



Being a better Netizen: MANRS @ DO

Mircea Ulinic

*Senior Network Engineer,
Network Development*

ESNOG 2023 - Madrid
May 18-19, 2023

Contents

Who is DigitalOcean?	03
What is MANRS?	04
Why become MANRS compliant?	05
Requirements	06-07
What we did	08-33
What's next	33-35



San Francisco

Toronto

New York

London


Amsterdam

Frankfurt

Bangalore

Singapore

Sydney



14 data centers in 8 global markets

An intro to MANRS

Mutually Agreed Norms for Routing Security (MANRS) is a global initiative, supported by the Internet Society, that provides crucial fixes to reduce the most common routing threats.

MANRS operates four key programs that target Network Operators, Internet Exchange Points, Cloud and CDN Providers and Equipment Vendors to help improve their routing security posture.

MANRS also offers Fellowship and Ambassador programs to engage motivated individuals of the Internet community. By participating you'll talk about routing security issues and provide valuable input to drive the state of the art among existing and new member organizations.



Mutually Agreed
Norms for Routing
Security



Operated by the
Internet Society since
2014



Four key programs for
participation

Why become MANRS Compliant?

**Our community is
bigger than us.**

- A Core Value at DigitalOcean

MANRS Cloud & CDN Program



Action 1:

Prevent propagation of incorrect routing information.

“... Whenever feasible, participants should check that the announcements originate from legitimate holders.”



Action 2:

Prevent traffic with illegitimate source IP addresses

“Implement anti-spoofing controls to prevent packets with illegitimate source IP address from leaving the network.”



Action 3:

Facilitate global operational communication and coordination

“Maintain globally accessible up-to-date contact information in PeeringDB and relevant RIR databases.”

MANRS Cloud & CDN Program



Action 4:

Facilitate validation of routing information on a global scale

"... routing information needs to be properly registered in public routing repositories...The two main types of repositories are IRRs and RPKI."



Action 5:

Encourage MANRS adoption

"A publicly available policy, a peering form or an email template with a recommendation to implement MANRS."

Action 1:

Filtering

Prevent propagation of incorrect routing information.

Action 1: Filtering

Challenges

DigitalOcean runs a medium-large global network that peers with hundreds of ASNs on many of the biggest peering fabrics in the world.

Analysis and automation is required to find a workable solution that provides appropriate knobs to control for our scale.



High cardinality of peering sessions



Varying hardware capacity



Automation required

Action 1: Filtering

Analysis

By collating IRR data related to our peers, we were able to look at the configuration impact of per-peer prefix filtering.

What is the sweet spot for making a meaningful improvement in routing security without compromising our current platform?



Generate prefix list length for each IRR object.




Aim to make a meaningful improvement to Routing Security



Don't compromise the network of today.

Action 1: Filtering

Analysis



Action 1: Filtering

Analysis

A cumulative graph provides a clearer perspective for determining where to draw the line to meet our needs today.

Limit the number of config lines to suite today's platform but pick a limit that can be increased in the future easily.



100% coverage would result in ~6.5M LoC on some routers.



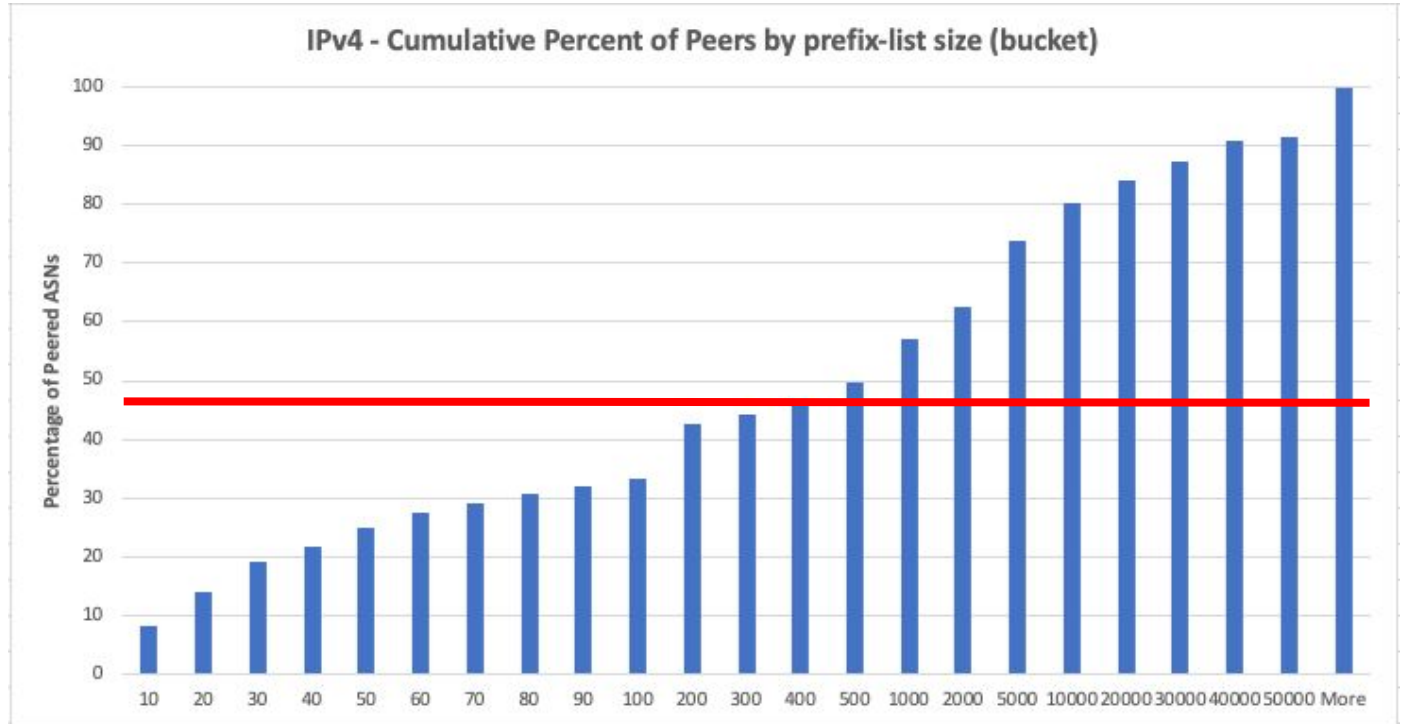
2M LoC ~ 95 sec apply times.



Picked a sensible point to maximise coverage within limitations.

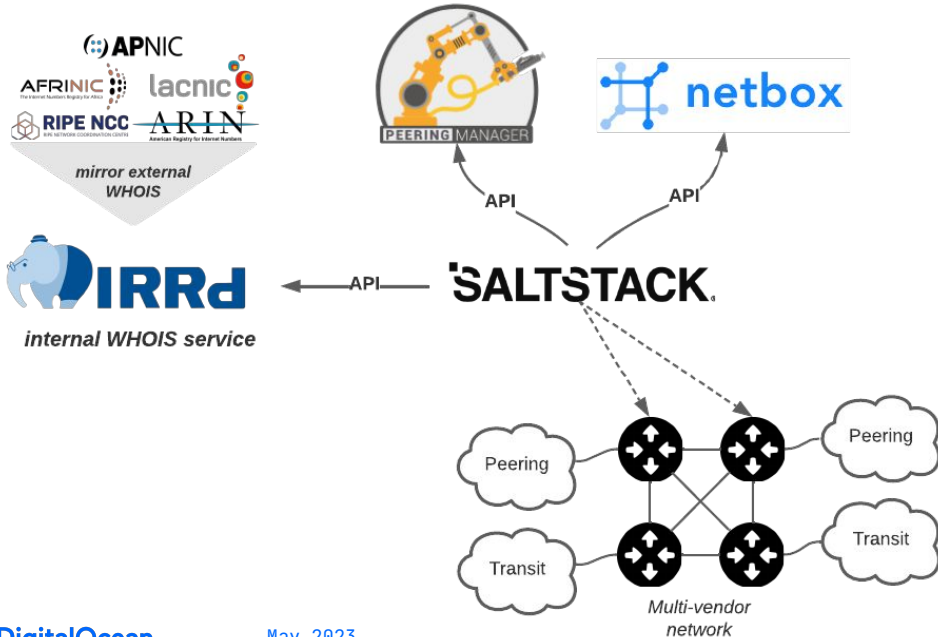
Action 1: Filtering

Analysis



Action 1: Filtering

The heavy lift...



IRRd, NetBox and Peering Manager for Source-of-Truth



SaltStack to build prefix-lists and templates.



Continual, automatic updates pushed to the network every 6 hours.

Action 1: Filtering

The outcome...

```
neighbor 192.0.2.1 {  
  description "Example Network Name";  
  import [ SCRUB-IMPORT-IPv4  
          IX-IMPORT  
          ABCIX-BILAT-IMPORT-IPv4  
          AS65535-IX-IMPORT-IPv4 ];  
  family inet {  
    unicast {  
      prefix-limit {...}  
    }  
  }  
  export ABC-PEER-EXPORT;  
  peer-as 65535;  
}
```



Policy per peer



Chained with other policies, optional completion early.



Easy to read and understand.

Action 1: Filtering

The outcome...

```
> show policy-options policy-statement \  
AS65535-IX-IMPORT-IPv4
```

```
term ALLOW-RPKI-VALID {...}  
term AS65535 {  
  from {  
    protocol bgp;  
    prefix-list AS65535v4; ←  
  }  
  then {  
    community add DO-IRR-VALID;  
    accept;  
  }  
}  
term DEFAULT-REJECT {...}
```



Filter on RPKI first



Filter by IRR second



Tag with useful communities as you go.

Action 1: Filtering

The outcome...

route-policy RP-AS65535-IMPORT-IPv4

apply RP-SCRUB-IN-IPV4

apply RP-IX-IMPORT

apply RP-ABCIX-BILAT-IMPORT-IPv4

if validation-state is valid then

pass

else if destination in AS65535v4 then

pass

else

drop

end-policy



Policy per peer



Re-use of standard policies.



Easy to read and understand.

Action 2:

Anti-Spoofing

Prevent traffic with illegitimate source IP addresses

We already prevent spoofing!

Given DigitalOcean runs such a huge number of workloads, bad actors and spoofed traffic isn't a new challenge. We already have several layers of protection to ensure that all traffic originating from DO is from legitimate sources.



Control from the hypervisor.



Internal detection tooling.



uRPF deployed on the edge of the network.

To be sure?

It's best practice to ensure that our mechanisms to prevent spoofing are actually working. When they aren't, we want to have a clear signal when they no longer are.

The CAIDA Spoofer project to the rescue! We run spoofer nodes in each of our DCs that attempt to send spoofed traffic to the public CAIDA endpoint. A prometheus exporter regularly queries the public CAIDA API and will alert us if spoofed traffic is received.



CAIDA Spoofer Project



Prometheus Exporter



Sensible alerting rules
with a playbook

Action 2: Anti-Spoofing

To be sure?

```
alerts:
- alert: CAIDA Session Received
  expr: caida_spoofers_session == 1
  labels:
    service: CAIDA
    severity: warning
    instance: "caida-spoofers::{{$labels.session}}"
  annotations:
    URL: <a
href="https://spoofers.caida.org/report.php?sessionId={{$labels.
session}}" target="_blank">Session {{$labels.session}}
Report</a>
    playbook: <a
href="https://doplaybooksite.tld/CAIDA+Spoofers+Servers"
target="_blank">CAIDA Spoofers</a>
```



CAIDA Spoofer Project



Prometheus Exporter



Sensible alerting rules
with a playbook

Action 3:

Coordination

Facilitate global operational communication and coordination

How to find us...

We keep our WHOIS data up-to-date as we on-board new IP space through a regularly used playbook. This ensures all the same data is present on all our prefixes:

OrgTechEmail: noc@digitalocean.com
OrgNOCEmail: noc@digitalocean.com
OrgAbuseEmail: abuse@digitalocean.com



Consistent WHOIS data through defined process



Accurate PeeringDB record



Monitored mailboxes

Action 3: Coordination

How to find us...

Because we rely on our peering partners to keep their PeeringDB record up-to-date for automation reasons, we should set the best example and do so as well.

DigitalOcean		Peering Policy Information			
Organization	DigitalOcean	Peering Policy	https://www.as14061.net/		
Also Known As	Digital Ocean	General Policy	Selective		
Long Name		Multiple Locations	Not Required		
Company Website	https://www.digitalocean.com	Ratio Requirement	No		
ASN	14061	Contract Requirement	Not Required		
IRR as-set/route-set	AS-14061	Contact Information			
		Role	Name	Phone	E-Mail
		Abuse	Abuse		abuse@digitalocean.com
		NOC	Network Operations		noc@digitalocean.com
		Policy	Peering		peering@digitalocean.com



Consistent WHOIS data through defined process



Accurate PeeringDB record



Monitored mailboxes

Action 4:

Global Validation

Facilitate validation of routing information on a global scale

We publish our routing data!

Given we allocate prefixes on a per-region basis, we need to ensure that the correct prefix lengths are kept up-to-date in our IRR objects. We use a scheduled “cron” job deployed to our internal application stack to ensure our IRR objects are accurate.



Automated IRR updates.



RPKI ROA coverage.



Automated alerting for non-compliance using NetBox reports.

Action 4: Global Validation

We publish our routing data!

We use covering ROAs with max-prefix-length populated to ensure we have valid ROAs for the prefixes we intend to advertise.

Routing completeness (IRR) ⁱ

<i>Unregistered</i>	0	0.0%
<i>Registered</i>	735	100.0%



■ Unregistered ■ Registered

Routing completeness (RPKI) ⁱ

<i>Valid</i>	724	98.5%
<i>Unknown</i>	11	1.5%
<i>Invalid</i>	0	0.0%



■ Valid ■ Unknown ■ Invalid



Automated IRR updates.



RPKI ROA coverage.



Automated alerting for non-compliance using NetBox reports.

Action 4: Global Validation

We publish our routing data!

NetBox reports are used to check for compliance and give us strong alerting signals when things aren't correct.

Report Results			
Time	Level	Object	Message
test_announce_roa			
2022-02-20T14:42:19.708340+00:00	Failure	69.55.48.0/24	ROA state for this prefix is not-found: No VRRP Covers the Route Prefix
test_aggregates_radb			
2022-02-20T14:42:00.073470+00:00	Info		Received 854 route objects from RADb
2022-02-20T14:42:00.155129+00:00	Success		all DigitalOcean aggregates have RADb objects



Automated IRR updates.



RPKI ROA coverage.



Automated alerting for non-compliance using NetBox reports.

Action 5:

Encourage Adoption

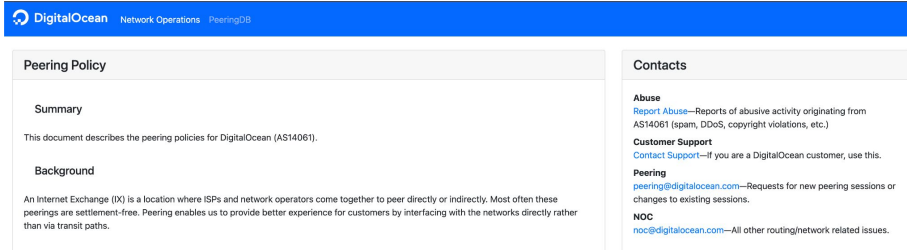
Encourage MANRS adoption

Action 5: Encourage Adoption

Simple as...

"...Peers are encouraged to implement Mutually Agreed Norms for Routing Security (MANRS) - <https://www.manrs.org>. "

<https://as14061.net/>



The screenshot shows the DigitalOcean PeeringDB website. The header includes the DigitalOcean logo and navigation links for Network Operations and PeeringDB. The main content is divided into two columns. The left column is titled 'Peering Policy' and contains a 'Summary' section with the text 'This document describes the peering policies for DigitalOcean (AS14061)' and a 'Background' section with the text 'An Internet Exchange (IX) is a location where ISPs and network operators come together to peer directly or indirectly. Most often these peerings are settlement-free. Peering enables us to provide better experience for customers by interfacing with the networks directly rather than via transit paths.' The right column is titled 'Contacts' and lists three categories: 'Abuse' with a link to 'Report Abuse', 'Customer Support' with a link to 'Contact Support', and 'Peering' with a link to 'peering@digitalocean.com'. There is also a 'NOC' link to 'noc@digitalocean.com'.



Updated our peering policy.



Encouraged adoption of MANRS.



Provided relevant links.

Compliance

December 2020

DigitalOcean Joins MANRS Initiative to Combat Routing Security Threats

Posted 2020-12-17 in [news](#)



By [Tim Raphael](#)

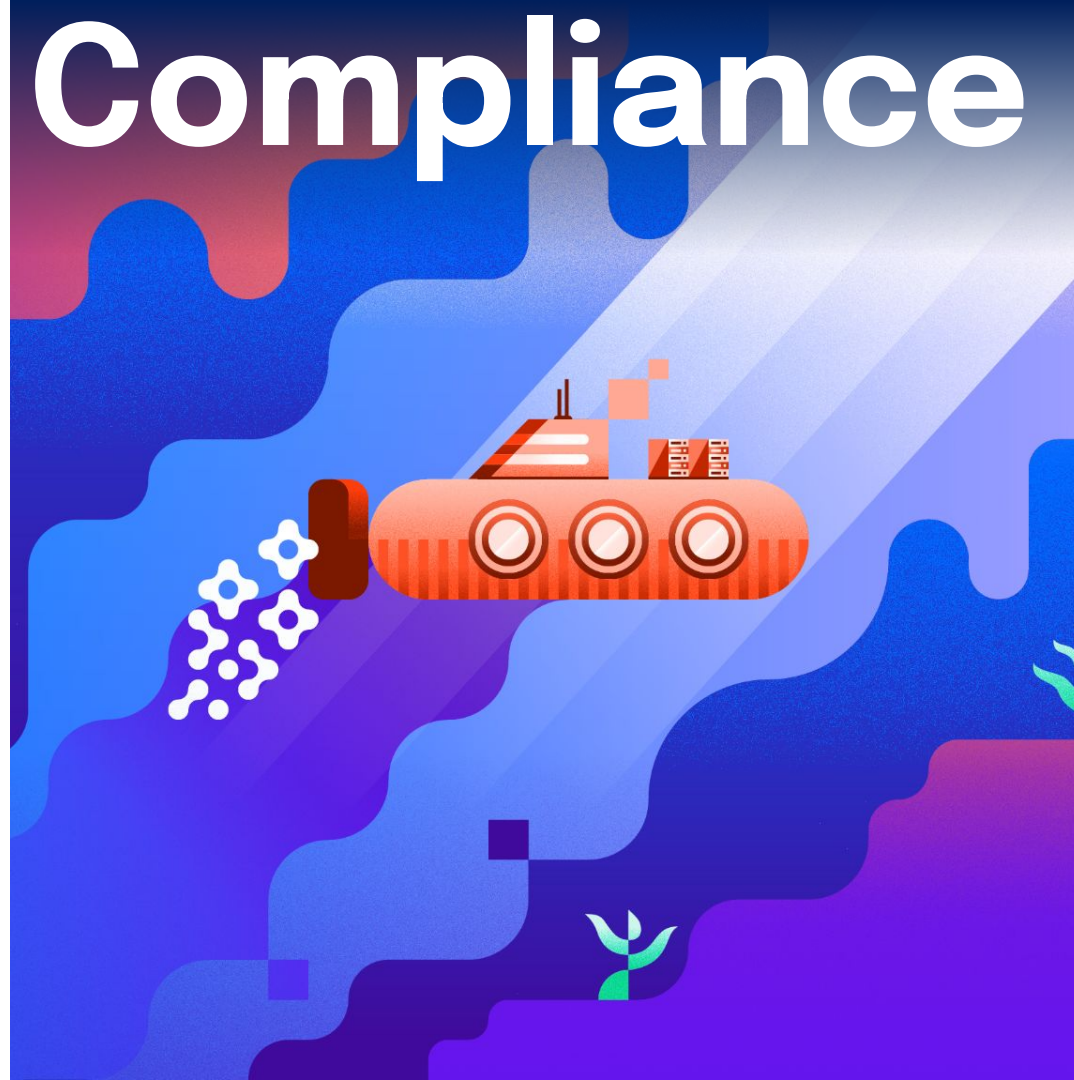
Today we are pleased to announce that DigitalOcean has joined the [Mutually Agreed Norms for Routing Security \(MANRS\) initiative for CDN and Cloud Providers](#) to reduce common routing security threats. The initiative, supported by the Internet Society, outlines actions network operators should take to improve the resilience and security of routing infrastructure.

<https://www.digitalocean.com/blog/digitalocean-joins-manrs>

DigitalOcean

May 2023

Compliance



What next?

Visibility is everything...

To help our peers we intend to launch an externally-facing looking glass that can help debug routing issues. To fit in with our other MANRS obligations we should ensure that RPKI status, route filtering status and various other aspects of routing policy are made clear with this tool.



Public Looking glass



Route filtering state



Route distribution policy

Improvement never ends...

While most of our processes are automated, there is always those few that aren't - we're continually aiming to improve automation coverage where it makes sense.

With the onset of a global network overhaul, new equipment gives us new capability to improve our filtering coverage.

Lastly, alerting and subsequent actions can always be improved as we experience new challenges and failure modes to learn from.



Increase automation coverage



Increase prefix-list coverage



Improve alerting

Thank you