June 2022

Status of RPKI in Australia and New Zealand

Government Services and Critical Infrastructure

Contributor
Terry Sweetser

MANRS

Table of Contents

Table of Figures

Project Task

Resource Public Key Infrastructure (RPKI) is a cryptographic method using the public key infrastructure (PKI) framework designed to secure the Internet's routing infrastructure, specifically the Border Gateway Protocol. This is the most important improvement that safeguards against **route mis-origination** attacks.

The data analyst will review how secure government services and other critical infrastructure are from route hijack/mis-origination attacks in Australia and New Zealand. The analysis will include the following:

- Network opacity testing from a central hosted site in Sydney, Australia.
- Reachability testing from ROA valid and invalid prefixes.
- Target GOV and EDU domains and other critical services that could be impacted by a lack of Route Origin Validation (ROV) compliance.
- Maximize routing diversity with four IXP connections and two national transit connections.

Introduction

In Australia and New Zealand, the Internet has become part of the critical infrastructure with governments and corporations now reliant on it for communications (Miron, 2015). The Australian Federal Government in particular is moving to a stance of protecting all essential digital services (*Protecting Critical Infrastructure and Systems of National Significance*, n.d.).

In this report, I focus on the integrity of routing information across the Internet in Australia and New Zealand. The primary concern of the data gathering was to see whether websites belonging to both public and private institutions can reject connections from clearly invalid sources of traffic. The risk to be highlighted is that these services could be accessed from untraceable sources by persons and organizations unknown and with little to no recourse by authorities to trace and track bad actors. Accidental and malicious route leaks are a hazard to the integrity of the Internet and the online safety of Internet users (Hall, 2012).

Another concern is the routing integrity of the host and transit networks between the end-user and those institutions and entities. I explore the idea that some of these websites could be hijacked and/or interfered with by persons unknown for malicious reasons (Wählisch et al., 2015).

I find that routing security is overall in a sub-optimal state in Australia and New Zealand. Problems arise with the acceptance of invalid routes but also with the acceptance of traffic from those invalid origins. There are exemplars where traffic and routes are dropped demonstrating what the best practice for routing security should look like for all network operators in the region.

Of concern in this sub-optimal environment is that businesses, government, and citizens are at greater risk of incidents where data is lost or stolen, or critical services interrupted.

Background

What is Internet Routing?

The Internet is literally coined from the phase “Internetwork of networks” (Hafner & Lyon, 1998).

Routing is the glue, the road network, or the “pipes” being used from point to point on all these networks to move information around the globe. Routers are devices programmed by network engineers to achieve that task. Routing is controlled logically, technically, and geopolitically by various organizations/networks across the globe. Organizations include (but is very much not limited to in any way) Network Information Centers, Internet Service Providers, Internet Exchange Points, and ICANN (Chander, 2017).

What is RPKI?

Resource Public Key Infrastructure (RPKI) is best described as a database used to secure the information used to route data on the Internet. Checking the origin of a route is a first step is ensuring that pathways through the Internet are valid. Other proposals and technical methods exist for validation of that path and the adjacencies of systems along that path (Bush, 2014; Chung et al., 2019; Lepinski & Kent, 2012; Levy, 2018).

Key concepts in Routing Security

IP Address	The Internet uses addresses to locate hosts and resources across the globe. Those addresses are unique, and ownership is recorded in central databases by various Network Information Centers.
Subnet	A range of IP addresses is referred to as a subnet due to the range being useful for adjacent devices to communicate directly on the same network using the same range. Expressions like CIDR and Prefix also refer to the same concept[1].
Bogons	IP addresses that are not allocated or not allocated for global Internet routing are referred as Bogons (*Routing Security Terms: Bogons, Vogons, and Martians*, n.d.).
Border Router	Between the various parts of the “Internetwork of networks” a device to join up the Internet is called a border router. Called thus as they form “borders” between networks.
Autonomous System Number (ASN)	An AS Number is allocated to each part of the Internet as operated by a service provider, network operator, exchange point, or telecommunications company. The number is unique globally and readily identifies the operating entity.
Border Gateway Protocol (BGP)	Between the various parts of the Internet, those border routers use this routing protocol to inform each other of pathways they provide across the Internet.
Route Origin Authorization (ROA)	This is a cryptographically signed record that relates the IP address range to the AS Number. The relationship expressed by this record is used to validate the origin of the route.

1 CIDR as per https://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing

Most Specific Announcement (MSA)	Larger subnets can be split across multiple end-users and sites. If those smaller networks are to be advertised by BGP into the global table the MSA value needs to be large enough to allow them to be validated across the Internet ("How To," 2019).
Domain Name System (DNS)	Also known as DNS, this is the global distributed database of records mapping names to IP addresses.
Route Origin Validation (ROV)	Border Routers can (but do not always) check incoming route information against the recorded ROA. Three possible outcomes exist: **VALID:** An ROA matches the origin ASN and IP address range. **INVALID:** There is no match, but a record exists for the subnet that defines a different set of one or more ASNs. When the subnet is too small for the public ROA record, that is also considered invalid. **UNKNOWN:** There is no match for the IP address range at all.

Figure 1: Key Concepts in Routing Security

What is a Route Leak?

"A route leak is the propagation of routing announcement(s) beyond their intended scope. That is, an announcement from an Autonomous System (AS) of a learned BGP route to another AS is in violation of the intended policies of the receiver, the sender, and/or one of the ASes along the preceding AS path." (Robachevsky & Christopher, 2016; Sriram et al., 2016)

One of the more famous route leaks was in 2008 when a Pakistani ISP accidentally routed all of YouTube to itself (Singel, 2008). The effective outcome of the route leak was to see YouTube traffic from across the globe make its way to one ISP in Pakistan. Route leaks tend to be accidental in nature, however some sovereign governments have attempted to censor the Internet with a route hijack (Padmanabhan et al., 2021).

What is a Route Hijack?

A hijack sounds nefarious: the outcome of a successful hijack places a bad actor in control of address space belonging to the victim. If the victim usually places resources online for the public or customers using that address space, then the bad actor is now presenting a false front to the public and customers relying on those resources. For instance, if the victim is a bank, then impersonating the bank would lead to substantial malfeasance. A very well covered hijack happened in 2018 to steal cryptocurrencies (Poinsignon, 2018).

Hijacks can also be used against organizations with bad actors using stolen address space to hide their activities and avoid detection by authorities. This scenario is a major premise for the testing conducted to support this report.

"Request For Comments"

The large cache of documents found at rfc-editor.org constitute the standards and practices of the global technical community and devices that form the routing infrastructure of the Internet.

BCP38/RFC2827

Best Current Practice number 38 describes a policy to be implemented by network operators to minimize the damage caused by spoofed IP addresses. Lack of proper outbound filters has meant that poorly managed networks can be a source of attacks, specifically denial of service (DoS) attacks that use fake source addresses.

MANRS

The global effort to ensure that each IP address is covered by a Route Origin Authorization is part of the work of the Internet Society with the Mutually Agreed Norms for Routing Security (*MANRS*, n.d.).

Methodology

Gathering data for this project involved capture of network information for 870,011 Australian domains (ending in .au) and 197,009 New Zealand domains (ending in .nz) with a control group of 38,232 domains from Pakistan with domains ending in .pk. The methodology was tested extensively over routes to Pakistan.

Domain Name System

2,610,020 total DNS records were captured as the first step in each test. 706,743 records were CNAMES or an alias to another domain name. While 271,330 records were IP address version six (IPv6), these destinations were not tested. That left 1,633,947 IP version four (IPv4) addresses recorded and eventually tested. The origins of each test were either a valid or an invalid ROA IPv4 address.

Traffic Source

Each tested domain was given one or more tests. The origin of the tests which I will refer as the valid and invalid sources were two different IP addresses from two adjoining address ranges. One had a published ROA that validated the source of the traffic: the valid source. The other was purposely set up with an invalid origin, failing to match the published ROA and hence becomes the invalid source herein.

Connectivity

The first step in each test was to make a TCP/443 connection to the target address. This test, if successful, would verify that a web server was running HTTPS (Hypertext Transfer Protocol Secure) and would reply to TCP traceroute to the host on port 443. There were occasions, for some targets, this test continuously failed for the invalid source with no response ever received. **Only successful attempts are recorded in traces.** Counts for targets only with valid source results clearly indicated the target was immune to ROA attacks.

Trace routes

For each successful connection, a path was recorded for the outgoing packets. In Internet terms, it is a traceroute which returns information about each hop as the test passes each router along the route. Note here that traces record the outbound route, source to destination. It is entirely possible for the return path to be different (Luckie et al., 2008).

Route paths

Each hop in the trace to the destination was mapped to an ASN. This is then vectorized to give a path of Autonomous Systems across the Internet. Note again, the path being explored is source to destination. The mapping used APIs like Team Cymru and ATLAS to get public data. While it was entirely possible to gather data from routing tables on Border Routers, this may not have been strictly true of the actual path taken, due to the policies of intermediate Autonomous Systems across the Internet.

Data Summary

Unique Traceroutes from all sources	2,333,647
Unique Domains all destinations	1,002,493
Tested top level country domains	AU NZ PK[2]
Tested second level domains	school.nz govt.nz asn.au org.au net.pk edu.au com.au gos.pk org.pk net.nz co.nz org.nz id.au com.pk net.au gov.au edu.pk

Figure 2: Main Data Summary Table (Source: Terry Sweetser)

Tests carried out per second-level and country-code top level domain:

Domain Suffix	Number of Domains	Number of Tests
au	107	219
nz	15544	34713
pk	18028	38189
asn.au	1675	4028
co.nz	140057	307579
com.au	728984	1720394
com.pk	12671	27707
edu.au	6133	14838
edu.pk	1755	3821
gos.pk	48	113
gov.au	2137	5084
govt.nz	592	1531
id.au	1233	3052
net.au	23239	56155
net.nz	2628	5950
net.pk	423	966
org.au	33597	80206
org.nz	10035	22808
org.pk	1401	3085
school.nz	2206	5209

Figure 3: Number of Domains and Tests (Source: Terry Sweetser)

2 PK is the Pakistan country code top level domain; it was used in the data for testing and verification of methodologies.

The result counts from each source per second-level and country-code top level domain:

Domain		Source (Number)		Source (Percentile)		
Country	Second	VALID	INVALID	VALID	INVALID	RATIO
au	com	1,108,717	611,677	64.4%	35.6%	55.2%
nz	co	213,361	94,218	69.4%	30.6%	44.2%
au	org	50,715	29,491	63.2%	36.8%	58.2%
au	net	35,148	21,007	62.6%	37.4%	59.8%
pk		26,382	11,807	69.1%	30.9%	44.8%
nz		24,128	10,585	69.5%	30.5%	43.9%
pk	com	18,325	9,382	66.1%	33.9%	51.2%
nz	org	15,295	7,513	67.1%	32.9%	49.1%
au	edu	9,150	5,688	61.7%	38.3%	62.2%
nz	net	3,890	2,060	65.4%	34.6%	53.0%
nz	school	3,416	1,793	65.6%	34.4%	52.5%
au	gov	3,390	1,694	66.7%	33.3%	50.0%
au	asn	2,530	1,498	62.8%	37.2%	59.2%
pk	edu	2,552	1,269	66.8%	33.2%	49.7%
au	id	1,822	1,230	59.7%	40.3%	67.5%
pk	org	2,056	1,029	66.6%	33.4%	50.0%
nz	govt	927	604	60.5%	39.5%	65.2%
pk	net	625	341	64.7%	35.3%	54.6%
au		162	57	74.0%	26.0%	35.2%
pk	gos	69	44	61.1%	38.9%	63.8%

Figure 4: Number of tests per top and second level domain segments with VALID and INVALID source counts (Source: Terry Sweetser)

Observation 1

The ratio of destinations allowing a connection from the invalid source is very high in one-third of all tests.

Diving into AU domains in the GOV and EDU second level:

Tests	Valid	Invalid	Domains	ASN	Ratio
1471	738	733	edu.au	55803	**99.3%**
1366	684	682	edu.au	45638	**99.7%**
1187	606	581	edu.au	15169	**95.9%**
1082	1082	0	edu.au	14618	**0.0%**
860	429	431	edu.au	45671	**100.5%**
833	435	398	gov.au	55532	**91.5%**
743	377	366	edu.au	4739	**97.1%**
553	553	0	edu.au	13335	**0.0%**
454	256	198	edu.au	38719	**77.3%**
384	384	0	gov.au	14618	**0.0%**
370	184	186	gov.au	20940	**101.1%**
312	164	148	edu.au	45768	**90.2%**
302	153	149	edu.au	9517	**97.4%**
295	142	153	edu.au	135543	**107.7%**
290	154	136	edu.au	9650	**88.3%**
285	136	149	gov.au	19551	**109.6%**
284	147	137	gov.au	56135	**93.2%**
247	111	136	edu.au	132680	**122.5%**
238	231	7	edu.au	16509	**3.0%**
236	236	0	gov.au	8075	**0.0%**
230	230	0	edu.au	8075	**0.0%**
221	221	0	gov.au	16509	**0.0%**
220	108	112	edu.au	27647	**103.7%**
214	111	103	edu.au	136557	**92.8%**
204	106	98	edu.au	26496	**92.5%**
192	92	100	edu.au	139344	**108.7%**
179	179	0	edu.au	20473	**0.0%**
171	171	0	gov.au	13335	**0.0%**
160	88	72	edu.au	7575	**81.8%**
160	95	65	edu.au	46606	**68.4%**

Figure 5: Test and result counts for GOV.AU and EDU.AU per ASN (Source: Terry Sweetser)

Observation 2

Breaking down the results by destination Autonomous System number immediately shows that some destinations are rejecting all connections from the invalid source.

Counts per ASN for domains:

Tests	Valid	Invalid	ASN	Ratio
257348	122247	135101	15169	**110.5%**
206004	103572	102432	45638	**98.9%**
202473	202473	0	13335	**0.0%**
183606	100734	82872	38719	**82.3%**
157599	79076	78523	55803	**99.3%**
97878	97877	1	14618	**0.0%**
78503	41811	36692	26496	**87.8%**
61059	60915	144	58182	**0.2%**
53412	26951	26461	45671	**98.2%**
47883	23598	24285	27647	**102.9%**
44980	25952	19028	46606	**73.3%**
40350	40350	0	53831	**0.0%**
37479	18773	18706	63956	**99.6%**
33200	32450	750	174	**2.3%**
33158	32705	453	16509	**1.4%**
32642	16378	16264	132680	**99.3%**
32449	16153	16296	16276	**100.9%**
23478	11719	11759	54113	**100.3%**
22418	10469	11949	14061	**114.1%**
21780	21780	0	20473	**0.0%**
20269	8875	11394	13768	**128.4%**
19087	17609	1478	03	**8.4%**
18465	9278	9187	133159	**99.0%**
18294	18286	8	8075	**0.0%**
18258	8636	9622	133618	**111.4%**
18055	9075	8980	136557	**99.0%**
17863	9271	8592	63949	**92.7%**
16704	16697	7	19527	**0.0%**
16701	8359	8342	18108	**99.8%**
15348	15348	0	45179	**0.0%**

3 AS0 represents subnets that fall outside normal public ownership data or were classed as bogons.

Observation 3

There are several well-known networks rejecting all invalid ROA traffic. Other operators at the same scale were accepting as many connections as possible from both valid and invalid sources.

Observation 4

1,141 networks (as ASNs) showed the behavior of rejecting connections from the invalid address space. 1,641 other networks allowed connections from the invalid source.

The Top Tens

Comparing networks with no penetration by invalid ROA with those with maximum penetration:

Tests	Valid	Invalid	ASN		Tests	Valid	Invalid	ASN
202473	202473	0	13335		257348	122247	135101	15169
40350	40350	0	53831		206004	103572	102432	45638
21780	21780	0	20473		183606	100734	82872	38719
15348	15348	0	45179		157599	79076	78523	55803
7348	7348	0	45459		78503	41811	36692	26496
4828	4828	0	9714		53412	26951	26461	45671
4067	4067	0	32787		47883	23598	24285	27647
2891	2891	0	209242		44980	25952	19028	46606
2139	2139	0	16552		37479	18773	18706	63956
1670	1670	0	9889		33200	32450	750	174

Figure 6: Top 10 Networks (Source: Terry Sweetser)

Traceroute Comparisons

To cyber.gov.au from the VALID source:

route_id	Name	Source	Hop	Router	Prefix	ASN	RIR
77	cyber.gov.au	VALID	1	103.162.142.65	103.162.142.0/24	141384	APNIC
2	cyber.gov.au	VALID	2	103.126.52.154	103.126.52.0/24	138466	APNIC
2	cyber.gov.au	VALID	3	103.126.52.163	103.126.52.0/24	138466	APNIC
3	cyber.gov.au	VALID	4	202.90.50.104	202.90.48.0/21	9336	APNIC
53	cyber.gov.au	VALID	5	203.153.18.96	203.153.16.0/22	38195	APNIC
485	cyber.gov.au	VALID	6	103.200.13.72	103.200.12.0/22	38195	
485	cyber.gov.au	VALID	7	103.200.13.64	103.200.12.0/22	38195	
485	cyber.gov.au	VALID	8	103.200.13.153	103.200.12.0/22	38195	
666	cyber.gov.au	VALID	9	202.177.40.22	202.177.40.0/24	0	
1944	cyber.gov.au	VALID	11	184.28.235.107	184.28.235.0/24	20940	RIPENCC

To cyber.gov.au from the INVALID source:

route_id	Name	Source	Hop	Router	Prefix	ASN	RIR
7	cyber.gov.au	INVALID	1	103.162.143.65	103.162.143.0/24	141384	APNIC
2	cyber.gov.au	INVALID	2	103.126.52.154	103.126.52.0/24	138466	APNIC
2	cyber.gov.au	INVALID	3	103.126.52.163	103.126.52.0/24	138466	APNIC
3	cyber.gov.au	INVALID	4	202.90.50.106	202.90.48.0/21	9336	APNIC
53	cyber.gov.au	INVALID	5	203.153.18.96	203.153.16.0/22	38195	APNIC
485	cyber.gov.au	INVALID	6	103.200.13.72	103.200.12.0/22	38195	
485	cyber.gov.au	INVALID	7	103.200.13.64	103.200.12.0/22	38195	
485	cyber.gov.au	INVALID	8	103.200.13.153	103.200.12.0/22	38195	
666	cyber.gov.au	INVALID	9	202.177.40.22	202.177.40.0/24	0	
1944	cyber.gov.au	INVALID	11	184.28.235.88	184.28.235.0/24	20940	RIPENCC

To gms.school.nz from the valid source:

route_id	Name	Source	Hop	Router	Prefix	ASN	RIR
77	gms.school.nz	VALID	1	103.162.142.65	103.162.142.0/24	141384	APNIC
2	gms.school.nz	VALID	2	103.126.52.154	103.126.52.0/24	138466	APNIC
2	gms.school.nz	VALID	3	103.126.52.163	103.126.52.0/24	138466	APNIC
3	gms.school.nz	VALID	4	202.90.50.104	202.90.48.0/21	9336	APNIC
53	gms.school.nz	VALID	5	203.153.18.96	203.153.16.0/22	38195	APNIC
71	gms.school.nz	VALID	6	154.54.45.121	154.48.0.0/12	174	ARIN
71	gms.school.nz	VALID	7	154.54.31.157	154.48.0.0/12	174	ARIN
485	gms.school.nz	VALID	8	103.200.13.168	103.200.12.0/22	38195	
72	gms.school.nz	VALID	9	62.115.179.16	62.115.0.0/16	1299	RIPENCC
72	gms.school.nz	VALID	11	62.115.123.136	62.115.0.0/16	1299	RIPENCC
72	gms.school.nz	VALID	13	62.115.113.85	62.115.0.0/16	1299	RIPENCC
72	gms.school.nz	VALID	14	62.115.188.199	62.115.0.0/16	1299	RIPENCC
593	gms.school.nz	VALID	15	206.123.64.17	206.123.64.0/18	30496	ARIN
594	gms.school.nz	VALID	16	207.210.229.6	207.210.192.0/18	30496	ARIN
594	gms.school.nz	VALID	17	207.210.229.50	207.210.192.0/18	30496	ARIN
1359	gms.school.nz	VALID	20	74.50.49.14	74.50.48.0/20	36024	ARIN

To gms.school.nz from the invalid source:

route_id	Name	Source	Hop	Router	Prefix	ASN	RIR
2	gms.school.nz	INVALID	2	103.126.52.154	103.126.52.0/24	138466	apnic
2	gms.school.nz	INVALID	3	103.126.52.163	103.126.52.0/24	138466	apnic
3	gms.school.nz	INVALID	4	202.90.50.104	202.90.48.0/21	9336	apnic
53	gms.school.nz	INVALID	5	203.153.18.96	203.153.16.0/22	38195	apnic
1359	gms.school.nz	INVALID	10	74.50.49.14	74.50.48.0/20	36024	arin
72	gms.school.nz	INVALID	14	62.115.188.199	62.115.0.0/16	1299	ripencc
593	gms.school.nz	INVALID	15	206.123.64.17	206.123.64.0/18	30496	arin
594	gms.school.nz	INVALID	16	207.210.229.6	207.210.192.0/18	30496	arin
594	gms.school.nz	INVALID	17	207.210.229.50	207.210.192.0/18	30496	arin
1359	gms.school.nz	INVALID	18	74.50.49.14	74.50.48.0/20	36024	arin
1359	gms.school.nz	INVALID	23	74.50.49.14	74.50.48.0/20	36024	arin

Observation 6

Where both valid and invalid connections succeed, the data distinctly indicates that not only does the destination allow the invalid connection, but the intermediate networks are also passing traffic indiscriminately.

Destination Route Validity

The cached IP address data from the test also captured the ROV of the destinations. 12,569 prefixes remain in the UNKNOWN category, whilst 1,113 were VALID, and 84 were INVALID. Of those 84 invalid prefixes, data shows 143,699 tests succeeding to those networks. Of those 143,699 tests, there are 101,251 “last hops” in the tests that show address space which is not routed and/or cannot be routed.

Examination of the data revealed the example of https://qld.gov.au on 111.118.196.23 and 111.118.196.29 in the subnet of 111.118.196.0/24. This site was accessed successfully from both valid and invalid sources. The ROV status of 111.118.192.0/21 covers the subnet with a VALID ROA, however insufficient public data existed to cover the 111.118.196.0/24 subnet in well know databases like ATLAS and Team Cymru.

Observation 7

Public data sources for mapping IP addresses to ASNs, like Team Cymru and ATLAS, are not 100% accurate.

Discussion

Data Quality

The code used to generate the data set went through several iterations to get reasonably accurate and useful data. Testing the code proved very useful in discovering more pitfalls in network testing and also validating previous findings from studies that have tried to examine similar problems with traceroute data. Choosing to strictly use a TCP trace to 443 after a SYN/ACK test provided the best results overall (Luckie et al., 2008).

Matching IP address ownership to various organizations and AS numbers proved to be trivial for 99% of the dataset and then very difficult for a minor number of IP addresses. Certain public databases, like RADB, were avoided due to the wildly inaccurate data found within them. Commercial sources were not used after finding in other work to be superficially accurate.

One notion totally avoided was the use of current BGP data to map an IP address to an ASN. What BGP contains and what should be there overlap, but BGP is not the source of truth that we would desire it to be.

So, traceroute being what it is, some packet loss and other issues like control packet (ICMP) rate limits and misconfigured firewalls meant that while the SYN/ACK test would succeed, a traceroute would be decimated.

Insecure BGP

Border Gateway Protocol has lacked critical security features since its invention. RPKI is a first major step towards having the routing information exchanged by border routers checked by a trusted source for validity. Even the simple act of getting an ROA assigned to your IP address allocations starts to provide a certain amount of protection. Despite the rewards not all routes are signed and not all borders are running ROV (Chung et al., 2019; Wählisch et al., 2015).

Data gathered for this report demonstrates that most addresses tested are in networks not signing (no ROA) and not validating (no ROV). With these measures missing, an ROA attack would be able to redirect traffic away from these networks. Further, if the adjacent networks are also weak on routing security, these networks are going to impact global routing in the event of a major accidental leak or intentional hijack.

Sovereign action is extremely unlikely in Australia and New Zealand, but a hijack could be executed for purely criminal purposes on or against these networks.

Poor Practices

Observations already detailed in this report clearly indicate that networks in Australia and New Zealand, as well as several international networks, are accepting traffic from address space with a clearly invalid Route-Origin.

Further, those not using ROV are also posing a risk to RPKI-enabled networks: they would accept a hijacked route from adjacent networks despite its status being invalid.

The data even contains examples of network operators who have valid ROA records in place and therefore afford some protection over their own network, yet allow traffic from a clearly invalid source to transit their network to/from other unprotected parties.

In the case where a network may not be accepting the invalid route but traffic still transits the network successfully, there is an explanation that they use a catch-all or default route to another network.

Lastly, the number of bogon IP addresses sending replies to traceroutes are significantly large and showed some networks are not engaged in basic source filtering hygiene.

Good Practices

For certain destination networks there was clear evidence of a policy to drop connections from IP addresses with invalid ROV results. The test network with the invalid origin remained in that state for the entire test period and yet certain networks never responded to a probe.

That indicates that these networks:

1. Do not accept BGP updates with an invalid ROA; and
2. Do not have a catch-all route to another network; and
3. Possibly also place the failed route origin into a dynamic access control list that rejects connections from it.

Network Security Implications

RPKI Uptake

Data gathered unfortunately clearly indicates that RPKI uptake is still low enough to enable hijacks and leaks to place many destinations at risk. Included in those destinations are many government and education services. Interruptions to these services would be of particular concern during the COVID-19 pandemic when isolation and work-from-home are countermeasures.

Same Origin Attacks

Even with ROA and ROV in place universally there is still the ploy of duplicating an origin's ASN and IP address range. RPKI by itself is not a security panacea. The role to be played by high uptake of RPKI in Australia and New Zealand would be to mitigate accidental leaks and malicious hijacks with a clearly invalid origin.

ROA Signing Errors

A common recurring problem is the mistake by operators not matching their ROA records to what they configure BGP to do on their network. It is not a good outcome to find the Internet will take your route due to an error in your input. However, from any security perspective, a status of invalid or mismatch in routing data should always result in failure ("RPKI Invalids Are Not Going Away," 2021).

Good MANRS

Deployment of RPKI is one of many useful and necessary methods for routing security. Filtering routes and traffic forms a large part of ensuring the routing data on any network is accurate, secure, and fit-for-purpose. Operators should consider joining the MANRS initiative and working towards a more secure Internet routing system.

Future Developments

While origin validation is now a reality on the BGP global table, work continues securing global Internet routing. Techniques such as BGPSec, ASPA, and FS-BGP are being discussed in academia and Internet Engineering Task Force meetings (Aelmans, 2020; Bush et al., 2019).

The main barriers to new standards are always whether the feature can be added ad-hoc or discretely with no disturbance to operators not running it. BGPSec is an example of "All Or Nothing" for securing BGP due to it being required across the entire Internet. ASPA, on the other hand, presents an opportunity for routing pathways to be discretely verified (like ROV) by users without interrupting global Internet operations.

Critical Infrastructure

The Internet has become critical infrastructure. Data processing centers and telecommunication networks combine to be what we would call the Internet and the cloud. Data gathered for this report

clearly shows that Route-Origin attacks can impact half of all domains in Australia and New Zealand directly by leaks and hijacks. That split of approximately half of all domains extends to government and education websites.

Conclusions

Routing security in Australia and New Zealand is in a poor state. Even with RPKI uptake increasing there are practices in various operator networks that at worst do not meet the basic hygiene requirements and at best still allow the transit of traffic from address space that has been clearly marked as invalid. Some of the big operators like Telstra and Vocus are doing the right things to protect their downstream network from such attacks but unless all major operators do the same these attacks will find their way to the destinations.

There are plausible scenarios where Australian and New Zealand corporations and citizens can be impersonated, suffer data loss or theft, or have critical services interrupted.

My key recommendations here are:

1. All organizations in the region sign their Route Origin Authorization on their network address space;
2. All transit providers, from the largest to the smallest, commence Route Origin Validation on all border routers handling public Internet traffic;
3. All Internet service providers dynamically build access control lists on their border routers to either reject, drop, or blackhole routes with an invalid ROA status.
4. All Internet service providers aspire to become MANRS participants and comply with the best practices of the MANRS initiative.

References

Aelmans, M. (2020). Global Routing Operations J. Snijders Internet-Draft NTT Intended status: Informational M. Stucchi Expires: October 26, 2020 Independent.

Bush, R. (2014). Origin validation operation based on the Resource Public Key Infrastructure (RPKI). IETF RFC7115 (January 2014).

Bush, R., Patel, K., Snijders, J., & Housley, R. (2019). Network Working Group A. Azimov Internet-Draft Yandex Intended status: Standards Track E. Uskov Expires: November 18, 2019 Qrator Labs.

Chander, A. (2017). Who runs the Internet? In Research Handbook on the Politics of International Law. Edward Elgar Publishing.

Chung, T., Aben, E., Bruijnzeels, T., Chandrasekaran, B., Choffnes, D., Levin, D., Maggs, B. M., Mislove, A., Rijswijk-Deij, R. van, & Rula, J. (2019). RPKI is coming of age: A longitudinal study of RPKI deployment and invalid route origins. 406–419.

Hafner, K., & Lyon, M. (1998). Where wizards stay up late: The origins of the Internet. Simon and Schuster.

Hall, C. (2012). Security of the Internet and the Known Unknowns. Communications of the ACM, 55(6), 35–37.

How to: Creating RPKI ROAs in MyAPNIC. (2019, September 11). APNIC Blog. https://blog.apnic.net/2019/09/11/how-to-creating-rpki-roas-in-myapnic/

Lepinski, M., & Kent, S. (2012). An infrastructure to support secure Internet routing.

Levy, M. J. (2018, September 19). RPKI - The required cryptographic upgrade to BGP routing. The Cloudflare Blog. http://blog.cloudflare.com/rpki/

Luckie, M., Hyun, Y., & Huffaker, B. (2008). Traceroute probe method and forward IP path inference. 311–324.

MANRS. (n.d.). MANRS. Retrieved November 22, 2021, from https://www.manrs.org/

Miron, W. (2015). Q&A: Should the Internet Be Considered Critical Infrastructure? Technology Innovation Management Review, 5(1), 37–40.

Padmanabhan, R., Filastò, A., Xynou, M., Raman, R. S., Middleton, K., Zhang, M., Madory, D., Roberts, M., & Dainotti, A. (2021). A multi-perspective view of Internet censorship in Myanmar. 27–36.

Poinsignon, L. (2018, April 24). BGP leaks and cryptocurrencies. The Cloudflare Blog. http://blog.cloudflare.com/bgp-leaks-and-crypto-currencies/

Protecting Critical Infrastructure and Systems of National Significance. (n.d.). Retrieved November 22, 2021, from https://www.homeaffairs.gov.au/reports-and-publications/submissions-and-discussion-papers/protecting-critical-infrastructure-systems

Robachevsky, A., & Christopher, D. (2016, June 23). RFC 7908 Defines BGP Route Leaks. Now, How Do We Prevent Them? Internet Society. https://www.Internetsociety.org/blog/2016/06/rfc-7908-defines-bgp-route-leaks-now-how-do-we-prevent-them/

Routing Security Terms: Bogons, Vogons, and Martians. (n.d.). Retrieved January 20, 2022, from https://www.manrs.org/2021/01/routing-security-terms-bogons-vogons-and-martians/

RPKI invalids are not going away. (2021, July 16). APNIC Blog. https://blog.apnic.net/2021/07/16/rpki-invalids-are-not-going-away/

Singel, R. (2008, February 25). Pakistan's Accidental YouTube Re-Routing Exposes Trust Flaw in Net. Wired. https://www.wired.com/2008/02/pakistans-accid/

Sriram, K., Montgomery, D., McPherson, D. R., Osterweil, E., & Dickson, B. (2016). Problem Definition and Classification of BGP Route Leaks (Request for Comments RFC 7908). Internet Engineering Task Force. https://doi.org/10.17487/RFC7908

Wählisch, M., Schmidt, R., Schmidt, T. C., Maennel, O., Uhlig, S., & Tyson, G. (2015). RiPKI: The tragic story of RPKI deployment in the Web ecosystem. 1–7.

Appendices