MANRS+: Assessing the feasibility of routing security compliance tests

Date
19 August 2025

Authors
Lisa Bruder
Moritz Müller

Page
1/6

Classification
Public

Contact
lisa.bruder@sidn.nl
moritz.muller@sidn.nl

Contact
T +31 (0)26 352 5500
support@sidn.nl
www.sidn.nl

Offices
Meander 501
6825 MD Arnhem
The Netherlands

Mailing address
PO Box 5022
6802 EA Arnhem
The Netherlands

Abstract

Mutually Agreed Norms for Routing Security (MANRS) is an initiative that aims to improve routing security on the internet. They try to achieve this by defining actions that organisations can take to reduce the risk of threats to the routing system. MANRS participants commit to implement this list of actions. The correct implementation of MANRS Actions is continuously measured by the MANRS Observatory. However, some of the current metrics are insufficient to confirm the implementation with a high level of assurance. More specifically, the implementation of controls that prevent propagation of incorrect routing announcements is hard to measure. We at SIDN Labs collaborated with Global Cyber Alliance, responsible for the secretariat and operational functions of MANRS, to assess how the conformance to these requirements could be checked with a higher assurance level. To achieve this, we built a prototype that allows the simulation of measurements of these security controls. The prototype is publicly available on our GitHub.[1]

With the help of our local testbed, we could not identify major roadblocks for deploying the tests also in the real world. However, we found several points that operators of a future measurement setup should take into consideration, affecting the test design and the test infrastructure.

SIDN Labs is the research team of SIDN, the company that manages the Netherlands' Internet extension, .nl. SIDN Labs develops, prototypes and evaluates new technologies and systems that enhance the security and stability of .nl, the DNS and the wider Internet. Visit us at www.sidn.nl and www.sidnlabs.nl.

[1] https://github.com/SIDN/manrs-prototype

1 Introduction

Routing on the internet is facilitated by the Border Gateway Protocol (BGP), which allows Autonomous Systems (ASes) to interchange routing information. However, because BGP does not have a mechanism to validate said shared information, it is vulnerable to known attacks, like route hijacking.[2]

To make networks and with that the entire internet less vulnerable to these threats, MANRS[3] recommends actions ranging from filtering of BGP announcements to implementing and providing open-source monitoring tools to the community. The set of actions is determined by the MANRS community and forms a baseline for good practice in the internet routing system. To become a MANRS participant, organisations must agree to implement the required actions in their own network.

However, it is not always easy to audit the conformance to these actions reliably and there is not enough incentive to become a participant of MANRS for many

[2] https://www.rfc-editor.org/rfc/rfc4593.html
[3] https://manrs.org/

parties. Because of this, MANRS is working on MANRS+,[4] an elevated tier of MANRS for Connectivity Providers (CPs) who offer transit services to their customers. MANRS+ demands from CPs to adhere to more detailed requirements and aims to validate conformance to them with a high assurance level through measurements and audits. These requirements are separated into different "domains", like routing security, DDoS attack mitigation and anti-spoofing protection. The MANRS+ working group developed the MANRS+ Controls Matrix[5] to define what the requirements for different domains should be and how they should be validated (Self-declared, Audited or Measured).

In this report, we describe how we can measure whether future MANRS+ participants have implemented controls in the "Routing Security" domain. After providing some more background, we will describe our approach and our local testbed that we've specifically developed for this purpose. Finally, we will discuss aspects that operators of a future MANRS+ measurement platform might want to take into consideration.

2 Background

In this project, we focused on the measurement of the correct implementation of four controls in the routing security (RS) domain of the control matrix:

- Route Origin Validation (ROV) (control ID: RS-01)
- Prefix Filtering of Customers (control ID: RS-02)
- Control a set of customer ASes (control ID: RS-03)
- Filtering of bogons (control ID: RS-06)

Firstly, we measure the correct configuration of Route Origin Validation (ROV). The tested AS should discard announcements that are invalid based on Route Origin Authorisation (ROA) objects. ROA objects are stored in the Resource Public Key Infrastructure (RPKI), a decentralised database which has its trust rooted in the 5 Regional Internet Registries (RIRs). ROAs allow internet resource holders to cryptographically prove that they have the right to use and announce a range of IP addresses (a prefix). Once ROA objects are retrieved and validated, the information in them can be used to filter BGP announcements. If a ROA for a prefix exists and the stated origin AS differs from the one in an announcement, the route should not be considered in the best path selection process of the router.

Secondly, the tested AS should filter announcements using a prefix list for the customer (RS-02) and a list of customer ASes that are permitted to originate prefixes (RS-03). This filtering can happen based on information retrieved for example from the Internet Routing Registry (IRR). The IRR consists of several distributed routing policy databases that contain objects with information on contact persons, IP address space, autonomous systems and routes among other things. Similarly to the RPKI, this allows networks to publicly document important public information about their network as well as use the available data to filter announcements they receive from other ASes. In contrast to the RPKI, validity of data in IRR databases is not cryptographically secured.

Finally, we measure the correct filtering of bogon announcements (RS-06). A bogon announcement is an announcement of a prefix that is not intended to be routed on the public internet. Examples for this are the private IPv4 and IPv6 ranges. The tested AS should not announce any received bogon prefixes to its peers.

3 Approach

For this prototype we focus solely on the measuring of routing security controls, validating controls through self-declaring and auditing is out of scope. Based on discussions with Global Cyber Alliance, we decided on an approach where we measure configured filters of a CP AS by setting up two peering sessions with it. A CP AS has three main types of BGP relationships: one with its peers, one with its transit providers, and one with its customers. It provides connectivity to its customers by forwarding their (legitimate) routes to its peers and transit providers and vice-versa. To evaluate the correct implementation of the earlier mentioned controls, we focus on the relationships of the CP AS with its peers and customers because we expect that the CP will share all routes with its peers, guaranteeing that we are able

[4] https://manrs.org/about/manrs-working-group/
[5] https://manrs.org/wp-content/uploads/2025/02/MANRS-Controls-20250204.pdf

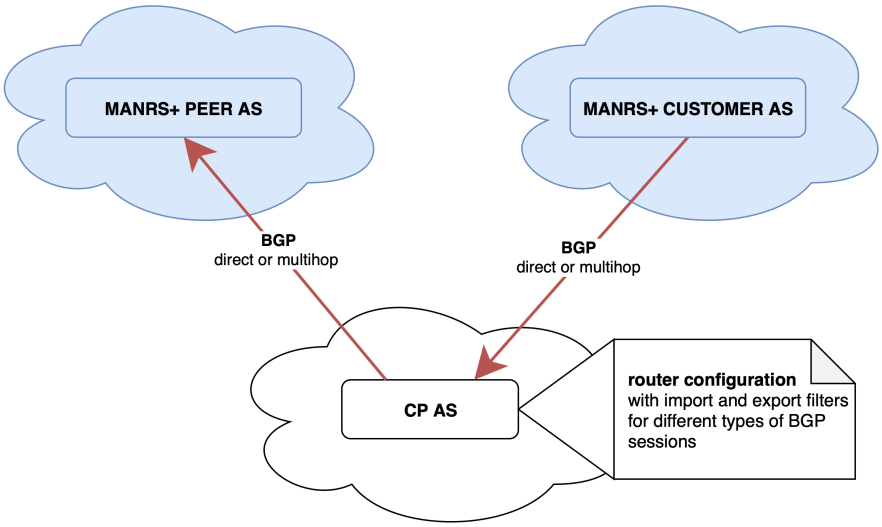

Figure 1: High-level overview of the eBGP peerings between the MANRS+ Peer AS, Customer AS and the CP AS.

to perform our tests.

By establishing two peering sessions with the CP AS – with the *MANRS+ Peer AS* representing the role of a peer AS and with the *MANRS+ Customer AS* representing the role of a customer AS – as shown in figure 1, we can observe how the CP AS processes announcements received from its customers. This can be done by announcing prefixes from the *MANRS+ Customer AS* and monitoring the announcements received by the *MANRS+ Peer AS*. If it receives routes that should be rejected according to the controls, this is an indication of improper filtering and configuration of the CP AS router.

4 Prototype implementation

To estimate the feasibility of this in a production setup, we set up a local prototype using containerlab.[6] Containerlab is a tool to orchestrate container-based networking labs. In contrast to tools like docker-compose, it allows easy configuration of links between containers and is therefore more suitable for our setup. In our prototype we use it to configure a topology that simulates the relationships between CP AS, *MANRS+ Customer AS* and *MANRS+ Peer AS*.

For the routers in our simulated ASes we use the BIRD[7] software router. It is well maintained and open source. This makes it accessible for everyone who wants to run the prototype. To validate and serve RPKI data to the CP AS router, we run a Routinator instance as a validating cache.

To simulate different types of production setups, we set up two topologies. One where the BGP sessions between the CP AS and the MANRS+ ASes are direct and one where they are configured as multi-hop BGP sessions. The latter is important since, on the internet, some tested CP ASes will not be directly connected to the MANRS+ ASes and remote peering over multi-hop BGP allows setting up sessions over intermediary hops. Next to the routers in our 3 local ASes, our topology also includes 2 supporting components: a local Internet Routing Registry (IRR) and a local RPKI. Running a local IRR and RPKI allows easy manipulation of data for demonstration of the different control test cases. We use Krill[8] and rsync to run a local publication server and IRRd,[9] Postgres and Redis to run a local IRR repository. To keep things simple, these are directly

[6] https://containerlab.dev/

[7] https://bird.network.cz/
[8] https://krill.docs.nlnetlabs.nl/en/stable/
[9] https://github.com/irrdnet/irrd/

connected to the CP AS router on a bridge and are therefore technically part of the CP AS. In a production scenario this is of course not the case.

To filter based on ROA objects, we configure the CP AS router to use Routinator[10] as a validating cache. Announcements where the origin AS and the prefix in the announcement match an existing ROA in the RPKI are marked as valid. If there is no ROA for the prefix in the RPKI, the announcement is marked as unknown and if there is a ROA connecting the prefix to a different AS, the announcement is marked invalid. We only accept routes with a status of valid or unknown in the path selection process. For filtering based on IRR data, we use bgpq4[11] to dynamically create prefix and AS lists based on route and AS-SET objects in our local IRR. The route object includes the prefix that the customer AS is authorised to announce. Only routes for prefixes in this list are accepted. The AS-SET contains the AS number of the customer AS and with that declares that it is allowed to originate prefixes. If an announcement is originated by an AS that is not in this list, the route is rejected. To validate proper filtering of bogon announcements, we create a static prefix list in the CP AS configuration with all bogon prefixes.

We automate the orchestration of the tests with bash scripts. If necessary for the test, objects are added to the RPKI or IRR. After that, filters in the configuration file of the CP AS router are either disabled or enabled depending on the test case. Finally, a prefix is announced from the *Customer AS* and the RIB of the *Peer AS* is monitored. Depending on the test case, the prefix should be announced or discarded at the CP AS and therefore should or should not show up in the RIB of the *Peer AS*. In the case of the control "Filtering of bogons" this means that there is one test case where the *Customer AS* announces a bogon and one where it announces a non-bogon prefix. The prefix should not be announced by the CP to the *Peer AS* in the first case and should be announced to it in the second. If the result is as expected, the test is marked as passed. Our repository on GitHub[12] includes scripts for all test cases and a more detailed documentation of them.

[10] https://github.com/NLnetLabs/routinator
[11] https://github.com/bgp/bgpq4
[12] https://github.com/SIDN/manrs-prototype

5 Considerations and requirements for production setup

From our simulations with the testbed, we could not identify any major problems that could hinder the measurement of MANRS+ compliancy using the proposed method. Nevertheless, when implementing the tests with future MANRS+ participants several aspects need to be taken into consideration that were not relevant for our testbed. We recommend a future operator of the production setup to consider the following points.

The considerations affect the design of the tests (considerations A.1 to A.6) and the test infrastructure (considerations B.1 to B.3).

5.1 Test design

Consideration A.1 – Test scheduling: Scheduling compliance tests in the real world with multiple CPs requires careful consideration.

- Before announcing the test prefixes towards the CPs, the information in the RPKI and IRR need to be updated accordingly. E.g. when announcing a prefix that should be discarded based on a ROV-invalid filter, the corresponding ROA needs to be updated (see *Consideration A.6*) and the announced prefix must be part of a valid AS-set in the IRR. Thereby, we can be sure that the CP dropped the prefix because of ROV and not because of invalid information in the IRR. The same goes for the other test cases: To validate prefix filtering based on information from the IRR, the prefix needs to have a status of ROV-unknown in the RPKI and must be part of a valid AS-set. To validate correct filtering based on an AS-set, the prefix must be either ROV-valid or ROV-unknown and valid according to information in the IRR.

- The test-operator should change the information in the RPKI and in the IRR long before the test prefix is announced towards the CPs. In the testbed, changes to the RPKI and IRR were propagated and reflected in the BGP control plane within seconds. On the internet, however, changes to the RPKI can

take an hour or longer[13] to propagate. Filters that are based on information in the IRR might propagate even slower (e.g. 24 hours in case of NTT[14]). When testing the controls in the real world, CPs should be given sufficient time to update their filters

- After the information in the RPKI and IRR has propagated, the test-operator can announce the prefix to the CPs. On the internet, BGP propagates updates usually in less than two minutes.[15] After two minutes, the operator should withdraw the test-prefix (see also *Consideration A.3*). Note that we expect that announcements between the MANRS+ ASes and the CPs will propagate faster due to their direct peering relationships. For this reason, the test prefix could be withdrawn even earlier, but real-world tests first need to validate this assumption.

Here, we expect that the test-operator performs tests one after another (e.g. first testing the propagation of a RPKI valid prefix, then a RPKI invalid prefix, then by announcing a prefix that belongs to the *Customer AS*, etc.). If multiple IPv4 and IPv6 prefixes are available, the test operator can also perform the tests in parallel (see also *Consideration B.1*).

Considerations A.2 - Testing multiple CPs in parallel: The operator of the *MANRS+ customer* can announce the same prefixes to multiple CPs at the same time. Then, the following aspects need to be taken into consideration:

- The test-operator should monitor incoming announcements at the *MANRS+ Peer* before the best path selection algorithm is applied (e.g. using the BGP Monitoring Protocol[16] (BMP)). This enables the operator to test the conformance of multiple CPs at the same time without the need for additional resources. See also considerations below. In contrast, in our testbed, we only monitor announcements by one CP at a time using the BIRD routing table.

- When monitoring incoming announcements at the *MANRS+ Peer*, the conformance test should only take announcements into account that follow the AS path "*AS-number CP AS; AS-number Customer AS"*. If another AS appears on the path, then we cannot determine the implementation of a control reliably.

Consideration A.3 – Limit impact of unintentional propagation: A failing test means that the CP has not implemented the appropriate controls. In order to prevent invalid announcements to propagate further, three countermeasures could be considered.

- First, the *MANRS+ Customer AS* should withdraw the invalid announcement as soon it has been observed at the *MANRS+ Peer AS*. In any case, announcements should be withdrawn after two minutes (see also *Consideration A.1*).

- Second, a CP might allow its customers to influence the propagation of their prefixes through BGP communities. In this case, the announcements of the *MANRS+ Customer AS* should contain BGP communities that instruct the CP to propagate the prefixes only towards the *MANRS+ Peer AS*.
- Third, in case a CP propagates test prefixes further, they may raise the awareness of other network operators and researchers. Therefore, MANRS+ should publicly document resources involved in the tests for transparency.

Note however, that even if an invalid announcement propagates to the wider internet its impact is limited. The invalid prefixes should only be used for testing purposes and no production traffic would be affected (see also *Consideration B.1*).

Consideration A.4 – Independent validation: Tests performed by MANRS+ are based on trust. A CP could treat announcements and peering relationships of the *MANRS+ Customer* and *Peer* differently (e.g. only applying MANRS+ filters on peering sessions with the *MANRS+ Peer* but not on sessions with regular peers

[13]https://link.springer.com/chapter/10.1007/978-3-031-28486-1_18
[14]https://www.gin.ntt.net/support-center/policies-procedures/routing-registry/
[15] https://ieeexplore.ieee.org/abstract/document/8861351
[16] https://datatracker.ietf.org/doc/html/rfc7854

and customers). Additional monitoring, similar to the measurements currently used for MANRS, could be deployed to monitor for the propagation for invalid routes independently.

Consideration A.5 – Filter direction: In our test setup, we validate whether a peer AS receives invalid announcements from the CP. Since we cannot assume that ROV is applied bi-directionally, additional tests could be implemented where the peer announces RPKI invalid announcements to the CP. If the *MANRS+ Customer AS* observes the invalid prefix, then the CP has not implemented the control.

Consideration A.6 – ROV test with AS0: Instead of relying on additional AS numbers for testing ROV, the operator of the test infrastructure could create ROAs using AS0. Thereby, the *MANRS+ Customer* can signal that the prefix should not be announced by any AS and thus, must be filtered by the CP.

5.2 Test infrastructure

Consideration B.1 – Required resources: We deployed our testbed locally, allowing us to use any IP addresses and AS numbers. When being deployed on the internet, the tests require at least the following resources:

- Three autonomous system numbers: One for the *MANRS+ Peer AS*, one for the *MANRS+ Customer AS* and one AS for the customer's customer.
- Routable IP addresses for the edge routers in the *MANRS+ Peer AS* and *MANRS+ Customer AS* to enable remote peering.
- At least two IP prefixes: One IPv4 prefix and one IPv6 prefix. While tests on IPv4 and IPv6 can run in parallel, tests of the different controls can only be performed one after another (see also *Consideration A.1*). The tests can run in parallel if more prefixes are available. IP prefixes of the size of a /24 (IPv4) and /48 (IPv6) are sufficient. To test attacks/leaks where a customer announces a more specific prefix, larger prefixes are required.
- Well-connected location to announce the ASes and prefixes: Ideally, the MANRS+ networks are located at well-connected locations, e.g. popular internet exchange points. In case a CP is not co-located with the MANRS+ networks it can utilize remote peering, see also *Consideration B.2*.

The used resources should only be used for testing the compliancy with MANRS+ controls.

Consideration B.2 – Multi-hop: In case a CP cannot directly peer with the MANRS+ networks, it might need to establish a multi-hop peering session with both of them. From our tests, we conclude that multi-hop peering sessions do not influence the tests.

Consideration B.3 – Unsolicited traffic: As soon as operators announce resources on the internet, they can expect traffic coming towards these resources (e.g. due to scanning). The operators of the MANRS+ networks should monitor traffic coming especially towards the *MANRS+ Customer* since it passes through the network of the CP. The goal is to ensure that the network of the CP is not utilized unnecessarily. Note, that the amount of unsolicited traffic is limited when the *Customer* withdraws the prefix right after the *Peer* has observed the update, or not later than after 2 minutes (see also *Consideration A.3*).

6 Conclusions and Future Work

We've tested the validation of four MANRS+ controls in our local testbed and could not identify major roadblocks for deploying the tests also in the real world. However, we found several points that operators of a future measurement setup should take into consideration. To validate our findings, the next obvious step is to carefully test the measurements on the internet.

Acknowledgements

We thank the GCA staff for providing their feedback on the report and the testbed.