# Enforcing RPKI-Based Routing Policy on the Data Plane at an Internet Exchange

Josh Bailey
Google
joshb@google.com

Dean Pemberton,
Andy Linton
School of Engineering and
Computer Science
Victoria University of
Wellington, New Zealand
{dean,asjl}@ecs.vuw.ac.nz

Cristel Pelsser,
Randy Bush
IIJ
cristel@iij.ad.jp,
randy@psg.com

## ABSTRACT

Over a decade of work has gone into securing the BGP routing control plane. Through all this, there has been an oft repeated refrain, "It is acknowledged that rigorous control plane verification does not in any way guarantee that packets follow the control plane." We describe what may be the first deployment of data plane enforcement of RPKI-based control plane validation. OpenFlow switches providing an exchange fabric and controlled by a Quagga BGP route server drop traffic for prefixes which have invalid origins without requiring any RPKI support by connected BGP peers.

## 1. INTRODUCTION

A design group and the IETF SIDR Working Group have been developing a data repository known as the Resource Public Key Infrastructure (RPKI) [1], and protocols [2] [3] to validate BGP protocol announcements. While there are mechanisms in place to provide strong authentication between any two BGP peers, they do not address the validity of the IP prefix advertisements themselves - whether a given AS is entitled to make an announcement for a given prefix. Examples of networks being originated by ASs which had no legitimate authority, either by accident or design to do so, have been occuring frequently over a long period of time [4].

RFC6481 in particular describes how to improve accidental mis-announcements of an IP prefix from the wrong Autonomous System (AS). This is referred to as RPKI-based origin validation. Routers use RPKI data to validate the BGP announcements they receive, and drop announcements for prefixes with incorrect origin ASes. It has been suggested however, that even though RPKI-based origin validation looks to cryptographically validate that a given autonomous system is authorised to originate a given prefix, there is no mechanism currently in place to ensure that this 'control plane' level assertion is enforced in the 'data plane'.


*HotSDN'14,* August 22, 2014, Chicago, IL, USA.


## 2. DATA PLANE ENFORCEMENT VIA SDN

In response to the common security problems associated with Multilateral Internet Exchanges [5], over the last couple of years [6], the CARDIGAN SDN Internet Exchange (SDX) in New Zealand has been experimenting with the use of OpenFlow controlled exchange switches to enforce Layer-3 BGP policy at Layer-2 in order to prevent providers from accidentially or intentionally using a neighbor as a default exit and similar erroneous forwarding behavior. Exchange peers (REANNZ, and peers on the WIX) use a Quagga-based [7] route server [draft-ietf-idr-ix-bgp-route-server-03.txt] to exchange routes among themselves. Quagga then programs the OpenFlow switches to only forward traffic for those routes, with all other traffic dropped by default.

## 3. IMPLEMENTING RPKI ON CARDIGAN

To support RPKI on CARDIGAN two changes were made. The first change was to replace stock Quagga with the rtrlib version [8] (which enables Quagga to be configured to validate routes via an RPC) along with a certificate cache. The second change was to Quagga configuration, discarding routes that are specifically invalid (as opposed to being of unknown authenticity).

## 4. RESULTS

In February 2014, [9] observes 48 invalid, 92 valid and 1875 unvalidated routes out of 2015 routes at the Wellington Internet Exchange. This corresponds to 2.38%, 4.57% and 93.05% of the routes, respectively. Invalid routes included those that did not have ROA records with exactly matching prefix lengths. Counters on the OpenFlow default drop rule did not increase after RPKI deployment demonstrating that the forwarding policy had no negative impact on traffic.

## 5. LIMITATIONS

CARDIGAN uses "top of rack" style OpenFlow switches which are extremely limited in the number of rules they can process. This puts a limit on the number of prefixes - less than 1000 - but is expected given the original intended use for the hardware. OpenFlow hardware with greater rule capacity (Eg, NoviFlow) and OpenFlow multi-table support would greatly improve scaling.

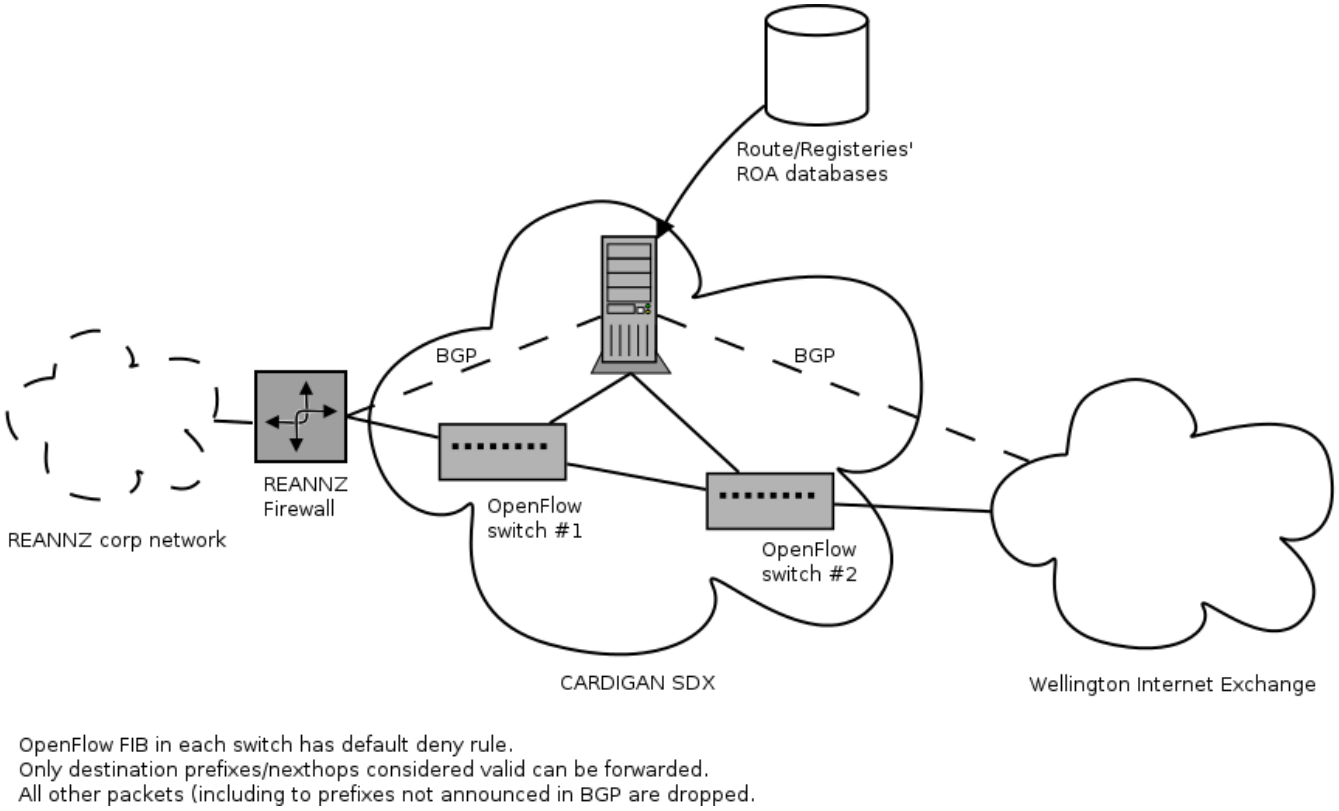


**Figure 1: CARDIGAN RPKI Overview**

## 6. CONCLUSION

RPKI on CARDIGAN uses SDN to directly address a number of existing operational problems on today's exchanges, in particular enforcing the consistency of routing announcements with forwarding. The scale of the deployment is primarily limited by the switch hardware. If hardware with more rules were available further traffic engineering and policy experiments could be attempted on the platform, for example uRPF (Unicast Reverse Path Forwarding).

## 7. REFERENCES


[1] G.˜Huston, R.˜Loomans, and G.˜Michaelson, "A Profile for Resource Certificate Repository Structure," RFC 6481 (Proposed Standard), Internet Engineering Task Force, Feb. 2012. [Online]. Available: http://www.ietf.org/rfc/rfc6481.txt

[2] R.˜Bush and R.˜Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol," RFC 6810 (Proposed Standard), Internet Engineering Task Force, Jan. 2013. [Online]. Available: http://www.ietf.org/rfc/rfc6810.txt

[3] P.˜Mohapatra, J.˜Scudder, D.˜Ward, R.˜Bush, and R.˜Austein, "BGP Prefix Origin Validation," RFC 6811 (Proposed Standard), Internet Engineering Task Force, Jan. 2013. [Online]. Available: http://www.ietf.org/rfc/rfc6811.txt

[4] (2008, March) Youtube hijacking: A RIPE NCC RIS case study. [Online]. Available: http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study

[5] M.˜Jager. (2012) Securing IXP connectivity. [Online]. Available: http://conference.apnic.net/__data/assets/pdf_file/0018/50706/apnic34-mike-jager-securing-ixp-connectivity_1346119861.pdf

[6] D.˜Pemberton. NZ scores first OpenFlow controlled connection to an IX. [Online]. Available: http://list.waikato.ac.nz/pipermail/nznog/2012-December/019635.html

[7] Quagga project. [Online]. Available: http://www.nongnu.org/quagga/

[8] GitHub: Quagga with RPKI-RTR prefix origin validation support. [Online]. Available: https://github.com/rtrlib/quagga-rtrlib

[9] M.˜Fincham. (2014, February) Rpki, nznog 2014. [Online]. Available: http://hotplate.co.nz/archive/nznog/2014/rpki/