Quickscan on routing measures to increase the security of the Internet

October 27, 2017

Ad Bresser

Ad.Bresser@gmail.com

+31 6 20 39 56 98

Abstract

Internet infrastructure (IP/BGP) routing is the crucial basis for all the Internet based services. This Quickscan lists all identified measures to improve the security and reliability of the Internet routing. The measures vary from BCP/RFC documented measures to architectural approaches and administration and assurance measures. The total number of measures is pretty large. To the routing community it is suggested to explore whether there are alternative approaches for solving some of the structural weaknesses.

Introduction

Internet and Cybersecurity get a lot of attention; especially the high visibility incidents and weaknesses. In order to create insight in the measures to improve the security of the Internet routing, 12 people were interviewed. These insights were enriched with desk research results and some explored architectural concepts. The created insights don't include any exploration of current weaknesses and are definitely not a risk assessment of any kind.

In order to facilitate the growth and correctness of this overview, a version is created on Github:

https://github.com/AdBresser/measures-to-increase-IP-security

Content:

Objective and scope

The objective of this document is to create insight in the measures taken to increase the security and stability of the Internet routing system. There is already a lot of information created and organized (like [MANRS]) around this subject. In this document it is structured in an accessible way and some architectural and business operations measures are added.

The focus is on possible measures in the IP and BGP routing domain. There is no guarantee that this overview is complete. It doesn't take into account the number of deployments and is most certainly not a risk assessment.

IP engineering measures

One of the major weaknesses of the IP protocol, is the possibility to spoof the IP source address of an IP packet. According to [CAIDA-Spoof], in around a third of the announced address space, prefixes and autonomous systems, the source address is spoofable.

Source IP address assurance

By setting up routers in networks properly, DDoS attacks based on spoofed IP source addresses, can be prevented. This way, spoofed packets are easily stopped. These packets are the source of many different forms of attacks. Unfortunately, filtering this is only effective if almost all networks set it up [CSAN2016].

The nearer the filters are applied to the origination of the spoofed traffic, the better the effects on the security and reliability of the hosts and the network will be. According to some measurements, up to 60% of the IP addresses used in attacks are bogon.

Ingress filtering - Customers

On the edge of a network, every source IP address from every incoming IP packet can be validated. If the IP source address is not the expected IP address or doesn't belong to the known prefix, the packet should be dropped. This is clearly described in [BCP38]. This can be applied in access networks, transit networks and server networks. The amount of deployments is unknown. The benefits are mainly for the Internet community as a whole. More specific description on how to prevent nodes attached to the same IP link from spoofing each other's IP addresses, can be found in [SAVI].

Ingress filtering - Transit & Peering

The IP source addresses of IP packets received from connectivity providers, can be validated. If the IP address is from your own prefixes, or from a private address range, the packet should be dropped. When the source IP address in the packet belongs to a customer, it can also be dropped, unless it is the IP address of a multi-homed customer. The amount of deployments in unknown. This filter is beneficial for your own infrastructure and your customers.

Egress filtering - Transit & Peering

Source IP addresses of IP packets send to connectivity providers, can be validated. They should only contain IP addresses from prefixes you or your customers own. Although this seems redundant compared to the Ingress filtering - Customers, it prevents the leakage of IP packets with source

addresses with from a private range, which is sometimes used for internal purposes. Moreover, it prevents IP packets with incorrect source IP addresses from wrongly configured or hacked servers you have in your network. The amount of deployments is unknown. this filter is beneficial for the Internet community and maintaining a profile of a credible network provider.

Egress filtering - Customers

The IP source addresses of IP packets send to customers can be validated. Those IP source addresses should not be from a private address range. It is an option to check for the IP range from the customer, but special care should be taken, because it could be valid traffic from multihomed customers or test traffic. The amount of implementation is unknown, but not a lot of implementations are expected. Benefit is for the customer, see also [BCP46]

Advanced cases and solutions

[BCP84] is an update of [BCP38] and gives some advanced options for implementing source address assurance, including the handling of multihomed situations.

In [RIPE431] examples are given on how to configure routers for filtering.

The filtering options described can best be applied on the external interface, but in some cases it is more logical to perform the filtering on a other interface. This is especially the case for egress filtering.

IPv6 considerations

IPv6 makes it somewhat easier to implement end2end security, because it supports end-to-end encryption.

As IPv6 is not yet widely deployed, it is likely that there are more (compared to IPv4) security weaknesses in implementations.

The available IPv6 address space will create opportunities to optimize security. E.g. in quarantine networks, were only the affected system can be blocked based on the IP address, instead of the blocking all systems that are behind one IPv4 address.

IPv6 will remove the need for IPv4 address space driven NAT solutions. NAT rewrites the source address of normal traffic to ensures reachability and is thus a natural source address assurance mechanism (which makes spoofing more complicated). IPv6 removes the need for NAT and could increase the spoofing opportunities, when [SAVI] is not implemented.

Although security by obscurity is never a sustainable security mechanism, the amount of IP addresses creates possibilities for this. This is only mentioned here for completeness and not advised as a security mechanism.

BGP engineering measures

The Border Gateway Protocol (BGP) is the protocol almost exclusively used in the Internet to exchange routing information between network domains. Due to this central nature, it is important to understand the security measures that can and should be deployed to prevent accidental or intentional routing disturbances.

Good risk analysis of BGP vulnerabilities can be found in [CSRIC3BGP] and [NIST800-54].

General BGP Operations and Security

General BGP Operations and Security is described in [BCP194]. It covers the following topics:

- Protection of the BGP Speaker
- Protection of BGP Sessions on TCP level and by TTL Security (GTSM)
- BGP Route Flap Dampening
- Maximum Prefixes on a Peering
- AS Path Filtering
- Next-Hop Filtering
- BGP Community Scrubbing
- Prefix Filtering:
  - Special-Purpose Prefixes
  - Unallocated Prefixes, based on IANA and RIR-Allocated Prefix Filters
  - Filters Created from Internet Routing Registries (IRRs)
  - Filters based on SIDR
  - Secure Inter-Domain Routing principles (BGP Prefix Origin Validation and BGPsec)
  - Prefixes That Are Too Specific
  - Filtering Prefixes Belonging to the Local AS and Downstreams
  - IXP LAN Prefixes for Network Security, PMTUD and the Loose uRPF Problem
  - The Default Route
  - In Full Routing Networks with Internet Peers, Customers and Upstream Providers
  - Prefix Filtering for Leaf Networks

Advanced BGP management

In [BGPJob] there are some BGP measures mentioned that are especially relevant for the more specialized network operators (there is some overlap with [BCP194]):

- Peer Locking aka "bignetworks filter"
- Dropping Bogon ASNs

Ingress:

- Dynamic maximum prefix settings
- Reject Bogon prefixes (RFC1918, etc)
- Reject Bogon ASNs (AS0 / AS23456 etc)
- Reject IXP prefixes (Some IXP subnets)
- Reject leakage with the Peerlock filter
- Match against IRR whitelist (only customers)
- Mark as customer route (or as peer route)
- Scrub internally significant BGP communities
- Apply Features – (blackholing, traffic engineering, etc, only for customers)

Egress:

- Reject Bogon prefixes
- Remove-private-AS
- Reject “bad” routes
- Accept peer routes (on customer session)
- Accept customer routes (on every session)
- Do prepending (if requested & applicable)
- Scrub internal communities
- Set next-hop-self
- Normalize Med

BGP origin validation

There is a form of attack against the routing information provided by BGP4. This is done through injection of misleading routing information into the routing system. The measures to protect against misleading routing information consists of several Resource Public Key Infrastructure (RPKI) based components (next to the already mentioned). The operation best practices are described in [BCP185].

Architecture concepts

There are several architecture concepts that can increase the security and reliability of the routing system.

Split destinations

It is possible to increase the security (e.g. mitigating DDOS attacks) by creating multiple instances of a destination for IP packets and manipulating the flow of IP traffic coming from the Internet. Such a solution can include separated servers, dual ASNs and separated peering policies. With the latter, peerings with trusted networks can be separated from worldwide connectivity. The advantage of this solution is that it is completely independent from other cooperation initiatives.

Transit & Peering

The continuity of Internet connectivity can be increased by using multiple Transit providers, connecting to multiple IXPs and implementing that on multiple locations.

With respect to the peering policy: a more open policy increases the robustness, but a full open policy creates the risk of connecting with unknown parties.

Selective peering

It is possible to selectively peer with networks that match a certain quality / security profile. Profiles can be based on region, black- and whitelist, [MANRS], information from RIPE DB and Peering DB. This typically is enabled / eased on route servers of IXP's.

Trusted networks

It is possible to build a solution consisting of a network (VLAN) of trusted networks, that can be activated by each participating network in case of an emergency situation. This will result in a temporary disconnect from the global internet, while ensuring connectivity to trusted networks. This can be completely separated from regular peering policies.

In the Netherlands the Trusted Network Initiative [TNI] was launched to counter the adverse effects of DDoS attacks [CSAN2016].

Direct connections

One of the most trustworthy solutions is of course the direct connection between networks.

IP measures for increased security

Port filtering

One of the most basic measures to increase the security is to perform filtering based on the TCP/UDP ports. This can be used to protect servers from being accessed on ports without any public service, but also in access networks to protect customers against attacks on ports that were not intended for global purposes. In access networks there is sometimes also a port 25 block to prevent spam by infected end-user devices. In these cases, the customer email is to be sent over ISP SMTP servers. This prevents the ISP from being blacklisted and thus completely loosing email connectivity.

Quarantine IP addresses

An IP address (typically associated with an user) can (temporally) be blocked from access to the majority of the Internet. This can be useful when the IP address is used for unwanted behaviour, like being a (potential) source of a DDOS. This can be the case when the control of a PC or server is taken over by an infection. Once this infection is removed, the IP address can be granted full access again. Such measures have proved to be very successful in not only reducing the unwanted behaviour, but also in the distribution of malware. The identification the relevant IP addresses, is of course crucial. In a case of an access network, this was done by forcing all email over an SMTP server (port 25 block) and profiling. In an access network were this was implemented, the zombie rate was significant lower.

Similar concepts also work in hosting environments to prevent / mitigate undesired traffic. It is not uncommon that a hosting customer that triggers abuse messages (e.g. by sending spam), receives a port 25 block until the issue is resolved.

In residential environments, typically only one IP address is issued per physical connection. More devices can use the connection by implementing concepts like NAT. When such a IP address is placed in quarantine, this means that all devices loose connectivity, when only one device is infected. This increases the awareness of the end-user (TV video streaming stops, when the adolescent laptop is infected), but can have unwanted side effects without being noticed (IoT device stops).

Shared DDOS scrubbing center

To mitigate the impact of DDOS attacks there are DDOS scrubbing solutions. They can be on-site or outsourced, but they can also be shared. The Dutch national anti-DDOS scrubbing center [NaWas], is a solution build on top of IXPs to make it accessible for all participants. It can be activated within seconds by BGP controlled rerouting the DDOS traffic to the scrubbing center and receiving the cleaned traffic.

Administration

IP address administration

Register your assigned IP range in the RIR, with up-to-date contact information, including abuse contact information. And regularly review if the information is still up to date.

ICANN - Identifier Technology Health Indicators (ITHI)

As described in [ITHI] ICANN is working on an initiative to improve the quality of (amongst other) IP address registration, by measuring the quality.

Register your network

As information is relevant it is a good idea to register your network in the [PeeringDB], as it facilitates the exchange of information related to Peering. It is a database of networks that are peering, where they are peering, and if they are likely to peer.

Register your routing policy

It is a very good idea to register you routing policy at [RADB] or your local RIR. There are some authorisation checks possible when describing your routing policy in your local RIR [RipeDbIRR]. Some connectivity providers (peering or transit) require you to do so. Next to publishing your own routing intentions, it supports the construction and maintaining of routing filters and router configurations and it assures diagnostic and information service for network management. Routing policies are described by using the Routing Policy Specification Language (RPSL) conventions. A good tutorial on using RPSL can be found in [RFC2650].

Assurance

Operational Security Requirements

In [RFC3871] a list of operational security requirements for the infrastructure of large Internet Service Provider (ISP) IP networks (routers and switches) is specified. It provides network operators a clear, concise way of communicating their security requirements to vendors. It covers functional requirements, documentation requirements and assurance requirements.

Functional Requirements:

- Device Management Requirements
- In-Band Management Requirements
- Out-of-Band (OoB) Management Requirements
- Configuration and Management Interface Requirements
- IP Stack Requirements
- Rate Limiting Requirements
- Basic Filtering Capabilities
- Packet Filtering Criteria
- Packet Filtering Counter Requirements
- Other Packet Filtering Requirements
- Event Logging Requirements
- Authentication, Authorization, and Accounting (AAA) Requirements
- Layer 2 Devices Must Meet Higher Layer Requirements
- Security Features Must Not Cause Operational Problems
- Security Features Should Have Minimal Performance Impact

Documentation Requirements:

- Identify Services That May Be Listening
- Document Service Defaults
- Document Service Activation Process
- Document Command Line Interface
- 'Console' Default Communication Profile Documented

Assurance Requirements:

- Identify Origin of IP Stack
- Identify Origin of Operating System

For detailed recommendation on how to protect the router's control plane, see [RFC6192].

Security audit

It is a good practice to conduct technical and procedural security audits. Those audits can be held with fixed time intervals, after a (major) change and after an incident.

Security certification

[ISO27000] is the standard that helps organizations keep information assets secure. Although this standard is not specific enough for networking equipment, implementation and certification eases audits (at least the procedural part of it).

Contingency planning

Because of the packet switching nature of IP, it has natural robustness. Still it can be very useful to regularly perform a contingency audit. An approach for such an analysis, is exploring several "what if" situations to identify areas of improvement. This is especially useful when the networks grows incremental.

CERT / CSIRT

A computer emergency response team (CERT) is an expert group that handles computer security incidents. Alternative names for such groups include computer emergency readiness team and computer security incident response team (CSIRT).

The CERTs are organized in [FIRST] (global Forum of Incident Response and Security Teams)

Security Operations Center

A Security Operations Center (SOC) is responsible for the detection and investigation of vulnerabilities in the operational infrastructure, the interpretation of cyber threats and the advising of countermeasures to remove existing risks. During disasters, the SOC acts as the Computer Emergency Response Team (CERT) [CSAN2016].

bgp.he.net

The Hurricane Electric BGP Toolkit [BGPHE], gives insight in, amongst others, actual routing, prefix and relations between networks. This can support quality and consistency checking.

National contingency boards

Internet service providers in the Netherlands have established the Dutch Continuity Board (DCB), which works on measures to minimise the impact of DDoS attacks on Dutch critical infrastructure and to make services that have been disrupted available again as soon as possible [CSAN2016].

National Detection Network

In The Netherlands investments have been made in the National Detection Network [NDN]. This is a partnership of the NCSC and other parties in the exchange of information on detected threats [CSAN2016].

Conclusions & recommendations

Protocols upon which the internet functions were originally designed to transport data as efficiently as possible and without too much attention to security [CSAN2017].

There are a lot of options available to increase the security of the Internet on routing level, but the number of implementations is not known.

Although the oldest Best Current Practice is from 2000 ([BCP38] - Network Ingress Filtering), still around one third of the source addresses is spoofable. This raises the question whether the current approaches for improving the security are effective, or that the approach should be sought on another level.

There are three major routing security challenges, that need attention: Source addresses, route injection and contact information (see [MANRS]). The issue with these vulnerabilities is that the cost of implementing the solutions by an IP operator doesn't benefit the operator directly; the benefits are for the other IP operators.

From discussions and observations, it has become clear that in general IP operators are not driving standardization and implementation of the security solutions. There is a group of good willing and acting, but there is no shared vision on how persuade the unwilling and laggards in implementing security measures.

Recommendations:

1. Assess risks of the known security vulnerabilities for the stability of the Internet, because the likelihood of exploits and impact of is not documented and communicated. (well, at least not found during this quickscan).
2. Explore number of good implementations. Current insights are based on self-assessments [MANRS] or external measurements [CAIDA-Spoof]. Good results of the latter could be based on non-security driven technology choices; like IPv4 NAT.
3. Create better understanding of the background of the problem of limited implementation of security measures and investigate the possible directions of solutions. IP operator organization and driving standardization and implementation, could be one of the solutions.

Annex: References

[BCP38] - Network Ingress Filtering (aka RFC2827)
Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing
https://tools.ietf.org/html/bcp38
Publication: May 2000

[BCP46] - Recommended Internet Service Provider Security Services and Procedures (aka RFC3013)
http://www.rfc-editor.org/bcp/bcp46.txt
Publication: November 2000

[BCP84] - Ingress Filtering for Multihomed Networks (aka RFC3704)
https://tools.ietf.org/html/bcp84
Publication: March 2004

[BCP185] - Origin Validation Operation, Based on the Resource Public Key Infrastructure (RPKI) (aka RFC7115)
https://www.rfc-editor.org/bcp/bcp185.txt
Publication: January 2014

[BCP194] - BGP Operations and Security (aka RFC7454)
https://tools.ietf.org/html/rfc7454
Publication: February 2015

[BGPHE] - BGP Toolkit
https://bgp.he.net/

[BGPJob] - Practical everyday BGP filtering with AS_PATH filters
Job Snijders - Peerlocking - NANOG67
https://www.nanog.org/sites/default/files/Snijders_Everyday_Practical_Bgp.pdf

[CAIDA-Spoof] - State of IP Spoofing
CAIDA - Center for Applied Internet Data Analysis - State of IP Spoofing
https://spoofer.caida.org/summary.php

[CSAN2016] - Cyber Security Assessment Netherlands 2016
Cyber Security Assessment Netherlands 2016: Professional criminals are an ever greater danger to digital security in the Netherlands
https://www.ncsc.nl/english/current-topics/Cyber+Security+Assessment+Netherlands/cyber-security-assessment-netherlands-2016.html
Publication: October 24, 2016

[CSAN2017] - Cyber Security Assessment Netherlands 2017
Cyber Security Assessment Netherlands 2017: Digital resilience is lagging behind the increasing threat
https://www.ncsc.nl/english/current-topics/Cyber+Security+Assessment+Netherlands/cyber-security-assessment-netherlands-2017.html
Publication: August 31, 2017

[CSRIC3BGP] - BGP Security Best Practices, FCC CSRIC III WG4 Final Report
http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG4_Report_March_%202013.pdf
Publication: March 2013

[ITHI] - ICANN Identifier Technology Health Indicators project
https://www.icann.org/ithi

[ISO27000] - ISO/IEC 27000 family - Information security management systems
https://www.iso.org/isoiec-27001-information-security.html

[FIRST] - Forum of Incident Response and Security Teams
https://www.first.org/

[MANRS] - Mutually Agreed Norms for Routing Security (aka Routing Manifesto)
https://www.routingmanifesto.org/
[NaWas] - Nationale anti-DDoS Wasstraat
https://nbip.nl/web/guest/nawas
https://slideshare.net/splend/hsb-nationale-anti-ddos-wasstraat-alex-bik
[NDN] - National Detection Network
https://www.ncsc.nl/english/Cooperation/national-detection-network.html
[NIST800-54] - Border Gateway Protocol Security - NIST Special Publication SP 800-54
http://csrc.nist.gov/publications/nistpubs/800-54/SP800-54.pdf
Publication: July 2007
[PeeringDB] - Peering DB
https://www.peeringdb.com
[RADB] - Routing Assets Database
http://www.radb.net/
[RFC3871] - Operational Security Requirements for Large Internet Service Provider (ISP) IP Network Infrastructure
https://tools.ietf.org/html/rfc3871
Publication: September 2004
[RFC2650] - Using RPSL in Practice
https://tools.ietf.org/html/rfc2650
Publication: August 1999
[RFC6192] - Protecting the Router Control Plane
https://tools.ietf.org/html/rfc6192
Publication: March 2011
[RIPE431] - RIPE Anti-Spoofing Task Force HOW-TO
http://www.ripe.net/ripe/docs/ripe-431
Publication: May 9, 2008
[RipeDbIRR] - Using the RIPE Database as an Internet Routing Registry
https://labs.ripe.net/Members/denis/using-the-ripe-database-as-an-internet-routing-registry
Publication: August 22, 2013
[SAC004] - Securing the Edge, Paul Vixie, ISC. Not referenced
http://www.icann.org/committees/security/sac004.txt
Publication: October 17, 2002
[SAVI] - Source Address Validation Improvement Framework (aka RFC7039)
https://tools.ietf.org/html/rfc7039
Publication: October 2013
[TNI] - Trusted Networks Initiative
https://tn-init.nl/
https://www.thehaguesecuritydelta.com/images/TNI_Info_Sheet_01-04-2015.pdf
http://procon.bg/article/trusted-networks-initiative-netherlands-response-ddos-attacks
https://www.thehaguesecuritydelta.com/projects/project/60-trusted-networks-initiative