Report on

Secure routing technologies

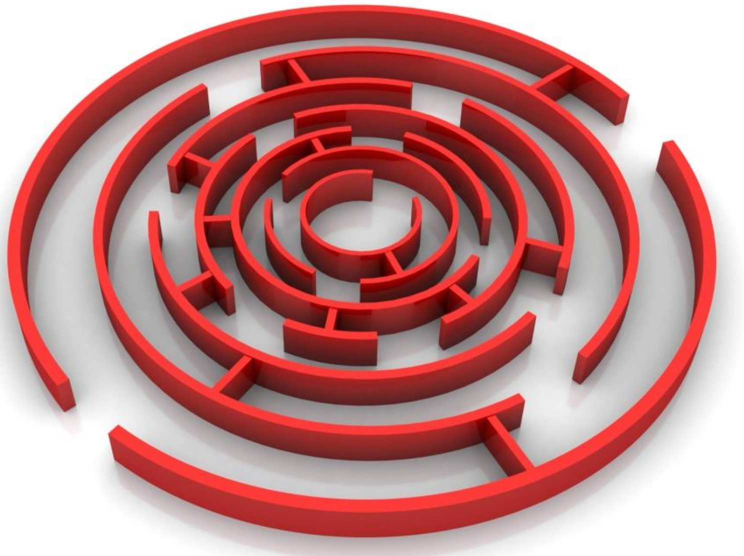

enisa
European Network and Information Security Agency

About ENISA: The European Network and Information Security Agency (ENISA) is an EU agency created to advance the functioning of the internal market. ENISA is a centre of excellence for the European Member States and European institutions in network and information security, giving advice and recommendations and acting as a switchboard for information on good practices. Moreover, the agency facilitates contacts between European institutions, the Member States, and private business and industry actors. Internet: *http://www.enisa.europa.eu/*

Contact details:

To contact ENISA or for enquiries regarding this study, please communicate with:

Technical Department, Security Tools and Architectures Section
Email: sta@enisa.europa.eu
Web: http://www.enisa.europa.eu/act/res/technologies/tech/routing/

This report is the result of the collaboration between the European Network and Information Security Agency (ENISA) and the Institute of Communications and Computer Systems (ICCS) http://www.iccs.gr/eng/.

Legal notice: Notice must be taken that this publication represents the views and interpretations of the editors, unless stated otherwise. This publication should not be construed to be an action of ENISA or the ENISA bodies unless adopted pursuant to ENISA Regulation (EC) No 460/2004. This publication does not necessarily represent the state-of the-art in secure routing technologies and it may be updated from time to time. Third-party sources are quoted as appropriate. ENISA is not responsible for the content of the external sources including external websites referenced in this publication. This publication is intended for educational and information purposes only. Neither ENISA nor any person acting on its behalf is responsible for the use that might be made of the information contained in this publication. Reproduction is authorised provided the source is acknowledged.

Table of Contents

Executive Summary

Reliable communications networks and services are now critical to public welfare and economic stability. Intentional attacks on the Internet, disruptions due to physical phenomena, software and hardware failures, and human mistakes all affect the proper functioning of public communications networks. The European Commission's communications have already highlighted the importance of network and information security and resilience. In this context, ENISA has approached the subject of securing the routing infrastructure by delivering two reports: An assessment report on the 'impact of deploying secure routing technologies on the network operators in the EU' and this report on 'secure routing technologies'.

The report addresses the issue of vulnerabilities in the routing protocols and related threats; attack objectives, mechanisms and the extent of their effects; mitigation measures; the operation of proposed secure routing protocols and the threats they are addressing; and the hurdles hindering their deployment. It also provides recommendations on the deployment of secure routing technologies.

Interdomain routing is mainly dictated by the Border Gateway Protocol (BGP), which was introduced in 1995. The channel over which BGP messages are exchanged between peers has to be secured to mitigate DoS and integrity attacks.

Adjacent peers should enable TCP-AO and the TTL security mechanisms.

BGP has not yet provided any protocol mechanism for the verification of resources (IP address prefixes, topology, etc) announced by entities (companies or ISPs). It is not aware of the existence of organizations and their respective address assignments. As a result, routing security incidents occur when an entity announces addresses without authorisation. This is called prefix hijacking (PH). Partial remedies to address PH are *filtering* and *route prefix monitoring*.

ISPs should protect customers with filtering 'in the import BGP' configuration clauses and the implementation of a prefix hijacking alarm solution around the Internet eXchange Points.

Additional mechanisms, like S-BGP and soBGP, for securing BGP are presented in this report. The solution for securing the interdomain routing is the Resource Public Key Infrastructure (RPKI) which is a mixture of the other mechanisms presented in the report and a PKI infrastructure. An RPKI roadmap should be fostered among equipment vendors, registries and ISPs for the deployment of the technology. Policy makers should address the governance issue of the authoritative trust anchors.

Introduction

Reliable communications networks and services are now critical to public welfare and economic stability. Intentional attacks on the Internet, disruptions due to physical phenomena, software and hardware failures and human mistakes all affect the proper functioning of public communications networks. Such disruptions reveal the increased dependency of our society on these networks and their services. Experience shows that neither individual providers nor a single country alone can effectively detect, prevent, and respond to these threats.

Communications from the European Commission[1][2][3] have already highlighted the importance of network and information security and resilience for the creation of a single European information space that will drive job creation, sustainability and social inclusion, and so contribute to the overall goals of the Europe 2020 strategy[4]. They have stressed the importance of dialogue, partnership and the empowerment of all stakeholders to properly address these threats. The updated Regulatory Framework Directives[5][6] include certain regulatory provisions for the improvement of the security and resilience of public eCommunications.

[1] *Communication from the Commission to the Council, the European Parliament, the European Economic and Social Committee and the Committee of the Regions: i2010 – A European Information Society for growth and employment /* COM/2005/0229 final */* http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2005:0229:FIN:EN:PDF

[2] *Communication from the Commission to the Council, the European Parliament, the European Economic and Social Committee and the Committee of the Regions: A strategy for a Secure Information Society – Dialogue, partnership and empowerment /* COM(2006) 251 */* http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2006:0251:FIN:EN:PDF

[3] *Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions: A Digital Agenda for Europe /* COM/2010/0245 final */* http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:52010DC0245:EN:NOT

[4] http://ec.europa.eu/eu2020/index_en.htm

[5] *Directive 2009/140/EC of the European Parliament and of the Council of 25 November 2009 amending Directives 2002/21/EC on a common regulatory framework for electronic communications networks and services, 2002/19/EC on access to, and interconnection of, electronic communications networks and associated facilities, and 2002/20/EC on the authorisation of electronic communications networks and services* http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:337:0037:0069:EN:PDF

[6] *Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009 amending Directive 2002/22/EC on universal service and users' rights relating to electronic communications networks and services, Directive 2002/58/EC concerning the processing of personal data and the protection of privacy in the electronic communications sector and Regulation (EC) No 2006/2004 on cooperation between national authorities responsible for the enforcement of consumer protection laws* http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:337:0011:0036:EN:PDF

Since 2008 ENISA has been executing a programme with the ultimate objective of collectively evaluating and improving the resilience of public eCommunications in Europe. The programme is comprised of four distinct phases.

The first step undertaken was an analysis of how national authorities implement current regulatory measures. This involved assessing how network and service providers of public communication networks ensure the availability and integrity of their networks and services, and evaluating whether existing technologies satisfy the needs and requirements of these providers. In this light an assessment of three key technologies (namely IP version 6, Multiprotocol Label Switching and DNS Security Extensions) was carried out regarding their potential to provide increased network resilience[7].

This analysis was carried out from two perspectives. The first consisted of analysing the characteristics of the selected technologies and their public communication network's resilience enhancing features[8]. In parallel, the effectiveness of these technologies as well as the problems and gaps that potentially could compromise the availability of networks and services was assessed through interviews with twelve network operators in the EU Member States[9].

Routing infrastructure is a critical infrastructure that needs to be addressed in order to secure public communication networks. Improving the resilience of a network is an issue of risk management which includes risk identification, evaluation and acceptance or mitigation.

In that context, ENISA has approached the subject of securing the routing infrastructure with two reports: an assessment report on the impact of deploying secure routing technologies on the network operators in the EU, and a report on the available technologies. In this respect this report is playing the role of support/background information.

In the following chapters the report on 'secure routing technologies' is presented. This report focuses on interdomain routing and addresses the following issues:

- vulnerabilities of the routing protocols and threats
- attack objectives, mechanism and extent of their effect
- existing mitigation measures
- secure routing protocols and their operation
- the threats they are addressing

[7] http://www.enisa.europa.eu/act/it/inf/tech

[8] http://www.enisa.europa.eu/act/it/library/deliverables/res-feat/at_download/fullReport

[9] http://www.enisa.europa.eu/act/it/library/deliverables/stock-tech-res/at_download/fullReport

- problems hindering their deployment.

Also, recommendations are provided on the deployment of secure routing technologies. The recommendations can be used by policy making bodies to provide directions or deliver advice to the industry and to address governance issue.

Interdomain routing

Routing in the Internet defines the path that IP packets take to get from their source to their destination. In general terms, this path, called a *route*, is a chain of routers and the links between them. Routers use a routing protocol to compute such paths. Routing protocols perform computations on data, called *reachability* information, to compute the desired routes.

In the earlier days of the Internet, routing was much simpler that it is today. At that time, the requirements were fairly simply, ie, shortest path routing. Routing was performed by a single administrative entity (NSFNET). Over time the Internet changed and became heavily commercialized by the internet service providers (ISPs), who had an interest in controlling their traffic for financial and political reasons. So the Border Gateway Protocol (BGP) [1][2] was introduced to accommodate routing control between ISPs.

In terms of BGP, ISPs are described as single entities, called *domains,* which have full control of their routing decisions. Routing inside a domain is dictated by an *intradomain* routing protocol. Routing between domains is mainly independent of the intradomain operation. BGP operates between domains by communicating route exchanges and, for that reason, it is commonly referred to as the *interdomain* routing protocol of the Internet.

After its initial deployment, subsequent incremental modifications of BGP defined further details of the protocol, such as the finite state machine, protocol messages and additional capabilities for route aggregation and classless interdomain routing (CIDR). The last change was implemented in order to cope with the routing table expansion problem. The current version of BGP, version 4, has been deployed extensively for over two decades. This study will address the security shortcomings of BGP as the *de-facto* interdomain routing protocol of the Internet.

Address management

A routing domain, equivalently referred as an *Autonomous System* (AS), represented by a unique numeric ID [3], announces destination IP address ranges (often in a prefix/notation form) to one or more neighbouring ASs. For instance, prefix 155.207.0.0/16 represents a 2^16 address block, belonging to AS 5470 (Aristotle University of Thessaloniki). ASs advertise the set of prefixes that they originate (ie, the addresses within their administrative domain).

While many organizations maintain their own AS, many others do not, and still others (typically connectivity providers) may maintain more than one. Each organization assigns its address space to the AS in which the addresses reside. Hence, *assignment* is the process whereby an organization gives an AS *the right to originate* a set of addresses. These addresses are configured into routers which subsequently advertise them via BGP.

The address space of the Internet is managed by the Internet Assigned Numbers Authority (IANA) [5]. IANA *delegates* large address blocks to the regional Internet registries (RIRs) [7][8] (ie, ARIN, APNIC, and RIPE NCC) which, in turn, allocate smaller blocks to local Internet registries (LIRs) or large ISPs. LIRs (which are typically ISPs or collections of ISPs represented at a country level) and large ISPs process the vast majority of address space assignments to ISPs and end-users.

Delegation and assignment on the Internet is currently an administrative process. There is no structure for validating claims to address ownership and assignment. BGP is not aware of the existence of organizations and their respective address assignments.

BGP operation

Since its debut, BGP has always been a *path vector* protocol. The term 'path vector' originates from the fact that BGP lists a sequence of AS numbers (commonly referred as the AS PATH) along the path over which an IP prefix has traversed. In Figure 1, AS5 U, the owner of 20.20.0.0/16, announces this prefix to its upstream providers. An AS, receiving this path, may choose to propagate it to some or all of its neighbours. An AS intending to propagate a received path, prepends its own AS to the AS path before announcing this to its neighbours. Thus in Figure 1, AS A has a path of (A,P,Q,R,U) to reach AS U. When multiple address prefixes overlap, the longest prefix match rule is used to break the tie. Thus, if a BGP routing table contains a path to reach prefix 20.20.0.0/24 as well as 20.20.0.0/16, a packet destined to 20.20.0.128 would choose the path of entry 20.20.0.0/24.

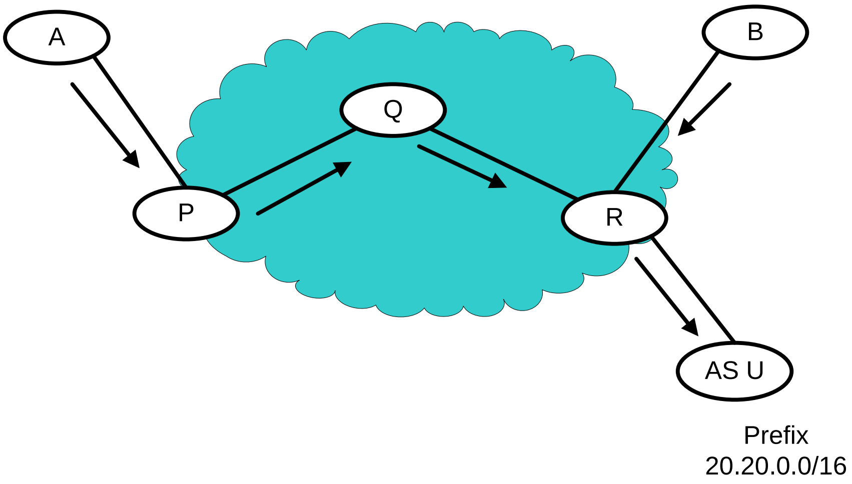

Figure 1: Valid origin AS U announces prefix 20.20.0.0/16

BGP is still (even after the 4th version) a relatively simple protocol in terms of message exchanges. Each AS announces prefixes that it has learned from its neighbours to other neighbours. BGP is an incremental protocol (commonly referred as a *delta protocol*) where, after a complete routing table is exchanged between neighbours, only changes to that information are communicated. Those changes may be new route advertisements, route withdrawals, or changes to route attributes.

The overall announcement procedure does not lead to overall flooding due to the *path vector* characteristics of the protocol which detects and removes loops (multiple occurrences or announcements by the same AS). The whole process of announcements stems from the egalitarian origin of BGP. Each AS is trusted *a priori* to behave in accordance with its specifications and to provide accurate routing information. In other words, BGP does not provide security.

Unfortunately, BGP's lack of security has allowed serious routing instabilities. One such failure occurred in February 2008 when the Government of Pakistan and its Pakistan Telecommunication Authority (PTA) directed the country's ISPs to block access to the website *www.youtube.com* from their users [9][10]. *On Sunday, 24 February 2008, Pakistan Telecom (AS17557) started an unauthorized announcement of the prefix 208.65.153.0/24. One of Pakistan Telecom's upstream providers, PCCW Global (AS3491) forwarded this announcement to the rest of the Internet, which resulted in the hijacking of YouTube traffic on a global scale* [9][10]. *On 8 April 2010, AS23724 (one of the data centres operated by China Telecom - China's largest ISP) originated about 37,000 unique prefixes* [11] *that were not assigned to them*. This is what we typically call a prefix hijack. Such misconfiguration may cause inconvenience or a communication failure. Fortunately enough, such problems do not occur often.

Addressing BGP's problem is a difficult task, not in terms of technology but in terms of *operational policy*. For example, changing the contents of the packets of the protocol or the mechanism of the propagation assumes a coordinated and simultaneous change to other ISPs which is not easily feasible. Due to this fact, most of the implemented modifications have focused in the *decision process* that the BGP uses to select routes. As a result the current protocol relies on the decision process and on the policies used to increase its associated security.

Routing security

The BGP protocol cannot verify the content of the messages it receives from a valid peer. By sending or accepting false information, an ISP can subvert a neighbour's routing goals, thus causing routers to overload and fail or to degrade service quality. False information can have a significant influence on routing in an AS, even if the source of the information is several AS hops away.

In the following figure, AS 110 announces the same prefix (20.20.0.0/16) to AS P. The decision mechanism of BGP on AS P would choose the prefix from AS 110 because the path vector is shorter (one hop) than the one (Q, R, U) originated from the valid origin. Thus traffic originating from AS A would not lead to AS U but to AS 110. This is commonly referred as prefix hijacking (PH).

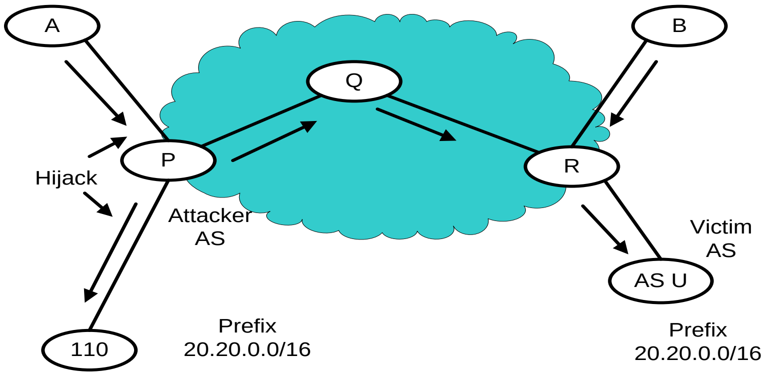

Figure 2: False origin AS110 announces prefix 20.20.0.0/16

The following are variations of the PH problem:

- ***Sub-prefix hijacking****:* when an AS originates the routes of a sub-prefix (ie, network prefix of smaller size than the hijacked prefix)
- ***Independent prefix hijacking:*** when an AS originates the routes of a prefix from an unused address space (also known as Bogon addresses)
- ***Man in the middle (MITM):*** when an AS allows an attacker to have traffic for certain destinations redirected to the attacker.

An ISP may exercise *filtering*, or what is commonly referred to as *filtering*, in order to tune the decision process and thus to protect itself in a limited way against the previous types of threats.

Filtering is implemented in terms of policy filters which express the routing behaviour of each AS. Filtering appears in three different flavours:

- *prefix filtering*, when an AS chooses to limit the routes (ie, prefixes) it announces or receives (eg, receive only 4.0.0.0/16)
- *AS filtering*, when an AS chooses to limit the paths of ASs it announces or receives (eg, allow ONLY the receipt of AS path sequence “`^<AS 110> <AS 5> <AS200>$`”
- *length filtering*, when an AS chooses to limit the length of routing updates (eg, do not receive routing updates > /25 length (255.255.255.64-256)).

Typically a tier-1 one provider exercises length filtering, tier-2 (or an ISP) employs prefix filtering and an end customer uses AS based filtering. Usually an ISP operates on one AS, though some ISPs may operate on multiple ASs for business reasons (eg, to provide more autonomy to an ISP’s backbones in the United States and Europe) or for historical reasons (eg, a recent merger of two ISPs). Non-ISP businesses (enterprises or campuses) may also operate their own ASs in order to gain the additional routing flexibility that arises from participating in the BGP protocol. Compared to enterprise networks, ISPs usually have more complex policies arising from the fact that they often have several downstream customers or they connect to certain customers in multiple geographic locations, or because they have complex traffic engineering goals; thus they deploy BGP on internal routers (rather than just on border routers).

Certainly there are a large number of different policies which could affect the security of the operation of the protocol. Routing policies express the routing behaviour of each AS.

The routing policies are typically described and documented using the Routing Policy Specification Language (RPSL) [12] which is a vendor independent language. In some cases RPSL statements are falsely documented which can lead to potential security problems. For instance, consider the policy of AS 110:

```
aut-num:       AS 110
as-name:       Autonomous system 110
import:        from AS 110
               accept ANY
export:        to AS110
               announce AS110
```

AS110 accepts anything from its upstream provider AS P. Although this might seem straightforward because AS110 is single-homed, such a statement hides potential security problems. The reason is that P, via its policy:

```
aut-num:          AS P
as-name:          Upstream of 110
import:           from  AS Q accept  ANY;
```

accepts anything from AS Q. AS P, via its policy, accepts anything from its Internet service providers. This situation lead us to consider the case where an AS on the Internet can inject a more specific route and how, via the BGP decision-making mechanism, this would be propagated all the way through, even inside AS 110!!

Such cases, although not typical (since single-homed institutions may not run BGP), indicate a potential problem. This is obvious in the case of most ISPs where, via their documented policy, they accept routes belonging to each client and announce anything to them. In the case of AS P, it is assumed that the peers of its upstream providers are *export-controlled* in the same way. Although this is valid for the peers of its upstream providers, *it is not the case of the upstream of its upstream providers*.

In general the potential vulnerabilities of BGP can be classified into three categories:

- insecure transmission of the messages carried inside the protocol (see section 0)
- improper routing policy (see section 0)
- unverified association of attributes (ie, addresses) with respect to an entity (ie, an AS) (see section 0)

Unsecure protocol operation due to insecure channel operation

BGP was designed without any specific transmission protocol in mind, as is the case with other routing protocols, especially those which operate at the intradomain level (ie, OSPF, IS-IS). BGP operates directly on top of Transmission Control Protocol (TCP), so there is no need for extra error correction and flow control from the routing protocol. Routers exchange BGP messages by establishing a TCP session that runs on port 179. The communication channel between two BGP speaking devices is vulnerable to the same kinds of attacks as any TCP communication between two TCP speaking devices.

Typically, two TCP speaking devices are vulnerable against integrity, confidentiality and denial of service (DoS) attacks. In the current BGP communication case we consider the cases of integrity and DoS attacks, as the issue of confidentiality is somewhat moot because routing policies are typically published on public repositories.

DoS attack

In a DoS attack, the attacker can tap the channel and modify the messages, making them incomprehensible to the receiver which could potentially lead to a router crash. The attacker could also send TCP RST to BGP listeners thus temporarily affecting the forwarding of packets between the BGP speakers.

Mitigation

The previous two types of attacks can be addressed by legacy techniques. The first technique is to apply IPsec [31] along the network level. This technique, although not BGP specific, could be applied selectively on the BGP traffic. IPsec was conceived in order to build virtual private networks (VPNs) on top of insecure internet channels. IPsec is an umbrella of protocols consisting of: a) the Internet Key Exchange (IKE) [17] for key negotiation, b) the IPsec Authentication Header (AH) [15] protocol, and c) IPsec Encapsulating Security Payload (ESP) [16]. All of the above mentioned protocols provide security according to the level of security required, which can vary from either authenticated transmission to very strong channel encryption.

Limited applicability

Although this technique requires zero changes in the protocol, it demands extra functionality in router capabilities. Ignoring the possible financial cost of this added functionality, which is almost common place in today’s PCs and small routers, one should not underestimate the additional computational effort required for encrypting the channel, especially in the case of routers with a full Internet routing table (FIRT). Such a table could host approximately 250K to 300K routing entries. In such a case, the first transmission of the table requires enormous computing capabilities and so an extra hardware security module should be considered.

Integrity attack

In this type of attack the objective is to insert bogus BGP messages. If the message added is a BGP update then a new prefix appears. If the message is a BGP withdrawal, a prefix previously reachable disappears.

Mitigation

The MD5 protection of BGP sessions via the TCP MD5 Signature Option (RFC 2385) [18] addressed this need. This mechanism defines a new TCP option for carrying an MD5 [19] digest in a TCP segment. This digest acts like a signature for that segment, incorporating information known only to the connection end points. This security mechanism is widely implemented across the routing software vendors who have BGP capabilities.

Another mechanism for message integrity relies on the generalized TTL security mechanism GTSM (RFC 3682) [20]. GTSM relies on the fact that the vast majority of protocol (in our case BGP) peerings are established between routers that are adjacent. Thus, most protocol peerings are either between directly connected interfaces or, in the worst case, between loopback and loopback, with static routes to loopbacks. Since TTL spoofing is considered nearly impossible, a mechanism based on an expected TTL value can provide a simple and reasonably robust defence from infrastructure attacks based on forged protocol packets. In the example presented in Figure 3, routers set the TTL on a packet to 255, which is decremented when it reaches a peer. If the router is configured such that no packets with a TTL of less than 254 will be accepted then remote adversaries attempting to inject a malicious info GTSM mechanism, will have their packets dropped.

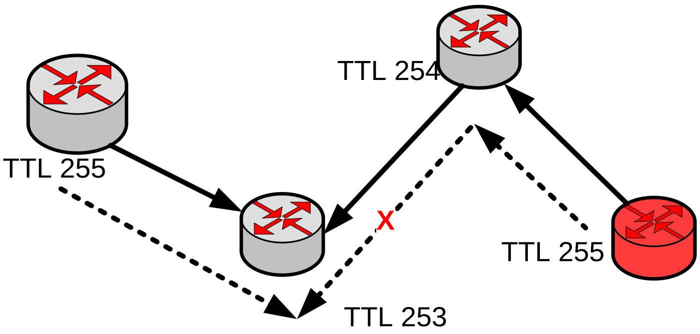

Figure 3: Application of the generalized TTL security mechanism

However, MD5 has been found to be vulnerable to collision attacks [43]. Recently the TCP Authentication Option (TCP-AO) [21], which makes the TCP MD5 obsolete, has been proposed as a standard. TCP-AO specifies stronger message authentication codes (MACs) to protect against replays even for long-lived TCP connections. TCP-AO is compatible with either a static master key tuple (MKT) configuration or an external, out-of-band MKT management mechanism; in either case, TCP-AO also protects connections when using the same MKT across repeated instances of a connection, using traffic keys derived from the MKT and coordinates MKT changes between endpoints. The result is intended to support the current infrastructure of long-lived connections (as used, eg, in BGP) and to support a larger set of MACs with minimal other system and operational changes.

Summary of applicable technologies

The BGP channel security mitigation techniques, as previously described, can be summarized in terms of their capabilities in the following table

	Integrity	Replay prevention	DoS prevention
MD5 integrity check	Yes	Yes	No
TCP-AO	Yes	Yes	No
GTSM	Yes	No	No
IPsec AH	Yes	Yes	Yes
IPsec ESP	Yes	Yes	Yes

Table 1: Applicability of BGP channel security solutions

Recommendation 1:

Adjacent peers should enable TCP-AO and the TTL security mechanisms. If TCP-AO is not available, MD5 should be used in the interim period. If DoS prevention constitutes a required characteristic of the desired protection, the IPsec AH should be considered. IPsec ESP solution should be considered only for a customer to ISP peering, as it poses a considerable computational burden in a wider context.

Improper routing policy and monitoring of prefix announcements

In a previous section it was shown that routing policy statements could become a potential security threat.

Vulnerability

One AS (the attacker) either announces a prefix that already exists but is with a different AS (victim) or announces a sub-prefix with a special address belonging to the victim.

It is impossible to combat prefix hijacking in the general case with filtering. It is possible though to increase the defences against it. Assume that, in the general case, every ISP has a number of single-homed and multi-homed[10] customers. The typical import routing policy for every ISP is to accept everything from its upstream provider. This should change. Every ISP should modify its import routing policy with respect to its single versus multi-homed customers. The routes corresponding to single-homed customers should never be accepted from upstream providers. Routes corresponding to multi-

[10] *In the present case we consider a customer as multi-homed when it peers concurrently with two or more different ASs.*

homed customers cannot be protected, as they can enter via an alternate AS path from the upstream provider.

Partial mitigation

Although complete protection against prefix hijacking is not feasible, it is possible to deploy a monitoring system which raises the alarm when prefixes belonging to multi-homed customers appear in the routing table. Anomaly- based detection solutions, such as the Prefix Hijack Alert System (PHAS) [22] and the Pretty Good BGP (PG-BGP)[25], work by collecting BGP routing data from a properly selected vantage point. There are a number of data sources of BGP routing information available for a BGP security solution such as Routeviews [30], RIPE-Routing Information Service (RIPE-RIS) [31], and Cooperative Association for Internet Data Analysis (CAIDA) [29].

Anomaly-based detection systems use data in multiple different ways. Some of them depend on data from RIRs as in the Nemesis tool [26], while others [22][24] use BGP trace data. Trace data is obtained from global BGP monitoring infrastructures (eg, RIPE-RIS, Routeviews) or a BGP speaker where the algorithm operates.

Prefix Hijack Alert System

The PHAS was the first approach which detected IP prefix hijacking. The idea behind this approach is simple: monitor BGP routing data and report any announcement of a new AS associated with an IP prefix or sub-prefix related to the prefix of multi-homed customers. The overall architecture of the system is shown in Figure 4.

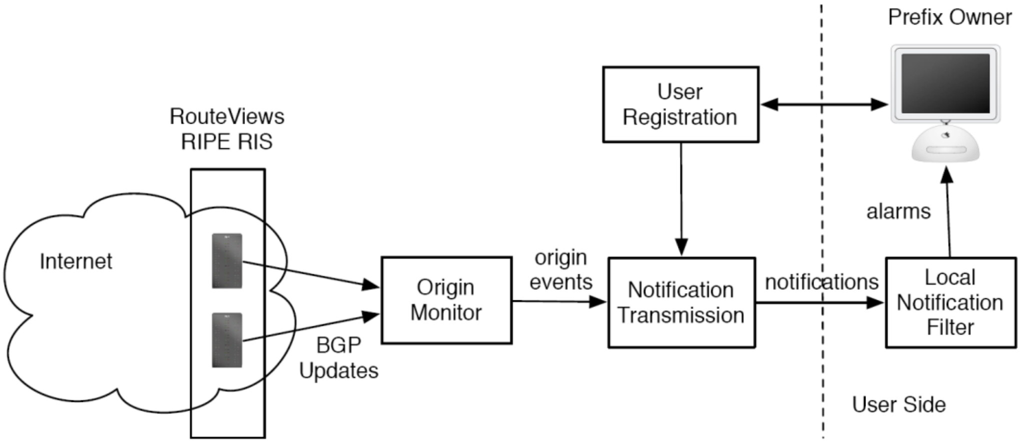

Figure 4: Components of PHAS [24]

PHAS does not associate a prefix with a true or false origin, thus it reports all origin changes to the prefix owner. However, not all origin set changes may be of interest. So the local notification filter could be tuned to limit unwanted alarms.

Pretty Good BGP

The Pretty Good BGP (PG-BGP) method operates with historical BGP trace data. In the first version of the algorithm, historical routing update data (of the form {prefix, origin AS}) and routing information base (RIB) entries are recorded, over the last h days (h = 10 days). Old routes (older than 10 days) are eliminated by the anomaly detector, if they are no longer active. A new update that is not in history records is considered suspicious and thus the update is propagated with lower local-pref value. The quarantine lasts for a period of s hours (eg, s = 24 hours); if the sub-prefix is not withdrawn during that time, then the update is propagated.

Nemecis system

Nemecis system [26][27] is a registry-based method driven by declarative RPSL policy data from the RIRs and IRRs. Following a routing update message {prefix, (AS)} Nemecis checks for the existence of prefix registration (ie, inetnum in RPSL-based RIRs or NetHandle in SWIP-based RIRs), AS registration (aut-num in RPSL and ASHandle in SWIP), and route objects in IRRs or RADb. A consistency check is performed between the declared objects in terms of their attributes (ie, organization, maintainer, email handle, etc). Where a full or partial consistency check fails, the algorithm can generate alerts.

Summary of applicable technologies

Table 2 outlines various prefix hijacking solutions and their capabilities. It is beyond the scope of this report to describe each one of them separately.

	Detection system	PH	Sub-prefix PH	Path spoofing	MITM
PHAS [22]	T	Y	Y	Limited	N
PG-BGP[25]	T	Y	Y	Y	Limited
Nemesis [26]	R	Y	Y	N	N
Qiu et al [24]	T	Y	Y	Y	N
Sriram et al [36]	T+R	Y	Y	Y	N
Krugel et al [34]	T	Y	N	Y	N
Hu et al [35]	T	Y	Y	Y	N

Table 2: Taxonomy of prefix hijacking solutions (PH: prefix hijacking, Y: yes, N: no, R: registry, T: trace data, MITM: man in the middle)

Recommendation 2.1:

Single-homed customers should be protected with filtering 'in the import BGP' configuration clauses.

Recommendation 2.2:

Inspect, and verify the import routing policy of every ISP at a national level. Differentiate the import routing policies for single v multi-homed customers. The protection of multi-homed customers can be improved by operating a prefix hijacking alarm solution around the Internet eXchange Points of every country properly notifying local ISPs.

Cryptographic validation of routing updates

The association of attributes (ie, addresses) with respect to an entity (ie, an AS) can be verified with the cryptographic validation of routing updates. Recent work within the standard bodies and in research has attempted to produce cryptographic frameworks for BGP security. A brief introduction to public key cryptography is warranted in order to explain the cryptographic validation of routing updates.

Public key cryptography

One commonly used security technology today is public key cryptography, which utilizes a pair of keys, A and B. Anything enciphered with key A can be deciphered only with key B and vice versa. In contrast to symmetric cryptography, knowledge of one key does not lead to discovery of the other key. Typically, key A is considered to be private and is never revealed, while key B is commonly published. Public key cryptography, besides encryption, provides the capability to validate integrity without encrypting the original message, just by generating a digital signature with a private key A. Any attempt to alter the message will be detectable because the signature will not match the content.

Individual members, the holders of pairs of keys, would have to exchange public keys with every other member in order to facilitate communication. Fortunately, the number of exchanges can be reduced if one entity, called a certification authority (CA), can be trusted in terms of certification criteria and procedures for the association between identity and public key ownership. In strictly technical terms, a *digital certificate* is a digitally signed public attestation by a certification authority that associates a subject's public key (B) with some attribute of that subject. In addition, digital certificates can be used to identify role membership or right-of-use authorities, which is relevant to resource certification.

Secure BGP

Secure BGP (S-BGP) was introduced as an extension to BGP to protect against false routing updates (called, in one word, UPDATEs) [37]. S-BGP applies strong authentication and authorization features to BGP based on public-key cryptography.

S-BGP introduces three major additions to BGP. First, a public key infrastructure (PKI) is introduced in the interdomain routing infrastructure to authorize prefix ownership and validate routes. The private keys are stored in S-BGP speakers, while the public keys are made available through a hierarchical PKI infrastructure. Second, it adds a new transitive attribute to BGP updates. That attribute verifies the authorization of routing UPDATEs, and avoids route modifications from intermediate S-BGP speakers. Third, IPSec can be applied, if routing confidentiality is required.

Address Attestations (AAs) and *Route Attestations* (RAs) are the two key components of S-BGP. An Address Attestation (AA) is produced by the owner of a prefix, and it is used by S-BGP nodes to verify the validity of advertisements of address prefixes from an originating AS. Route Attestations (RAs), on the other hand, are added by S-BGP routers in UPDATEs, authorizing a neighbouring AS to propagate the route contained in that UPDATE. S-BGP uses a PKI infrastructure to verify AAs and RAs.

RAs flow with UPDATE messages through a sequence of S-BGP routers. Each S-BGP router along the path validates the integrity of an UPDATE before signing it and passing it to its neighbours. The result is an onion style attestation that includes signatures from all the routers along the path (Figure 5).

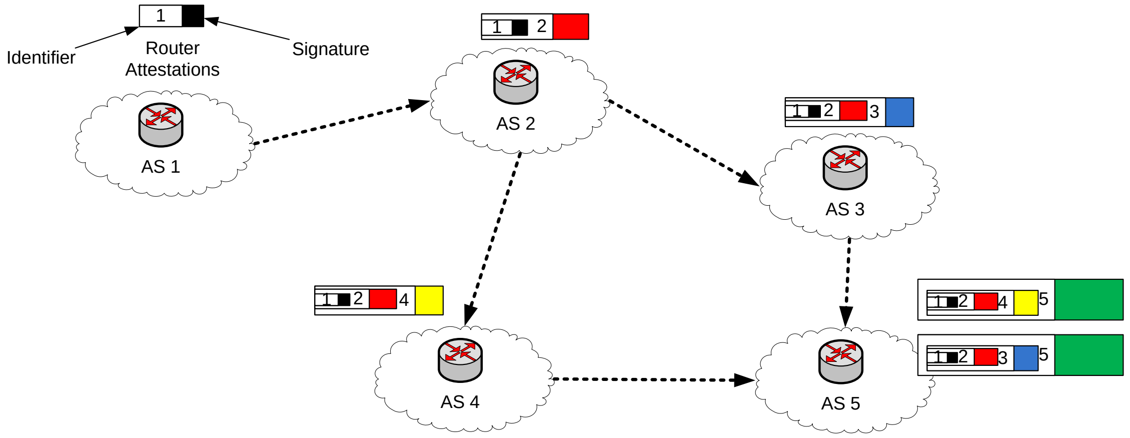

Figure 5: Route attestations in S-BGP.

Secure origin BGP

Secure origin BGP (soBGP) was introduced as a lightweight alternative to S-BGP, mainly by researchers at Cisco Systems [39]. The objective of soBGP was to verify two issues of routing information, namely that an AS is the authoritative owner of a given prefix and to verify that that the advertising AS has at least one valid (in terms of policy and topology) path to that destination. soBGP utilizes three types of certificate for the required verification. The *entity certificate* establishes the identity and public key of an AS. The *authorization certificate* verifies the assignment and delegation of IP address blocks, and it is used to validate prefix ownership. The *policy certificate*, on the other hand, authenticates in accordance with AS or pre-prefix policies and AS connectivity information and it is used to verify the validity of a route. soBGP routers use a topology database to validate received routes. soBGP utilizes a web-of-trust model, instead of a hierarchical PKI, for certificate validation, relying on the existing relations between ISPs.

The major difference between the S-BGP and soBGP is the nature of the RA. RAs are dynamic for s-BGP while they are static for soBGP. soBGP provides various deployment options [38], such as verify before accept, accept and verify afterwards, thus applying a trade-off between routing convergence and security.

Interdomain routing validation

Interdomain routing validation (IRV) also performs path and origin verification. This method was originally proposed by Goodell et al [40]. IRV isolates the authentication component from the BGP protocol, by introducing a separate companion protocol called IRV. With IRV, each AS employs one or more IRV servers. For every BGP UPDATE message, the corresponding IRV server for every AS in the AS path is contacted, to verify both the origin and the routing path of the received UPDATE.

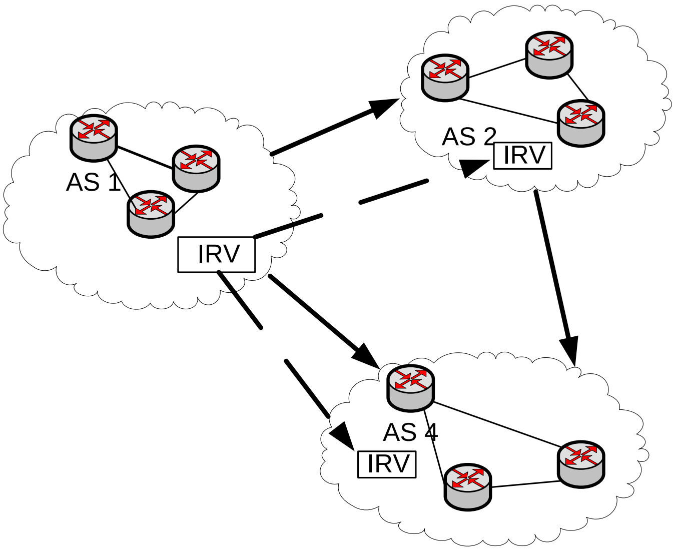

Figure 6: ASs running the IRV protocol query the appropriate authorities for validation of received routing data. IRV validators are independent of routers within an AS.

Any data item exchanged over the BGP protocol can be validated by the IRV. It is up to an AS to provide reliable data. IRV servers operate similar to routing registries, but manage information only from the parent AS. The operational model of IRV is more distributed than any of the ones previously described because ASs retain control over the validated data, and hence may provide more fresh and accurate data. The trust point is relocated from the registry to the AS for accurate assertions.

Deployment obstacles

S-BGP and soBGP are not evolutionary in terms of BGP operations. Routers with the former functionality cannot cooperate with the new functionality. Routers with the new functionality have great difficulty in verifying paths with partial deployment. IRV on the other hand needs no additional BGP functionality; it requires improved functionality of a registry in order to validate the correctness of routing information.

Summary of applicable technologies

None of the previously described protocols are in use today, either due to the increased computational requirements of the cryptographic functionality or due to increased requirements from registries for correctly communicating address ownership and delegation, which is a necessary first condition for implementing real origin authentication solutions.

A comparison of the previously described cryptographic solutions for BGP is given in Table 3. The solutions are compared in terms of authentication services: topology, path and origin.

Solution	Topology check	Path check	Origin check
S-BGP	Strong	Strong	Strong
soBGP	Strong	None	Strong
IRV	Strong	Moderate	Strong

Table 3: Comparison of security BGP solutions

IETF Framework for secure interdomain routing

IETF is the authoritative technology body for changes governing the operation and routing of the IP protocol. BGP was standardized by an IETF working group, hence the evolution of BGP was addressed by IETF as a topic for securing interdomain routing by following a typical 'IETF way' approach. The first step was to initiate a BoF session and then, given the increased interest in the community, the Routing Protocol Security Requirements (RPSEC) Working Group (WG) was established. The WG has considered various vulnerabilities in today's interdomain routing system and has produced a set of requirements that should be addressed—without necessarily specifying the solution.

The set of requirements is complete and is described in the draft-ietf-rpsec-bgpsecrec-10. Once the description of requirements had achieved a suitable level of consensus within the IETF community, it was possible to start working on tentative solutions. Various proposals [37][38][39] have started to emerge.

The Keying and Authentication for Routing Protocols (KARP) was chartered by IETF with the task of improving the communication security of the packets on the wire used by the routing protocols and not the security of the protocol itself. One of the WG's objectives is to submit a specification document for BGP to the IESG to be considered as a proposed standard by April 2011.

The IETF secure inter-domain routing (SIDR) Working Group started in April 2006 to work on basic security questions regarding the validity of routing information, e.g., prefix AS origination, accurate AS identification, and validating address prefix and AS number. The scope of work in the SIDR WG is to formulate an extensible architecture for routing security. Given the complexity, both technical and policy-wise, the SIDR WG process is expected to take 10 to 15 years before there will be a general uptake of these technologies. In addition, the SIDR WG has been working with a number of stakeholders on the specification of the resource public key infrastructure (RPKI).

Before describing the possible solution which appears to be emerging, it is important to provide some definitions.

Resource certificates

A resource certificate is a conventional X.509 certificate that adheres to the PKIX profile (RFC 5280) [42] supplemented with a certificate extension. This extension, which is asserted to be critical, typically lists a collection of IP number resources (IPv4 addresses, IPv6 addresses, and AS numbers) (RFC 3779) [41]. Due to this critical extension, these resource certificates cannot be used in a conventional manner for identity verification or web-server assurance.

These certificates attest the allocation of associated number resources and not the subject identifier; namely that the certificate's issuer has granted to the entity represented by the certificate's subject the right-of-use of the associated set of IP number resources listed in the certificate's extension. The right-of-use concept mirrors the resource allocation framework in operation today, where the IP address space is governed by IANA. Certificates enable the validation of assertions related to resource allocations by any third party (relying party).

For instance, assuming that an entity (ie, a company) receives an address allocation block from a particular regional Internet registry (RIR), only that RIR can issue a resource certificate for the entity which includes its public key and the allocated number resources. Anything signed by the end entity's private key, whether it is a routing update protocol message in a new BGP variant or an administrative request to an ISP to route a prefix or an assertion of a right-of-use of a number resource, can be validated through the RIR's issued certificate which contains the matching public key and the IP number resource that are enumerated in this certificate. An issued resource certificate can be verified in the framework of a Resource Public Key Infrastructure (RPKI).

Signed attestations and authorities

The overall objective of digital certificates—and resource certificates in particular—in the context of secure interdomain routing is to build a web of transitive trust which allows a relying party to verify the validity of routing protocol messages. In the context of BGP, the validation of the authenticity of route objects is among the objectives.

Route origin attestation (ROA) is a signed artifact (document) which proves that an autonomous system (AS) has been given permission by an IP address block holder to advertise routes to one or more prefixes within that block. The message of an ROA, for example, might state the following: 'ISP 100 permits AS 65004 to originate a route for the prefix 192.4.100.0/24.' The message is signed using a cryptographic message signature (CMS) (RFC3852) by the address holder which is represented by an end entity (EE) certificate in the ROA.

The resulting object is published in the RPKI as a routing origin authorization (ROA). A relying party can validate a ROA by verifying the three embedded structures: a) that the digital signature of the ROA is valid, b) that the resources in the associated EE certificate encompass the prefixes specified in the document, and c) that the EE certificate itself is valid in the context of the RPKI.

Another type of signed attestation verifies that there is an inter-domain adjacency between the announcing AS (local AS) and those ASs adjacent to it. If an autonomous system advertises extra autonomous systems with their associative routes, it should sign an analogous statement, called an adjacency attestation (AA).

It should be mentioned that the concepts described so far are a work in progress in the SIDR working group's agenda of study. The working group has not yet reached any consensus regarding the decision to advance these proposals further along the Internet standards process.

Resource public key infrastructure

The resource public key infrastructure (RPKI) describes the structure of a framework used by resource certificates. The objective of the RPKI is to construct a robust hierarchy which allows relying parties to validate assertions about IP addresses and AS numbers and their use.

The structure of the RPKI is related to the existing framework of address allocation worldwide, namely, IANA manages the number resources at a very high level and only provides a registry of currently allocated and unallocated address blocks. IANA does not allocate resources directly to end users. Instead, it allocates blocks of number resources to the regional internet registries (RIRs). The RIRs perform the next level of distribution: allocating number resources to local internet registries (LIRs), or to national internet registries (NIRs), and end users. NIRs make allocations to LIRs and end users, and LIRs allocate resources to end users (Figure 7).

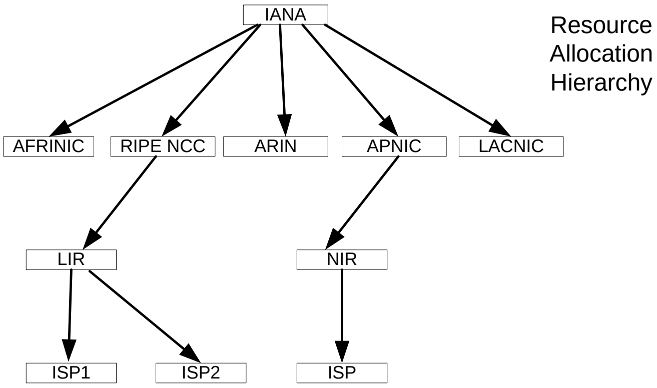

Figure 7: Address distribution hierarchy for the Internet

The RPKI framework establishes an equivalent allocation hierarchy based on public key cryptography. A possible interpretation is that IANA manages the root of RPKI with a self-signed certificate and issues subordinate CA certificates with an extension describing the addressing resource allocated to RIRs.

Any entity which is allowed to make further allocations of resources to other parties must be capable of issuing resource certificates which correspond to these allocations. Similarly, any entity holder that wishes to attest the usage of number resources needs to create a signed *attestation* and issue an end entity (EE) certificate which performs the digital signing operation of the attestation. For that reason, all issued certificates that correspond to allocations have the capability to issue a subordinate CA to create further subordinate EE certificates that correspond to the generation of digital signatures on attestations.

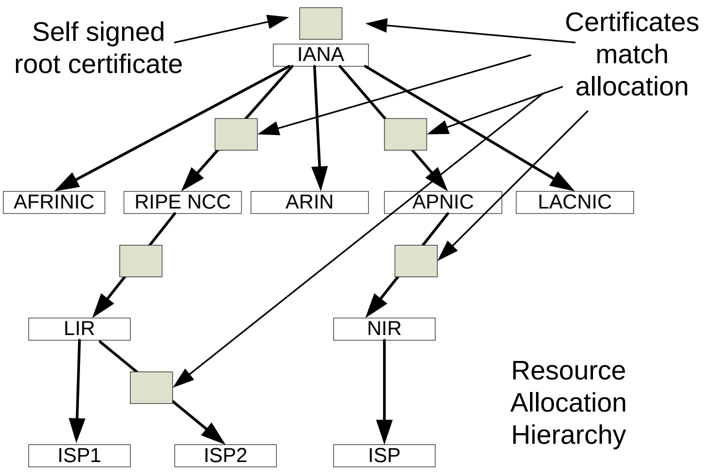
Self signed root certificate
Certificates match allocation
IANA
AFRINIC
RIPE NCC
ARIN
APNIC
LACNIC
LIR
NIR
Resource Allocation Hierarchy
ISP1
ISP2
ISP

Figure 8: RPKI resource certificate hierarchy

The validation of resource certificates is performed similarly to conventional PKI. Certificate revocation lists (CRLs) are used to control the validity of issued certificates, and every CA certificate in the RPKI must issue a CRL.

Resource certificates and attestations are considered to be public documents stored in openly accessible repositories. The availability of all the repositories is fundamental to the security of the public Internet's interdomain routing system. For that reason RPKI should periodically scan all the repositories to construct a view that is as complete as possible thus becoming a trust anchor (TA). In a routing context (inside an AS), a relying party (the local AS) may validate certificates and attestations using a stored replica of the available objects. A relying party is able to verify that the stored replica is complete using a newly defined object called a manifest. The manifest allows a synchronization comparison, to ensure that a locally managed cache of the RPKI has not changed since a previous synchronization operation.

RPKI is evolutionary in terms of BGP operations and overrides the deployment obstacles of other proposed solutions. Routers with the former functionality will be able to communicate with routers that implement RPKI. An RPKI roadmap should be fostered among vendors, registries and ISPs for the availability of the equipment, the creation of the repositories and the deployment of the technology.

The use of a single authoritative trust anchor has been heavily argued between the stakeholders. Although there have been statements (namely from the IAB[11] and RIRs) in favour the single authoritative trust anchor hierarchy for technical reasons, other members of the community believe that those reasons are not essentially technical but rather political. The proponents specify that in a case of a partial misconfiguration there will be no single correct view and thus a single root hierarchy is needed. The opponents insist that trust cannot be applied equally at any given time thus making a choice inevitable, as such; there is no reason for a single authoritative trust anchor hierarchy. Early trials will aid the resolution of key operational and policy issues, such as the hierarchy of the verification of the framework which has not been settled yet.

[11] *IAB statement on RPKI http://www.ietf.org/mail-archive/web/ietf-announce/current/msg07028.html*

Glossary of Terms

APNIC A regional Internet registry (RIR) that allocates IP and AS numbers in the Asia Pacific region.

ARIN A regional Internet registry (RIR) that allocates IP and AS numbers in the North-American region and parts of the Caribbean.

AS An autonomous system is a collection of connected Internet Protocol (IP) routing prefixes under the control of one or more network operators that presents a common, clearly defined routing policy to the Internet.

BGP The Border Gateway Protocol is the core routing protocol of the Internet. It maintains a table of IP networks or prefixes which designate network reachability among autonomous systems (AS).

Bogon address The term ‘bogon’ (hacker slang derived from ‘bogus’) refers to an IP address that is reserved but not yet allocated by IANA or some other Internet registry. Addresses that have not been allocated to legitimate users should never be routed, and packets that appear to come from these addresses are most likely forged.

Byzantine robustness See http://en.wikipedia.org/wiki/Byzantine_fault_tolerance for a comprehensive definition and explanation of robustness against Byzantine failures.

CA A certificate authority or certification authority (CA) is an entity that issues digital certificates for use by other parties. It is an example of a trusted third party. CAs are characteristic of many public key infrastructure (PKI) schemes.

DNS/DNSSEC The Domain Name System is a hierarchical naming system for computers, services, or any resource connected to the Internet. Most importantly, it translates domain names that are meaningful to humans into the numerical (binary) identifiers associated with networking equipment for the purpose of locating and addressing these devices worldwide. DNSSEC is a suite of specifications for securing certain kinds of information provided by the Domain Name System.

DoS/DDoS A denial-of-service attack (DoS attack) or distributed denial-of-service attack (DDoS attack) is an attempt to make a computer resource unavailable to its intended users.

IAB The Internet Architecture Board is chartered both as a committee of the Internet Engineering Task Force (IETF) and as an advisory body of the Internet Society (ISOC).

IANA The Internet Assigned Numbers Authority (IANA) is the entity that oversees global IP address allocation, AS number allocation, root zone management for the Domain Name System (DNS), media types, and other Internet Protocol related assignments.

IETF The Internet Engineering Task Force (IETF) develops and promotes Internet standards, cooperating closely with the W3C and ISO/IEC standards bodies and dealing in particular with the standards of the TCP/IP and Internet protocol suite. It is an open standards organization, with no formal membership or membership requirements.

Internet exchange An Internet exchange point (IX or IXP) is a physical infrastructure through which Internet service providers (ISPs) exchange Internet traffic between their networks (autonomous systems).

IP hijack IP hijacking (sometimes referred to as BGP hijacking or prefix hijacking) is the illegitimate take over of groups of IP addresses by corrupting Internet routing tables.

IPsec Internet Protocol Security (IPsec) is a protocol suite for securing Internet Protocol (IP) communications by authenticating and encrypting each IP packet of a data stream.

IRR The Internet routing registry consists of several databases where network operators publish their routing policies and routing announcements so that other network operators can use this data.

ISP An Internet service provider (ISP) is a company that offers its customers access to the Internet.

KARP The IETF Keying and Authentication for Routing Protocols working group is tasked with improving the communication security of the packets on the wire used by the routing protocols. This working group is concerned with message authentication, packet integrity, and denial of service (DoS) protection. See also https://datatracker.ietf.org/wg/karp/.

MD5 Message-Digest algorithm 5 is a widely used cryptographic hash function with a 128-bit hash value. MD5 has been employed in a wide variety of security applications, and is also commonly used to check the integrity of files

PKI A public key infrastructure (PKI) is a set of hardware, software, people, policies, and procedures needed to create, manage, distribute, use, store, and revoke digital certificates. A PKI is an arrangement that binds public keys with respective user identities by means of a certificate authority (CA).

RADb The Routing Assets Database is an IRR run by Merit Network.

RFC A request for comments (RFC) is a memorandum published by the Internet Engineering Task Force (IETF) describing methods, behaviours, research, or innovations applicable to the working of the Internet and Internet-connected systems. The IETF adopts some of the proposals published as RFCs as Internet standards.

RIR A regional Internet registry is an organization overseeing the allocation and registration of Internet number resources within a particular region of the world. Resources include IP addresses (both IPv4 and IPv6) and autonomous system numbers (for use in BGP routing)

RIPE NCC A regional Internet registry (RIR) that allocates IP and AS numbers in the European, Middle East, and Central Asian region.

ROA A route origin authorisation is a digitally signed object that provides a means of verifying that an IP address block holder has authorised an autonomous system (AS) to originate routes to one or more prefixes within the address block.

Route aggregation/disaggregation The Border Gateway Protocol allows the aggregation of specific routes into one route. Route aggregation can be used to decrease the size of the BGP routing tables. This helps in speeding up the convergence time and improves network performance. Route disaggregation is the reverse process, where a route is split into two or more specific routes, and hence increases the size of the BGP routing tables.

RPKI A resource public key infrastructure system can be used to certify autonomous system (AS) numbers and IP addresses allocations in order to substantially improve the security of the routing system.

SIDR The IETF Secure Inter-Domain Routing working group works on the formulation of an extensible architecture for an inter-domain routing security framework. This framework must be capable of supporting incremental additions of functional components. See also https://datatracker.ietf.org/wg/sidr/.

Tier 1/2/3 network A Tier 1 network is a transit-free network that does not pay settlements to any other network to reach any other portion of the Internet. Therefore, in order to be a Tier 1 network, a network must peer with every other Tier 1 network. A Tier 2 network peers with some networks, but still purchases IP transit or pays settlements to reach at least some portion of the Internet. A Tier 3 network solely purchases transit from other networks to reach the Internet.

References

[1] Y Rekhter, T Li, *A border gateway protocol 4*, IETF RFC 1771, March 1995

[2] S Halabi, *Internet Routing Architectures* (2nd Edition), Cisco Press

[3] IANA, *Autonomous System Numbers*, March 2003

[4] IANA, *Internet Protocol V4 Address Space*, http://www.iana.org/assignments/ipv4-address-space

[5] IANA, *The Internet Assigned Numbers Authority*, May 2003. http://www.iana.org/

[6] ICANN, *The Internet Corporation for Assigned Names and Numbers*, May 2003. http://www.icann.org/

[7] *The Regional Internet Registry Policy Development Process*, http://www.isoc.org/briefings/010/

[8] RFC 2050, *Internet Registry IP Allocation Guidelines*, http://www.faqs.org/rfcs/rfc2050.html

[9] *YouTube Hijacking*: *A RIPE NCC RIS case study*, http://www.ripe.net/news/study-youtube-hijacking.html

[10] BBC News: *Pakistan blocks YouTube website*, http://news.bbc.co.uk/2/hi/south asia/7261727.stm.

[11] *Chinese ISP hijacks the Internet*, http://bgpmon.net/blog/?p=282

[12] RFC 2280 - *Routing Policy Specification Language* (RPSL), http://www.faqs.org/rfcs/rfc2280.htm

[13] RFC 2401 - *Security Architecture for the Internet Protocol*, http://www.faqs.org/rfcs/rfc2401.html

[14] RFC 4301 - *Security Architecture for the Internet Protocol*, http://www.faqs.org/rfcs/rfc4301

[15] RFC 4302 - *IP Authentication Header* (AH), http://www.faqs.org/rfcs/rfc4302

[16] RFC 4303 - *IP Encapsulating Security Payload* (ESP), http://www.faqs.org/rfcs/rfc4303

[17] RFC 4306 - *Internet Key Exchange* (IKEv2) *Protocol*, http://www.faqs.org/rfcs/rfc4306

[18] RFC 2385 - *Protection of BGP Sessions via the TCP MD5 Signature Option*, http://www.faqs.org/rfcs/rfc2385

[19] RFC 1321 - *The MD5 Message-Digest Algorithm*, http://www.faqs.org/rfcs/rfc1321

[20] RFC 3682 - *The Generalized TTL Security Mechanism* (GTSM), http://www.faqs.org/rfcs/rfc3682

[21] *TCP Authentication Option* http://www.ietf.org/internet-drafts/draft-ietf-tcpm-tcp-auth-opt-11.txt

[22] M Lad, D Massey, D Pei, Y Wu, B Zhang, and L Zhang. *PHAS: A Prefix Hijack Alert System*. In Proc. USENIX Security Symposium, 2006.

[23] *University of Oregon Route Views Project*, http://www.routeviews.org/

[24] J Qiu, L Gao, S Ranjan, and A Nucci, *Detecting bogus BGP route information: Going beyond prefix hijacking*, SecureComm 2007

[25] J Karlin, S Forrest, and J Rexford, *Pretty Good BGP: Improving BGP by Cautiously Adopting Routes*, IEEE ICNP 2006, Santa Barbara, CA, USA, Nov. 2006

[26] G Siganos and M Faloutsos, *A Blueprint for Improving the Robustness of Internet Routing*,

Security ‘06, 2006.

[27] G Siganos and M Faloutsos, *Analyzing BGP policies: methodology and tools*, IEEE Infocom, 2004.

[28] BGPmon, http://bgpmon.net/

[29] CAIDA, http://www.caida.org/

[30] RouteViews, http://www.routeviews.org/

[31] RIPE-RIS, http://www.ripe.net/ris/

[32] Merit Network Routing Assets Database, http://www.radb.net/

[33] IETF Working Group Secure Inter-Domain Routing (SIDR),Routing Protocols Security(RPSEC)

[34] C Krugel, D Mutz, W K Robertson, and F Valeur, *Topology-Based Detection of Anomalous BGP Messages*, in RAID, 2003, pp 17–35

[35] Xin Hu and Z Morley Mao, *Accurate Real-time Identification of IP Prefix Hijacking*, IEEE Security and Privacy, Oakland, 2007

[36] K Sriram, O Borchert, O Kim, and P Gleichmann, and D Montgomery, *A Comparative Analysis of BGP Anomaly Detection and Robustness Algorithms*, CATCH _09, Washington D.C., March 3-4, 2009

[37] K Seo, C Lynn, and S Kent, *Public-key infrastructure for the secure border gateway protocol (S-BGP)*, in IEEE DARPA Information Survivability Conference and Exposition II, Anaheim, CA, Jun 2001

[38] R White, *Securing BGP Through Secure Origin BGP*, The Internet Protocol Journal - Volume 6, Number 3

[39] J Ng, *Extensions to BGP to Support Secure Origin BGP (soBGP)*, Internet Draft, Apr 2004

[40] G Goodell et al, *Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing*, in Proceedings of Symposium on Network and Distributed Systems Security, Feb 2003

[41] RFC 3779 - X.509, *Extensions for IP Addresses and AS Identifiers*, http://www.faqs.org/rfcs/rfc3779

[42] RFC 5280 - Internet X.509, *Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*, http://www.faqs.org/rfcs/rfc3682

[43] US-CERT Vulnerability Note VU#836068 - MD5 vulnerable to collision attacks, http://www.kb.cert.org/vuls/id/836068