “BGP Security Vulnerabilities Analysis,” *IETF RFC 4272*, Jan. 2006.
- Pakistan hijacks [Online]. Available: YouTube http://www.renesys.com/2008/02/pakistan-hijacks-youtube-1/.
- A. Heffernan, “Protection of BGP Sessions via the TCP MD5 Signature Option,” *IETF RFC 2385*, Aug. 1998.
- M. Lepinski and S. Kent, “An Infrastructure to Support Secure Internet Routing,” IETF RFC 6480, Feb. 2012.
- G. Huston and G. Michaelson, “Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs),” IETF RFC 6482, Feb. 2012.
- R. Bush and R. Austein, “The Resource Public Key Infrastructure (RPKI) to Router Protocol,” IETF RFC 6810, Jan. 2013.
- Resource Public Key Infrastructure (RPKI) [Online]. Available: https://www.arin.net/resources/rpki/index.html.
- M. Lepinski, S. Kent, and D. Kong, “A Profile for Route Origin Authorizations (ROAs),” *IETF RFC 6482*, Feb. 2012.