Internet Routing with MANRS

MANRS

By Fred Baker

Executive Summary

The Internet as we know it is approximately 35 years old. The Border Gateway Protocol (BGP), the primary backbone routing protocol, was designed before we knew much about security; we have been changing and fixing it as our knowledge has increased. Improvements include the establishment of specific data to help operators identify what routing data is valid and what is not, as well as specific operational practices to address known attacks. To date, the biggest problem with those capabilities is that many operators remain either unaware or unconvinced that their participation is needed if matters are to improve.

Charles Dudley Warner famously said, “Everybody complains about the weather, but nobody does anything about it.” In the summer of 2014, a group of companies made the same observation about Internet routing—and then did something about it. The result is the Mutually Agreed Norms for Routing Security,[1] or MANRS, a global initiative designed to collaboratively provide crucial fixes to reduce the most common routing threats. And the acronym is no accident, the authors are making a point: it is good etiquette, good *manners*, to say trustworthy things when speaking to one’s neighbour. MANRS actions result in trustworthy Internet routing, a reasonable basis for a business.

At the writing of this report, more than 100 Internet service providers (ISPs)[2] and Internet exchange points (IXPs)[3], comprising hundreds of Autonomous Systems (ASs) in many countries, have agreed to take the following four actions:

1. Filtering route origins
2. Anti-spoofing of source addresses in Internet traffic
3. Coordination of actions
4. Global validation of routing announcements

Individually, the four steps are quite straightforward. While they require some effort, that effort is neither difficult nor expensive for most implementations. The issues those steps address, however, are costly from both insurance and public relations perspectives. For example, if a company’s traffic is misrouted to someone who harvests access credentials and uses them to hack the company or its customers, hundreds of millions of dollars in damage could be accrued. What’s more, the company misrouting the traffic could be found culpable. More than simply good route hygiene or cheap insurance, these recommended actions might be all that stands between your network and the financial and public relations nightmare of a security breach.

Governments are taking notice of routing issues. For example, in the United States, the National Institute of Standards and Technology (NIST) and the Department of Homeland Security (DHS) recently reinforced the importance of the MANRS project by publishing a

draft US standard[4,5] along the same lines. In addition to Resource Public Key Infrastructure (RPKI)[6] and route origin validation (ROV)[7], the NIST–DHS standard calls for path validation via Border Gateway Protocol Security (BGPsec), which is a possible future building block for routing security. US Government networks and companies that contract with governments in the United States should anticipate fulfilling requirements similar to those outlined in the MANRS agreement.

This paper explores some of the issues surrounding routing security and provides examples of both implementation approaches and where those approaches have been used to successfully prevent or mitigate attacks. It includes the following sections:

- Section 2: Introduction. A description of the issues around Internet routing and why we need to address the security of it.
- Section 3: The Four MANRS Actions. An outline of the four MANRS actions and how they may be cost-efficiently and effectively carried out.
- Section 4: Conclusion. Possible next steps.
- References. For those interested in more detail, this paper references a number of reports and relevant online commentary.

Introduction

In the public square, an attack is any event in which one party carries out forceful or aggressive opposition to another. We know what an attack is in war or between students on a playground. In the Internet, attacks are carried out with other weapons—often to embarrass someone, commit theft, or to undermine the effectiveness of a service.

As many as 14,000 Internet routing incidents happen annually,[8,9] and nearly 30,000 denial-of-service (DoS) attacks of various kinds occur daily.[10] One recent report describes a two-hour hijack in which data from the encrypted messenger app Telegram was redirected through networks in Iran.[11,12,13]

Customers trust that their ISPs and IXPs will connect them to those entities with whom they want to communicate. Routing incidents, such as accepting or propagating a false prefix, are a fundamental service failure in that they connect their customers to someone else.

Internet networks have the ability to defend their prefixes—they can document and claim the prefixes allocated to them for inclusion in other networks' filters, as well as build and apply such filters. With this in mind, one could argue that a network that announces a prefix and doesn't defend it or that receives and uses or propagates a prefix without verifying it is an accessory to the attack.

On the surface, propagating or receiving a prefix without verifying it or failing to defend a prefix they have announced might not appear to affect an ISP or its customers. But each is, in fact, both directly and indirectly affected. For example, networks that announce false prefixes frequently find themselves carrying far more traffic than they are capable of handling. In June 2015, a major route leak by Telekom Malaysia overloaded its infrastructure with users and resulted in a massive slowdown on its networks. Because all Internet networks are connected by routing, those networks that propagated prefixes from Telekom Malaysia during the route leak also contributed to the "significant packet loss and Internet slowdown in all parts of the world".[14]

Routing incidents have real-world costs—from failing to provide services to one's customers to the costs of mitigation and response. During the Route 53 attack of April 2018,[15] an ISP announced a hijacked prefix to Equinix that then passed the prefix to a variety of other networks. Those networks that filtered their routing did not propagate the prefix further, and their customers remained unharmed. Those networks that did propagate the hijacked prefix subjected their customers to incorrect connections and monetary losses.

The Route 53 attack and countless others illustrate that actions performed to protect the Internet also protect the networks that implement the actions by limiting the risk of costly routing incidents and improving the network operators' bottom lines.

MANRS offers a set of simple actions designed to help network operators improve the security of the global routing system. Supported by the Internet Society, MANRS is a global initiative comprising like-minded network operators, IXP operators, and enterprises with the common goal of preventing route hijacking and certain kinds of DoS attacks via tools developed by the Internet Engineering Task Force (IETF) and the greater operational community.

The Four MANRS Actions

Network operator members of MANRS agree to perform the following four actions designed to improve routing security.

- **Filtering**: to ensure the correctness of their own announcements and of those from their customers to adjacent networks with prefix and AS-path granularity.
- **Anti-spoofing**: to enable source address validation for at least single-homed stub customer networks, their end-users, and infrastructure.
- **Coordination**: to maintain globally accessible, up-to-date contact information.
- **Global validation**: to publish their data, so others may validate routing information on a global scale.

Most networks find these actions simple to implement. CloudFlare's deployment[16] demonstrates that, even for larger networks, the actions are feasible.

Filtering

Users of BGP[17] routing can't believe what they are told unless it is validated. This is why MANRS' first action is filtering—and why it is a crucial first step.

BGP's primary mechanisms include the announcement of prefixes and the filtering of received announcements; a secondary mechanism addresses route selection. During route selection, a router announces to its neighbor that it knows how to reach a set of prefixes with a certain set of attributes, including an Autonomous System (AS) path, a list of the autonomous systems that an announcement passes through in order to reach the network in question. The first network in an AS path is the one originating the prefixes. Prefixes are allocated to an originating AS by an administration—a regional Internet registry (RIR), a nonprofit corporation that administers and registers *Internet* Protocol (IP) address space and AS numbers within a defined *region,* or the party to which the RIR allocated the prefix. Prefixes are announced by the party to which they were allocated. Hence, a second AS can determine if the AS that originated the announcement was, in fact, the AS authorized to do so—this is the essence of ROV. In some cases, a prefix is a more-specific version of an announced prefix, which also makes it valid.

When prefixes are not correctly documented or not correctly announced

Example 1. BGPmon[18] recently reported a situation in which Hyatt Corporation was allocated a prefix, and Hyatt Germany, which has a different AS number, announced a subset of the prefix as its own. Most would agree that this is a matter of etiquette; the prefix was allocated to a company, and a wholly-owned subsidiary used it. It might not be legally wrong, but, technically it's not right either. Most important, other networks cannot verify that it is correct. If a non-Hyatt organisation behaved in this way, it would be undeniably both technically and legally wrong.

Example 2. BitCanal[19,20] is a serial hijacker. They reportedly coopt address space allocated to other parties without the parties' permission, and then sell the use of that address space to others, thereby denying its use to those for whom it has been allocated. In many cases, the upstream networks to whom BitCanal connects are unaware of BitCanal's actions, and they disconnect BitCanal as a customer when they learn the truth. But it appears BitCanal is persistent—when one upstream network disconnects them, they find another. Beyond the legal and organisational ramifications, these kind of routing abuses propagate an unreliable Internet in which users cannot trust that the party with whom they are conversing is the one they think it is.

BitCanal is neither the first nor will they be the last company to hijack prefixes. Reports of rerouting so-called big-name traffic to ASs in Russia[21], China[22], Belarus[23], and other countries can be found by a simple web search. Although surprisingly common, they are egregious abuses.

Securing the Internet from companies like BitCanal raises the questions, what is the provable identity of a routing peer? and what is it authorized to announce? The MANRS-compliant operator polices its own announcements to ensure that it is advertising only its own prefixes and those for whom it is contracted.

Example 3. In 2012, Dodo, an Australian ISP, had two links to the Telstra network: a primary link and a secondary link. When the primary link failed, Dodo announced all of the routes it had learned from Telstra on the primary link back to Telstra on the secondary link, which marked all the routes as customer routes.[24] For 30 minutes, approximately 1,400 prefixes—10% of all Australian IPv4 prefixes—were unable to communicate internationally. If either Dodo or Telstra had implemented a correct filter (e.g., "is it my route?" or "do I already have a preferred route?"), the leak would have been avoided—either Dodo would not have announced the routes or Telstra would not have accepted them.

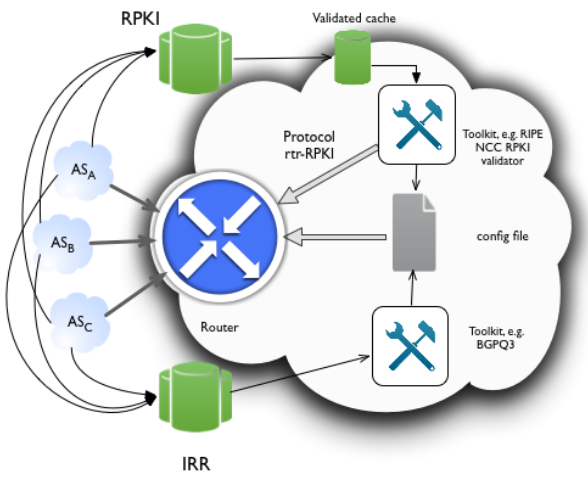

Figure 1 The ISP filtering toolkit

For most companies, filtering is as simple as converting the Internet Routing Registries (IRRs) of its customers and, in some cases, their customers, to filters on the customers' edge routers[25] or route arbiters. The filters verify announcements against published ROAs with an RPKI service or compare them to published lists of prefixes allocated by the RIRs (figure 1). An unvalidated prefix is a liability that could cost a company more than the cost of validating its routes. Most companies that exchange BGP routes with

customers or peers can require their customers to maintain IRRs or ROA data. The IRRs would then be downloaded, converted to filters, and applied in a background task; ROAs would be downloaded, verified, and stored in a server that could be queried to validate announcements. The phrase, *a small matter of programming,* is an old joke in network security circles, but in this case, keeping companies and their customers secure is just that: a small matter of programming.

Anti-spoofing

Data attacks fall into several categories, including hacking, phishing/spearphishing, and DoS/Distributed DoS (when aided by a botnet).[26] While many ISPs block spoofed source addresses, many still don't[27]—70% of networks surveyed for a 2017 Association for Computing Machinery (ACM) report continue to experience these attacks.[28] According to a second 2017 ACM report, DoS attacks alone number more than 28,700 a day.[29] Help Net Security, in independent site focusing on information security, reports that the most frequent attacks are User Datagram Protocol (UDP) floods (59.7%), Transmission Control Protocol (TCP) SYN floods (3.3%), and Internet Control Message Protocol (ICMP) floods (0.9%); and that high volume attacks are ramped up via Memcached reflection, Simple Service Discovery Protocol (SSDP) reflection, and Connection-less Lightweight Directory Access Protocol (CLDAP).[30] Far from inevitable, an Akamai threat advisory[31] notes that CLDAP attacks, as well as many other reflection and amplification attack vectors, "would not be possible if proper ingress filtering was in place".

Reflection attacks, such as memcached reflection, SSDP reflection, and CLDAP reflection attacks, are DDoS attacks in which a source sends a request using a spoofed source address in order to trick a service into responding legitimately to an unsuspecting party.[32] In these cases, both the service's and the unsuspecting party's ability to perform their intended functions is reduced from the attack on their available bandwidth or computational capabilities. In addition, trust in the service itself is also compromised.

MANRS' second action is designed to decrease spoof attacks by preventing hosts from sending packets with spoofed source IP addresses.

The most general solution to spoofing is the one recommended by both MANRS and Akamai: the ISP or other network filters arriving subscriber IP traffic for rational source addresses (figure 2). BCP 38 Ingress Filtering[33] is a basic mechanism developed by the IETF to prevent attacks that target a BGP session using a spoofed source address. In short, the service provider of a given subscriber ensures that its subscriber only uses addresses that have been allocated to it. In addition to simple access control lists (ACLs), BCP 84[34] describes multiple implementations of a technology for ingress/egress filtering, called Unicast Reverse Path Forwarding (uRPF). The premise in uRPF is that it is reasonable to receive a packet from a direction, if one might also send a packet to that address in that direction.

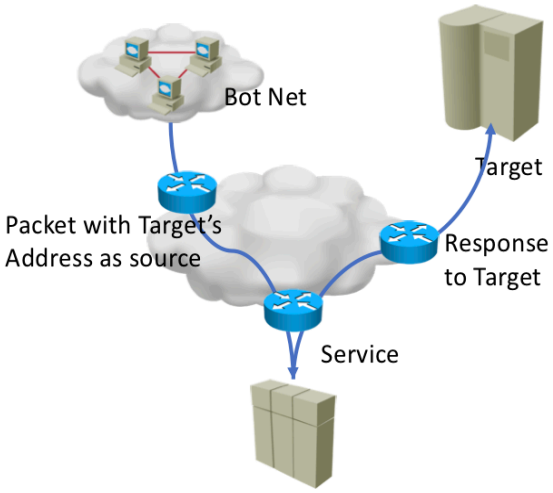

Figure 2 Diagram of a Reflection Attack

Loose and strict source address filtering

Originally proposed by Cisco, the strict or loose uRPF[35] model uses a forwarding information base (FIB) that discards packets received on an interface other than the next hop for transmission to the address. This filter relies on the router to choose the next hop interface to forward to the address and requires that the prefix be in the FIB. It is only applicable to symmetrical routing.

Feasible Path uRPF

Juniper's model, feasible path uRPF[36], builds on the BGP Routing Information Base (RIB). If a prefix is announced to a router on a given link, whether or not routing chooses to use the BGP neighbor as its next hop, the hosts that use source addresses in that prefix may, if the router has a prefix in its RIB, originate traffic that uses that link. This filter allows for asymmetric routing.

Enhanced Feasible Path uRPF

A problem exists, however, with uRPF in strict, loose, and feasible path variants. Imagine that a network is disaggregating its routes—traffic to its prefix arrives, but the prefix is subdivided and advertised in pieces to several upstream networks. There is nothing to stop the network from sending traffic from its entire prefix to each upstream and receiving traffic to the various disaggregated prefixes via those networks. The effect of feasible path uRPF is to enforce route symmetry; any given upstream will block traffic from subprefixes advertised to other upstreams.

Enhanced feasible path uRPF[37] relies on the prefix being associated with the source ASN using IRRs or RPKI. In this scenario, traffic from the prefix in the IRR or RPKI—the entire allocated prefix—is acceptable as a source prefix to all upstreams, even if the prefixes themselves are not advertised upstream.

Using uRPF in real configurations

ISPs that don't filter frequently comment about the cost of filtering.[38] But the truth is that uRPF, which does a route lookup on the source address, costs no more than a route lookup on the destination address—it is simply a different route table. And configuring uRPF is simple: for Cisco equipment, the *ip verify unicast reverse-path*[39] interface configuration command accomplishes it, as does the *set from route-filter*[40] command on Juniper equipment. For interfaces that face simple subscribers, it just makes sense to include the configuration in the default customer configuration.

Coordination

The simplest and most common way of addressing problems in Internet routing involves the telephone: the operator who notices the issue calls someone he or she thinks is able to fix it. Often, that someone's telephone number is already in the operator's contact list as either a personal or corporate relationship. Operators commonly coordinate their responses to issues using the mailing lists of their Network Operator Groups. For example, during the Route 53 Attack, Nordu.Net[41] commented on an attack on MyEtherWallet[42], looked through its own logs, and asked others to do the same.

In general, the exact means of communication is up to the network operator or company. MANRS-compliant operators are available for other operators to contact them, when appropriate. For the majority of errors, this is usually sufficient.

Global Validation

While MANRS has a limited scope—it requires validation of routing announcements only for an operator's own networks and customers—it contributes to the ability of relying parties to validate by asking each member to maintain its routing information in relevant public databases. The biggest obstacle to validating any announcement—be it from a customer, a peer, or a transit provider—is incomplete or inaccurate routing data. MANRS solves this problem by clearly identifying what prefixes are originated by which network.

When errors are more complicated than average, finding a solution requires an efficient and reliable means of validating routing information. Because BGP, unlike BGPsec, does not directly require route, route origin, or path validation, route validation must be done out of band. The MANRS Best Current Operational Practice (BCOP)[43] details the information that must be available for validation and the means for doing so, including use of three essential tools: IRRs, RPKI, and proprietary databases.

History proves that global validation helps. During the Route 53 attack in 2018[44] some ISPs propagated the route, others did not. The ISPS that filtered the rogue announcement performed a valuable service for themselves and their customers. Unfortunately, Hurricane Electric was one of the 12 out of 15 AS paths in the hijacked route that failed to perform that service. As a result, malicious advertisements were shared with many of their worldwide BGP peers, causing a massive rerouting of traffic towards a malicious DNS server.

Internet Routing Registries (IRRs)

IRRs are public databases that contain routing information, such as network routing policies and expected routing announcements, in the form of objects represented by the Routing Policy Specification Language (RPSL)[45]. In concept, anyone can maintain such a registry, and several do. For example, almost every RIR maintains a registry for its members, and the Merit Corporation maintains an aggregate IRR database derived from these RIR databases. Tools[46,47] exist to use registries to add defensive capabilities to routing configurations; if a BGP peer announces a prefix that it shouldn't, filters derived from the IRR database filter the incorrect announcements, thereby protecting both the network that implements the filters and anyone they would pass the routes to. The concept is that each AS files an RPSL description of its list of prefixes, and its peers can obtain that from trusted registries.

In the aforementioned case of Hyatt Corporation and Hyatt Germany (page 4), if Hyatt has an IRR entry for its prefix, it might add the AS numbers of each of its subsidiaries that are using them. When Hyatt Germany announces the prefix, the AS detecting a fault updates its own IRR database and the change becomes acceptable.

IRRs also have disadvantages. They are incomplete and not all of the data in them is reliable. Perhaps most important, they work best when a filtering AS is near the network edge, particularly when serving customers who, in turn, provide no transit service to others. If an AS is further from the network edge, the IRR needs to do more work—it will receive prefixes from customers whose AS paths end in the AS number of the networks that are downstream. If an algorithmic analysis of the IRR's RIB indicates that it could route to an AS downstream, it will add that AS's IRR to its list of filters and manage routing upstream. The enemy of that algorithm is scale, however. At some point it becomes too complex to efficiently manage or apply.

As noted, IRRs help validate routing by associating prefixes with prefix lengths (and highlighting prefixes that are too long) and originating AS numbers. With enhanced uRPF,

they can also help with BCP 38 filtering—the ISP now knows what prefixes are associated with ASs downstream.

IRRs are mostly applicable to implementation of the first action of MANRS: protecting customer cones and their own networks. A more scalable solution can be found in origin validation, a more-recently developed technology based on RPKI.

Resource Public Key Infrastructure[48]

In routing, an RPKI infrastructure is another way of securing prefix exchanges. By securing the channel used for exchanging information (e.g., TCP MD5, TCP Authentication, or IPsec) by mutual authentication or encryption, the data is secured. In DNS, DNSSEC signs (and, therefore, secures) the resource records that are exchanged; it doesn't matter how the records were obtained. DNSSEC is an example of securing DNS data. Unlike BGPSEC, in which the originating AS cryptographically signs the announcements it sends, in the case of ROV RPKI is used to cryptographically sign records similar to IRR entries, so that anyone who receives them can see that the routing policy of the originating AS allows it to be announced. An RPKI infrastructure is far more scalable than an IRR, but it comes with costs, including taking the actual validation offline.

As in an IRR entry, an RPKI Route Origin Authorization[49] (ROA) details the prefix, its acceptable set of lengths (indicated as its minimum and maximum length), and the set of AS numbers from which it might be announced. Upon validation, a ROA can determine the following about an announcement:

- If it matches ROA's prefix, prefix length, and originating AS – it is valid
- If it matches the prefix but not the prefix length, or the originating AS - it is invalid
- If no ROA is found for a prefix - its validity is unknown

NIST has closely studied RPKI deployment and its effectiveness.[50] The web page of the laboratory's Advanced Network Technologies Division offers two key results of its observations: as of this writing, while most routes are not found and those that are found are valid, approximately 10% of prefixes have ROA descriptions that to some companies announce routes that RPKI rejects as invalid. If the companies that use RPKI to validate routes wanted, they could protect themselves and their customers from questionable routes by rejecting the routes RPKI found invalid. Similarly, if ASs wanted, they could create ROAs with their favorite RIR and make the technology more effective.

As noted, two responsibilities must be fulfilled: the data must be available, which falls on every party to whom a prefix has been allocated, and the data must be used to validate, and occasionally reject, announcements. The company with the unprotected asset and the AS that falls for the hijack are both accessories to a crime. Validation requires taking action, and databases and open-source software exist to do it—the MANRS actions are accessible to most, if not all, companies.

Proprietary databases

Some ISPs require customers to enter the equivalent of IRR information into proprietary databases, and then block from them all routing and traffic that doesn't match that information. Because of the nature of closed proprietary sources, these databases do not facilitate origin validation on a global scale. A more effective solution would be for the ISP to ask customers to point them to their published IRR or ROA.

If the use of a proprietary database is a requirement, such companies can import IRR or RPKI data and use their databases in a more secure fashion. For this reason, companies seeking a wider range of marketing choices will often turn to public registries.

Conclusion

Many say that networking is unnecessarily complex and costly, not because the infrastructure cannot be protected, but because it isn't protected. We, as an industry, are capable of verifiably identifying the prefixes we use and preventing them from being misused—either in routing or as a source address when accessing a service. And it is our responsibility to do so.

Fulfilling that responsibility starts with the following next steps:

1. There are no disadvantages to a network filing IRR and RPKI data to protect its assets. RIRs should require members to file IRR and RPKI data when possible, and companies advertising prefixes in BGP should register valid IRR registries and RPKI ROAs.
2. For networks in which deployment is simple (e.g., networks that serve SOHO customers or enterprise customers), enforcement of IRRs and RPKI in routing and BCP 38 in data transmission is a low-cost way for a network to protect both itself and the networks that connect to it. They should deploy uRPF for each downstream customer and validate the routes that customers advertise to it.
3. Open- and closed-source development communities should clean routes via tools that are inexpensive to operate and that integrate RPKI and IRR technology with the data in each RIR's delegated-<registry>-latest allocation files.

As the examples in this paper and innumerable others every day illustrate, it is significantly easier and less costly to implement MANRS than to lose your customers' data and your network's good name to a security breach.

References

1 https://www.manrs.org/
2 https://www.manrs.org/participants/
3 https://www.manrs.org/participants/ixps/
4 https://www.zdnet.com/article/standard-to-protect-against-bgp-hijack-attacks-gets-first-official-draft/
5 https://www.nccoe.nist.gov/sites/default/files/library/sp1800/sidr-piir-nist-sp1800-14-draft.pdf
6 Resource Public Key Infrastructure, also known as Resource Certification, is a specialized public key infrastructure framework designed to secure the Internet's routing infrastructure.
7 Route Original Validation describes route filtering in order to ensure that the routes received match RPKI-certified specifications.
8 https://blog.apnic.net/2018/06/22/worldwide-survey-engages-network-operators-on-bgp-hijacking/
9 https://blog.apnic.net/2018/01/24/14000-incidents-routing-security-2017/
10 https://www.securityweek.com/internet-sees-nearly-30000-distinct-dos-attacks-each-day-study
11 https://www.cyberscoop.com/telegram-iran-bgp-hijacking/
12 https://www.theregister.co.uk/2018/08/01/bgp_route_leak_telegram_iran/
13 https://spacewatch.global/2018/08/iran-hijacks-global-telegram-app-traffic-prior-to-national-protest/
14 https://bgpmon.net/massive-route-leak-cause-internet-slowdown/
15 https://www.internetsociety.org/blog/2018/04/amazons-route-53-bgp-hijack/
16 https://blog.cloudflare.com/rpki/
17 Border Gateway Protocol is a routing protocol used to direct traffic across the Internet. Use of BGP enables autonomous systems to build routes according to business policies or agreements.
18 BGP Monitoring System (BGPmon) gathers BGP data from connected peers and provides real-time routing information. https://www.bgpmon.io/
19 https://www.bleepingcomputer.com/news/security/internet-transit-providers-disconnect-infamous-bgp-hijack-factory/
20 https://blogs.oracle.com/internetintelligence/shutting-down-the-bgp-hijack-factory
21 https://arstechnica.com/information-technology/2017/12/suspicious-event-routes-traffic-for-big-name-sites-through-russia/
22 https://www.washingtontimes.com/news/2010/nov/15/internet-traffic-was-routed-via-chinese-servers/
23 https://www.wired.com/2013/12/bgp-hijacking-belarus-iceland/
24 https://bgpmon.net/how-the-internet-in-australia-went-down-under/
25 https://ripe76.ripe.net/presentations/43-RIPE76_IRR101_Job_Snijders.pdf
26 https://www.cisecurity.org/wp-content/uploads/2017/03/Guide-to-DDoS-Attacks-November-2017.pdf
27 https://spoofer.caida.org/summary.php
28 https://conferences.sigcomm.org/imc/2017/papers/imc17-final24.pdf
29 https://www.securityweek.com/internet-sees-nearly-30000-distinct-dos-attacks-each-day-study
30 https://www.helpnetsecurity.com/2018/08/15/ddos-attacks-outside-business-hours/
31 https://www.akamai.com/us/en/about/our-thinking/threat-advisories/connection-less-lightweight-directory-access-protocol-reflection-ddos-threat-advisory.jsp
32 https://resources.arbornetworks.com/wp-content/uploads/NTPReflection_Final.pdf
33 https://tools.ietf.org/html/rfc2827
34 https://tools.ietf.org/html/rfc3704
35 https://www.cisco.com/c/en/us/about/security-center/unicast-reverse-path-forwarding.html
36 https://www.juniper.net/documentation/en_US/junos/topics/concept/unicast-rpf-understanding.html
37 https://datatracker.ietf.org/doc/draft-ietf-opsec-urpf-improvements
38 https://conferences.sigcomm.org/imc/2017/papers/imc17-final24.pdf

39 https://www.cisco.com/c/en/us/td/docs/ios/12_2/security/configuration/guide/fsecur_c/scfrpf.html
40 https://www.juniper.net/documentation/en_US/junos/topics/topic-map/unicast-rpf.html
41 https://mailman.nanog.org/pipermail/nanog/2018-April/095105.html
42 https://techcrunch.com/2018/04/24/myetherwallet-hit-by-dns-attack/
43 https://www.manrs.org/guide/global-validation/
44 https://www.internetsociety.org/blog/2018/04/amazons-route-53-bgp-hijack/
45 https://tools.ietf.org/html/rfc2622
46 https://sourceforge.net/projects/irrpt/
47 https://peering.readthedocs.io/en/latest/PrefixLists.html
48 https://www.apnic.net/get-ip/faqs/rpki/
49 https://www.ripe.net/manage-ips-and-asns/resource-management/certification/resource-certification-roa-management
50 https://rpki-monitor.antd.nist.gov/