March 2013

WORKING GROUP 6
Secure BGP Deployment
Final Report

Table of Contents

1. Results in Brief

In this section, we briefly review the mission statement for the working group, and summarize the group's four main recommendations.

1.1 Mission Statement

The Communications Security, Reliability, and Interoperability Council (CSRIC) is a federal advisory committee established by the Federal Communications Commission (FCC) to provide recommendations to the FCC regarding best practices and actions the Commission may take to ensure optimal operability, security, reliability, and resilience of communications systems. The systems at issue include telecommunications, media, and public safety communications systems. The FCC created ten Working Groups to develop information for CSRIC, and each of the groups was given a charter of responsibilities relating to a topic area.

Working Group 6 (WG6) on Secure BGP Deployment was established in August 2011 with the following mission statement: "The Border Gateway Protocol (BGP) is used for inter-domain routing on the Internet. BGP relies on mutual-trust among operators of gateway routers to ensure the integrity of the Internet routing infrastructure. Over the years, this trust has been compromised on a number of occasions, both accidentally and maliciously, revealing fundamental weaknesses of this critical infrastructure. This Working Group will recommend the framework for industry regarding incremental adoption of secure routing procedures and protocols based on existing work in industry and research. The framework will include specific technical procedures and protocols. The framework will be proposed in a way suitable for opt-in by Internet Service Providers (ISPs) in order to create incentives for a wider scale, incremental ISP deployment of secure BGP protocols and practices in a market-driven, cost-effective manner."

1.2 Executive Summary

The Border Gateway Protocol (BGP) is the glue that holds the disparate parts of the Internet together, by allowing independently-administered networks (called Autonomous Systems, or ASes) to announce reachability to IP address blocks (called prefixes) to neighboring networks. Like many of the protocols underlying the Internet, BGP is vulnerable to accidental misconfigurations, malicious attacks, and software bugs, which can cause the spread of "bogus" routing information throughout the Internet. When ASes inadvertently select bogus routes, they may direct data traffic into "blackholes" (that drop the traffic) or detour the traffic over circuitous paths that traverse unexpected ASes that may snoop on the data. Bogus routes arise relatively often, with high-profile incidents affecting many Internet services happening several times a year.

The fundamental notion in all BGP security solutions is distinguishing "legitimate BGP routes" from "bogus BGP routes". Loosely speaking, a BGP route announcement is legitimate if the IP prefix and BGP route attributes it contains are "consistent with the *intent* of the network address holder and transited network providers". Intent can be expressed as the routing-policy specifications of the ASes along the path, and otherwise is open to interpretation. This working group compared a range of security solutions designed to prevent the propagation of bogus routes. The comparison spanned a variety of dimensions, including technical maturity, implementation cost, trust models and governance, security benefits and residual threats, new attack surfaces, complexity and design trade-offs, the feasibility of incremental deployment, and the impact on the autonomy of network operators.

The phrase "to secure BGP routing" used in this report means "*to add to BGP, or to the operation of BGP, mechanisms preventing or limiting propagation of routes with bogus attributes*". Over the years, many solutions have been proposed for BGP security, and we expect more will be explored in the future. In our work, we do not address previously well-studied aspects of BGP security involving session security, router access security, and techniques for defending internal networks, in part because CSRIC Working Group 4 addressed some of these issues. We focus on solutions that are becoming available in the relatively near term, rather than solutions still in the preliminary stages of investigation or standardization. As such, we focus primarily on the Resource Public Key Infrastructure (RPKI) and the use of RPKI data for "origin validation," since this proposed solution has some initial traction in standards bodies and deployment by Regional Internet Registries (RIRs). We also briefly discuss other solutions under active investigation and experimentation in the body of the report.

Since RPKI is still in the early stages of deployment, network operators have not had much time to experiment and learn from early implementations. As such, our report focuses on providing high-level guidance concerning participation in RPKI, and a high-level analysis of risks of the RPKI. In particular, the Working Group has the following four recommendations, which are described in greater depth, and with corresponding background information, in the main sections of this document:

1 **Accurate records about Internet number resource holders:** All techniques for improving the security of inter-domain routing rely on authoritative, accurate, and timely information about which ASes are authorized to originate routes for each IP address block. To improve the accuracy of records in commonly used Internet Routing Registries (IRRs), AS operators should ensure their IRR records are accurate, complete, and up-to-date. Yet, IRRs have inherent limitations. In addition to maintaining their IRR records, AS operators should begin using systems such as Resource Public Key Infrastructure (RPKI) to generate certificates and Route Origin Authorizations (ROAs). To prevent missing or conflicting information in the RPKI, we encourage the establishment of a single, global "root of trust" for the RPKI. At this stage experimentation and research with alternative origin authorization techniques is not discouraged.

2 **Cautious, staged deployment of RPKI origin validation:** The RPKI provides information that AS operators can use to detect and prevent the spread of bogus origin AS information in BGP. Yet, ASes should retain autonomy in setting their policies for selecting and disseminating routes, including a decision whether and how to use the RPKI data. AS operators should follow a cautious and staged deployment of RPKI, starting by using RPKI data in an out-of-band fashion as one of several ingredients in detecting suspicious routes and constructing their route filters. Any future *fully automated* use of RPKI data in filtering "invalid" routes should come only after AS operators are highly confident in the reliability and timeliness of the RPKI data.

3 **Mitigating risks inherent in the RPKI**: Any new technology raises certain risks, including the possibility of misconfiguration or manipulation. To aid in diagnosing and fixing operational problems, we recommend that every RPKI repository employ accurate and effective human contact records, and that corresponding addressing authorities keep their allocation data up to date and synchronized with corresponding RPKI data. That is, it should be possible to easily correlate RPKI data with existing resource allocation data including the identities of resource holders that are decoupled from RPKI records. We also recommend that organizations responsible for operating RPKI databases and managing certificates make available tools to detect possible configuration errors and expiring certificates, as well as flag suspicious changes that may stem from abusive manipulation of the data. The policies of RPKI operators should allow open access to the RPKI record databases and permit general dissemination of the data and derived results (e.g., lists of validated origins), without formal approval.

4 I**mproving BGP security metrics and measurements**: Quantitative analysis of BGP measurement data is crucial for informing decisions about which security solutions to deploy. The BGP security community should evaluate the security metrics used in previous measurement studies, and extend them where necessary. The community should perform continuous monitoring and analysis of BGP, to quantify the importance of the problem and identify any changes in the frequency, severity or likelihood of security incidents. As RPKI deployments proceed, the community should track participation in the RPKI and gauge its effectiveness in preventing the spread of erroneous routing information.

While unanimity in recommendations was an objective from the outset, the conclusions of the report are not necessarily shared by all group members. Although Level 3 Communications participated in the CSIRC III WG 6 Task Force, it does not agree with the balance of recommendations expressed in this paper.

We also note that we intentionally stop short of attaching concrete time frames to our recommendations. Evolving threats and a better understanding of the operational challenges of RPKI would clearly have a major influence on the appropriate times to adopt the new technology, and preserving the autonomy of individual Network Providers and Autonomous System operators is critically important.

2. CSRIC Structure

In this section, we briefly summarize where Working Group 6 "Secure BGP Deployment" fits in the overall structure of CSRIC III working groups, and list the members of the working group.

2.1 CSRIC Structure

Working Group 6 on Secure BGP Deployment is a working group under the FCC's CSRIC III (Communications Security, Reliability, and Operability Council III), under the following structure:

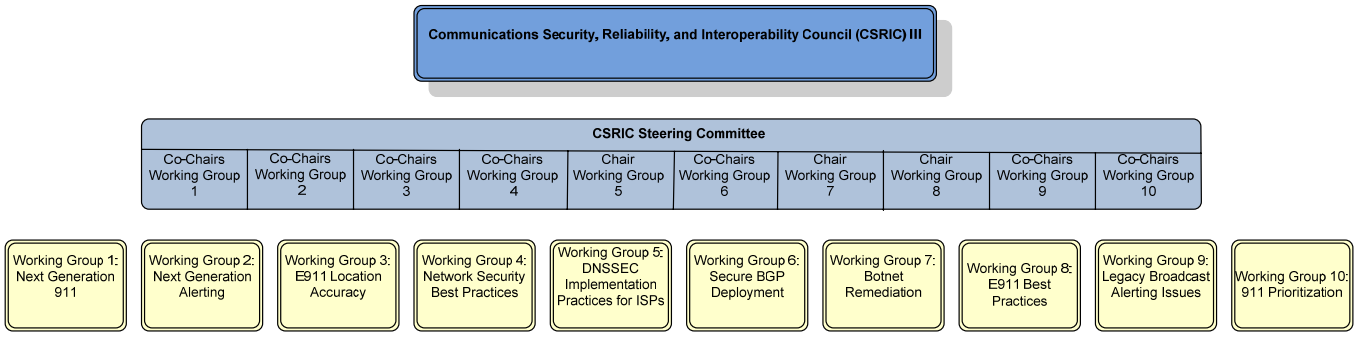

Figure 1 – CSRIC III Organization Chart

2.2 Working Group 6 Team Members

Name	Organization
Andy Ogielski, Co-Chair	Renesys
Jennifer Rexford, Co-Chair	Princeton University
Shane Amante	Level3
Daniel Awduche	Verizon
Ron Bonica	Juniper
Jay Borkenhagen	AT&T
Belinda Carpenter	Sprint
Martin Dolly	ATIS/AT&T
Mike Geller	ATIS/Cisco

Sharon Goldberg	Boston University
Adam Golodner	Cisco
John Griffin	TeleCommunication Systems
Kyle Hambright	Las Vegas Metro Police
Lars Harvey	Internet Identity
Michael Kelsen	Time Warner Cable
Ed Kern	Cisco
Padma Krishnaswamy	Battelle Memorial Institute
Eric Lent	Comcast
Doug Maughan	DHS S&T
Danny McPherson	Verisign
Doug Montgomery	NIST
Christopher Morrow	Google
Sandra Murphy	SPARTA
Mats Nilsson	ATIS/Ericsson
Eric Osterweil	Verisign Labs
Mary Retka	CenturyLink
Isil Sebuktekin	Applied Communication Sciences
Ted Seely	Sprint
Greg Sharp	Internet Identity
Tony Tauber	Comcast
David Ward	Cisco
William Wells	TeleCommunication Systems

Table 1 – List of Working Group Members

3. BGP Security Background

In this section, we briefly present background on interdomain routing security. The way BGP disseminates routing information from one AS to another makes it vulnerable to misconfiguration, malicious attacks, and software bugs. To protect against bogus routes, many network operators already follow Best Current Practices (BCPs) to defend their networks. Yet, establishing "ground truth" for which ASes can originate each IP prefix would make it easier for ASes to identify bogus routes that originate several AS hops away. As such, we discuss proposed solutions for establishing ground truth and using this information to detect or prevent the spread of bogus routes. In the appendices, we provide a more complete overview of BGP, and a glossary of key terminology used in this report.

3.1 Interdomain Routing Security Problems

The current implementation of inter-domain routing in the Internet relies on the Border Gateway Protocol version 4 (BGP v4). For purposes of BGP routing, all locally-owned networks (such as in a corporation) are grouped by their owners into administrative domains referred to as Autonomous Systems (ASes). Within each AS, IP routing among its internal networks need not be compatible with internal routing employed by other ASes; in fact, distinct ASes may internally use quite incompatible routing protocols. In contrast, BGP is a single standard protocol that provides *routing between distinct ASes,* making BGP a critical part of holding the Internet together. Therefore, immense care must be taken to preserve *robustness* of the BGP routing system when adding any new functionality, such as security extensions.

Each Autonomous System is assigned a unique numerical identifier (AS Number, or ASN). The BGP protocol allows each AS to advertise reachability for IP address blocks (or prefixes) in BGP protocol messages. In BGP, each AS learns routes to destination prefixes from neighboring BGP speakers, either internal BGP within the same AS, or external BGP speakers within remote ASNs, and chooses to use, discard, modify, and/or propagate those routes to other internal and external BGP neighbors according to the BGP path selection algorithms and local policies.

Either by accident or malicious attack, BGP routing messages with bogus attributes can be injected into the system by one BGP-speaking router, and easily propagate from one AS network to another, rapidly leading to serious global consequences ranging from "traffic blackholing" (where traffic never reaches its destination) to "traffic detouring" (where traffic is diverted to, or away from, a particular transit network). In addition, peculiar routing messages (either syntactically malformed, or correctly formatted but with unusual attributes) sometimes trigger bugs in certain BGP software implementations, causing instabilities such as reboot or failure, including cascading failures that can lead to wide-area meltdown[1]. Most security solutions focus on detecting bogus routing messages, rather than protecting against software bugs.

[1] J. Cowie and A. Hobgood, The Curious Incident of 7 Nov 2011, http://www.renesys.com/tech/presentations/pdf/nanog54-cowie.pdf

How common are BGP security incidents? It is very difficult to provide a definitive answer for the general problem. Here we focus on a few categories of incidents that are addressed by currently considered BGP security extensions.

First consider unauthorized route origination, popularly called "route hijacking". In such incidents, routes to a given network prefix (or many prefixes) or their more specifics are originated by an AS that has not received authorization to do so from the legitimate owner(s) of these prefixes.

How common are such hijacking incidents? Confirmed large scale mis-originations involving many routes concurrently 'hijacked' from a large number of distinct origin ASes and propagated worldwide happen on average about once or twice per month, lasting from a few minutes to a few hours per incident[2]. Some generate large publicity[3] while others pass with little public notice[4] regardless of the scale of the actual impact. Such large hijack events are in principle relatively easy to detect, because one assumes that there can be *no legitimate intent* for such a large change. Still, without additional information it is generally impossible to tell which such events resulted from accidental misconfigurations or errors, and which were malicious in nature.

Much smaller potential hijack events (impacting one AS and one or a few prefixes) are not noticeable by their size alone. Changes to long-standing, stable route originations might provide some clue of a problem, but in general such incidents could result from a variety of intentional (e.g., traffic engineering) or un-intentional (e.g., misconfiguration, purposeful attack) causes. True route-hijacks that affect reachability to popular (e.g., the widely-publicized YouTube incident[5]) or business-critical destinations are often quickly detected by end users or the service providers themselves. Upper bounds on the number of such incidents depend on the details of detection criteria: Recent research[6] found 4,000 candidate events in 2009. This order of magnitude is roughly consistent with the number of multiple-origin prefixes tracked by Team Cymru[7]. It is important to stress that many origin changes are *legitimate*, and some are executed on a very short notice for security purposes, for instance denial-of-service (DoS) attack defense by a security vendor such as Prolexic, Verisign, and others.

The fact that the scale and rate of legitimate route origination changes can vary widely limits the utility of statistical anomaly detection techniques for security purposes. The nature of the problem requires some means of establishing accurate "ground truth" about who owns what address resources and which autonomous systems are authorized to announce routes to them.

---

[2] Varun Khare, Qing Ju, and Beichuan Zhang, Concurrent Prefix Hijacks: Occurrence and Impacts, in *Internet Measurement Conference*, November 2012. http://www-net.cs.umass.edu/imc2012/papers/p29.pdf

[3] James Cowie, http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml

[4] A. Popescu, B. J. Premore, T. Underwood, The Anatomy of a Leak: AS9121, http://www.renesys.com/tech/presentations/pdf/renesys-nanog34.pdf

[5] Martin A. Brown, http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml

[6] Varun Khare et al., op. cit.

[7] http://www.team-cymru.org/Monitoring/BGP

How common are other BGP security incidents? Routes with bogus AS_PATH attribute are seen more frequently, but unlike bogus origin AS they need not deny reachability to affected network prefixes. One particularly long-running study[8] has been recording advertisements of routes with AS_PATH containing three consecutive "tier-1" AS numbers. Such AS_PATHs are considered bogus (and likely unintentional), since they would cause the middle AS to transit traffic between the other two at no cost and most likely against its business policy. This study collected 9.6 million such anomalous route announcements since 24 January 2008, but their actual impact was not quantified.[9]

Beyond the AS_PATH, BGP route announcements have a number of other attributes; each designed to influence the decision of the next-hop BGP router how to process such a route. In principle, each attribute may be bogus for purposes of non-legitimate manipulation of routing decisions. We did not find a systematic study of the prevalence of other bogus attributes.

Regardless of the frequency of incidents caused by bogus routes, there are justified concerns among Internet service providers, critical infrastructure operators and authorities charged with communication network regulation and security that *the very fact that such vulnerabilities exist in global inter-domain routing system presents a significant danger*, and efforts should be made to design and implement effective and internationally acceptable mitigation mechanisms. The foundation of current thinking about BGP security is that one can create databases containing sufficient information about *who is authorized to route what and when*, and that network operators can use such databases to distinguish legitimate routes from bogus routes. The actual defense mechanism is filtering of routes by BGP routers, with bogus routes being rejected. However, only a small portion of such a complete solution has been addressed to date. At a high level, any implementation of such a strategy can be broken into its major components:

- A database of "ground truth" for storing information about allocation of resources (IP address blocks, AS numbers) and routing policies for entities participating in BGP routing.
- Mechanism for distribution of assertions about legitimacy of routes to network operators.
- Mechanisms for use of such assertions by network operators for the creation of route filters.

Today's approach is based on principles compiled in Best Current Practices (BCPs), as described in more detail in the next subsection.

[8] http://puck.nether.net/bgp/leakinfo.cgi
[9] https://fileshare.tools.isoc.org/robachevsky/public/2012-RRMWS/RRMWS-Jared-Mauch.pptx

3.2 Best Current Practices (BCPs)

The security and robustness of the Internet routing system depends on availability of accurate information about the authorized holders of IP address blocks and ASNs. This information provides a basis for declarations of routing policies and constraints. Examples include declarations that a given ASN is authorized to originate or transit a route to a given IP Address block. In this subsection we discuss the state of existing systems (i.e., Regional Internet Registry (RIR) "whois" databases and multiple Internet Routing Registries (IRRs)), and their use in BCPs.

In today's Internet, interdomain routing security relies on AS operators configuring their BGP routers to filter at least some categories of bogus routes (e.g., routes with a bogus origin AS or AS_PATH). However, route filtering is predominantly deployed near the "edge" of the Internet, that is, on the BGP sessions between Internet Service Providers (ISPs) and their customers, and, recursively, their customers' customers. Enforcement of the legitimacy of routing information that originated several AS hops away, especially on links to AS neighbors who are not customers but transit peers, is much more difficult and not typically implemented.

The reason for preference for route filtering on customer sessions but not on other BGP sessions is that a network service provider can reasonably accurately verify which routes received from a customer AS have legitimate values of origin AS or AS_PATH, and which are bogus. Such determination is generally not possible for routes received from peer ASes because accurate information about what constitutes legitimate origin or AS_PATH for a non-customer network several AS hops away is generally not available.

For customer sessions, an ISP can prevent propagation of bogus routes by "whitelisting" allowed routes. This is done by maintaining a local mapping of customer-assigned prefixes to each customer ASN in order to provide some protection against accidental or intentional route mis-origination by customers. To validate that the customer organization has actually been assigned the IP address blocks (prefixes) they wish to announce, an ISP can consult the "whois" database of a RIR. In addition, an ISP can recursively build whitelists for prefixes that may be originated by a customer's customers, to allow the customer to propagate routes to these prefixes as well. The ISP may require each customer to provide a list of the AS numbers of its downstream transit customers, or use IRRs to make this information available.

Some ISPs may employ only AS_PATH filters and not prefix filters. This practice is considered inadequate, since some common routing configuration errors can allow accidental re-origination (vs. re-distribution) of BGP routes learned from another network. In that case, the customer's ASN may show up as the origin *AS* of those routes and, again, in the AS_PATH.

The success of these filtering techniques depends on the accuracy, completeness, and timeliness of the IP allocation and assignment record keeping, as well as on diligence of ISPs in constructing current and complete route filters from that information base and associated with the corresponding customer / peer AS(es). Inaccurate records manifest themselves in various

ways. One simple way is that the organization name input to the RIR records, being a free-form field, does not clearly match the name provided by the customer, leading to ambiguity in the verification process. For instance, the RIR records may refer to "ABC Corp" but the ISP customer record may identify the customer as "Acme Beer Company Inc." Thus some human judgment by ISP staff is inevitably used in this process. Another source of errors are often introduced by either typographical or dictation errors, or transposed digits. For instance 192.10.2.0/24 rendered as 192.100.2.0/24.

Route filters can be generated and maintained either manually or automatically. One common automation technique is to use information stored in an IRR (Internet Routing Registry) as a means of generating filters (see www.irr.net for a listing of IRRs). Several IRRs are operated by the RIRs (ARIN, APNIC, RIPE) and the others are operated by various ISPs and other third parties, with varying degree of mirroring. The IRR records may be used as a way to check the list of prefixes a customer is asking the ISP to accept via BGP. Alternately, the customer may indicate the proper set of objects to query in the IRR for the purpose of automating filter generation. Examples include (i) Origin AS (i.e., query the IRR repository database for prefixes with a given value in the "origin" field) and (ii) AS-set (i.e., iterate through a set of ASNs or nested AS-sets keying off each AS number as an origin value). Each approach produces a list of ASNs and address blocks which can form a "whitelist" for the router configuration.

IRRs can be employed as part of an automated configuration scheme; a typical method is to run queries at regular intervals (e.g., some number of hours or days apart). Any unsupervised, fully automated scheme should provide for checking for potential errors in IRR records. The results of each run would compare the resulting list with the previous list, and then update route filters on routers if necessary. Since this method relies on published data, the integrity of this data is important. The IRRs use RPSL (Route Policy Specification Language) authentication and authorization model (see RFC 2725). The RPSL authorization generally follows the address-assignment scheme but is not always implemented (including by ARIN, the RIR serving the North American region). Most IRRs support one or more types of transactional authentication which can give a hint about who added a given record, but not an *authorization* scheme to judge the authority of a user to manipulate a record about a given resource.

The remediation of BGP routing security failures, such as stopping of route hijacking (malicious or accidental), or suppression of AS_PATH manipulation, necessarily involves communications among distinct network providers, often in different countries or on different continents. It is currently exclusively based on ad-hoc communication channels (such as Internet Relay Chat, mailing lists, email, etc.) within the global community of experienced network engineers operating BGP routers in the largest networks. Their devotion to maintaining global reachability has been demonstrated many times. However, there are serious concerns that such a remediation mechanism may not scale to meet the reliability needs of global businesses and other organizations as the size of the Internet continues to grow.

In summary, current BGP security practices suffer from a number of limitations, many of which can be traced to the absence of complete and accurate “ground truth” about authorized holders

of number resources and in particular who is authorized to announce which routes. A partial list of such shortcomings would include[10] [11] [12]:

1. Lack of globally unique, meaningful resource holder names and authentication procedures for IRRs
2. Lack of enforcement of an authorization model for registry changes
3. Insecure updates to the IRRs leading to proxy registrations, inaccurate data, etc.
4. No mechanisms to expire stale records, leading inconsistent data

Further improvements to BGP security practices have to rely on having a common notion of high quality, up-to-date "ground truth" about which AS(es) can announce routes for each IP prefix, and will depend on network operators actually using this information to detect or prevent the spread of invalid routing information. However, the IRRs will continue to play an important role in BGP security practices in the foreseeable future, and are not mutually exclusive with RPKI or similar approaches. Improving quality of the IRR infrastructure continues to be an important goal.

3.3 Cryptographic Approaches to BGP Security

A more recent strategy to enhance BGP security employs cryptographic techniques based on a public key infrastructure for certification of resource ownership and validation of route attributes. Organizations participating in BGP routing can acquire cryptographically protected certificates that prove they have been assigned a given Internet number resource, such as an IP address block or Autonomous System number, and develop a robust infrastructure for certificate handling. The use of a specially designed public key infrastructure known as the Resource Public Key Infrastructure (RPKI) currently has momentum, as it provides a way for third parties to formally verify assertions related to Internet number resource holdings, and can enable resource holders to bind prefix and origin AS information in a manner that can be formally validated by third parties. There are also other proposals for securing BGP with other cryptographically protected infrastructures, such as using DNSSEC to represent route origination information in the reverse DNS[13][14], but these are still in the experimental and research stage preceding possible IETF standardization in the future.

[10] D. McPherson, S. Amante, E. Osterweil, L. Blunk, IRR & Routing Policy Configuration Considerations, Internet-Draft (2012)
[11] A. Khan, Hyun-chul Kim, Ted "Taekyoung" Kwon, How Complete and Accurate is the Internet Routing Registry (IRR)?, In 4th CAIDA-WIDE-CASFI Joint Measurement Workshop, http://www.caida.org/workshops/wide-casfi/1112/slides/caidawidecasfi1112_akhan.pdf
[12] Richard Steenbergen, What's Wrong with IRR , http://www.nanog.org/meetings/nanog44/presentations/Tuesday/RAS_irrdata_N44.pdf
[13] J. Gersch et al., ROVER BGP Route Origin Verification via DNS, http://www.nanog.org/meetings/nanog55/presentations/Tuesday/Gersch.pdf
[14] S. Amante et al., Resource Certification Using DNS + IRRs, http://www.nanog.org/meetings/nanog55/presentations/Monday/Amante.pdf

The IETF's Secure Inter Domain Routing (SIDR) working group is addressing two specific categories of threats to the global BGP routing infrastructure:

- Bogus origin AS: when an AS originates a BGP route for an address block that has not been authorized by the legitimate network address block holder (route hijacking);
- Bogus AS_PATH attribute that can result from unauthorized insertion or deletion of ASNs to falsify the recorded sequence of ASNs transited along the route, to encourage or discourage other ASes from selecting the route.

Origin validation is amenable to incremental deployment, since the technology does not require any changes to the BGP protocol itself. In addition, origin validation protects against bogus routes caused by accidental misconfiguration (e.g., an AS inadvertently originating the wrong IP prefix), and is represents a first step toward more complete security solutions. As such, the rest of this report focuses primarily on establishing "ground truth" for which ASes can originate each IP prefix, to enable origin validation.

While origin validation handles the simplest BGP security incidents, route leaks[15] and more sophisticated attacks that manipulate the sequence of ASes in the BGP AS_PATH attribute may still be effective, as long as the attacker ensures that the origin AS matches an authorized origin for the prefix in question. While the SIDR working group is designing and standardizing techniques for cryptographic protection of the path information, the standardization process for path validation is not yet complete, and implementations of the technology are not yet available. In terms of cost, path validation would require software upgrades to the routers to perform the necessary cryptographic operations. Router hardware upgrades may be necessary as well, since the signing and validating of routes requires additional processing and memory. Current discussions indicate that the future of BGP security extensions to cryptographic path validation is not yet settled.

3.3.1 Resource Public Key Infrastructure (RPKI)

The Resource Public Key Infrastructure (RPKI)[16] provides a way for the holders of Internet number resources (that is, IP addresses and AS numbers) to formally demonstrate their right to use these resources, to bind prefix and origin AS information, and for third parties to verify assertions made about these resources. The RPKI uses restricted X.509 v3 certificate profiles, with a number of specially designed extended attributes chosen for routing security purposes.

The RPKI is a hierarchy of certificate authorities intended to map onto the existing administrative hierarchy for allocating the number resources (from ICANN/IANA to RIRs, RIRs to NIRs or LIRs, NIRs or LIRs to their customers and so on). These certificate authorities can sign other certificates and attestations, or revoke the issued certificates if they wish. Certificates will also expire after a period of time if they are not explicitly renewed. The primary purpose of

[15] D. McPherson, S. Amante , E. Osterweil, Route Leaks & MITM Attacks Against BGPSEC, Internet-Draft (2012)
[16] See RFCs 6480-6488, 6492, 6493.

the RPKI is to validate the authority to use IP number resources, not the identity of individuals or organizations responsible for these resources. For IP address blocks, the address block holder would sign a Route Origin Authorization (ROA) certificate authorizing a designated ASN to originate a route for the IP address block.

All Regional Internet Registries (RIRs) are now offering RPKI services to their members; in North America, the American Registry for Internet Numbers (ARIN) started a pilot RPKI service in 2009, and an initial operational service in 2012. To date, all five global RIRs offer what is known as a "hosted" model. While most RIRs have plans to roll out a full delegated model, that has not happened to date.

Like the currently operating IRRs, the RPKI relies on participation of network operators to populate the repository databases. Yet, the RPKI has several advantages over existing IRRs that can help make the data more accurate and complete. It is relevant that RPKI could be used to improve the existing IRRs, which could quite rapidly improve interdomain routing security by the virtue of the IRRs role in Best Current Practices.

- First, today's IRRs do not have a common, strong authorization model. As a result, change control and authorization in many IRRs is difficult to enforce or verify, leading to questionable data. In contrast, the RPKI has a stronger trust model, where only a party whose authorization is verifiable via the chain of trust can make declarations on behalf of the holder of a specific number resource.
- Second, IRRs operated by parties other than RIRs have no reliable means to validate if an IRR registrant is the authorized holder of the number resource being registered. But this can remedied by RPKI.
- Third, the existing routing registries have no mechanism to automatically delete, or even to identify, stale information. In contrast, the RPKI can remove expired information, or revoke certificates when number resources are re-assigned or removed from allocation as the necessary information and mechanisms are implicit in the data objects and associated infrastructure operations.

The RPKI is not a replacement for the whois databases or all of the capabilities of IRRs. In particular, the identifying information (abuse contacts, etc) and information needed for conduct of business (organization names and addresses, etc), will continue to be needed but is not provided by the RPKI. Additionally, the IRRs index their data in a way that is amenable to enumeration of prefixes with a common origin attribute, which is key to their use for generating prefix-filters on BGP sessions. In short, RPKI has not been designed for storing routing policies for use in building customer filter trees filtering like the IRRs.

3.3.2 Origin Validation Using RPKI Data

A ROA authorizes an AS to originate BGP update messages for one or more address blocks. Thus any BGP update message received can be compared to the set of validated ROAs (i.e., those whose signatures have a valid certificate chain leading to the configured trust anchor). This process, known as "prefix origin validation," results in an update being categorized as:

- ***valid*****:** the origin AS matches a valid ROA for the IP prefix;
- ***invalid*****:** the IP prefix has valid ROAs, but none match the originating AS;
- ***unknown*****:** the IP prefix is not described in any valid ROA, and is not a more specific prefix of any valid ROA.

While the meaning of the first two categories should be obvious, the "unknown" result exists to accommodate partial and incremental deployment scenarios in which the RPKI data to classify the route is not in the system.

How an AS or BGP router uses RPKI information (e.g., to build offline filters, influence local preferences, or to flag certain routes) is always a matter of local implementation and routing policy. For routes that are found to have invalid origins, example reactions might include:

- generate a notification to network operations staff;
- assign a lower preference to the route in best path selection; or,
- reject the route.

These example policies demonstrate a range of responses from strictly advisory to directly affecting the reachability of an address block. It is expected that ISPs would adopt use of RPKI techniques in a staged manner, progressing along this range of actions as trust and confidence in the underlying RPKI and registration information systems increase.

4. Recommendations

Any deployment of RPKI-based origin validation should be based on a "ground truth" of which ASes can announce each IP prefix. This section discusses four main ways to achieve that goal. First, we discuss ways to improve the accuracy and completeness of public and certified records of numbered resources. Next, we outline how ASes can pursue a staged deployment of origin validation, based on participation in RPKI and "out of band" usage of RPKI data as part of existing operational practices. Then we outline some of the new risks introduced by hierarchical resource certification systems, as well as possible ways to mitigate those risks. Last, we identify security metrics and associated measurement techniques that can track the deployment of RPKI and origin validation, detect configuration mistakes and intentional manipulations of the RPKI, and generally raise the confidence of network operators in the accuracy and robustness of the RPKI.

4.1 Accurate Records About Number Resource Holders

All techniques for improving the security of interdomain routing rely on authoritative, accurate, and verifiable information about which Autonomous Systems are authorized to announce routes to which IP Address Blocks. To have a firm notion of "ground truth," AS operators should publish accurate and timely information about their IPv4 and IPv6 address blocks in Internet Routing Registries (IRRs). More accurate IRRs would naturally improve on existing routing security practices during the incremental deployment of RPKI, while also capturing important information (such as information about organization name of a resource holder and human points of contact) that may not be present in RPKI records. To have truly verifiable information about numbered resources, North American AS operators can certify their numbered resources using the RPKI service offered by ARIN. To prevent conflicting RPKI records under different RIRs, AS operators and the larger BGP security community should encourage implementation of the single global "root of trust" envisioned for the RPKI. Otherwise, different parts of the RPKI could store inconsistent records concerning the same numbered resources, and that would compromise the reliability of (and confidence in) the RPKI.

AS operators should ensure their Internet Routing Registry (IRR) records are public, complete, and up-to-date: Having a common, public notion of "ground truth" for identifying bogus routing information is a prerequisite for all BGP security solutions. We recommend that every AS operator ensure its IRR records are public, complete, and up-to-date. ISPs that assign portions of their address space to customers should register and maintain accuracy of these address assignments in IRRs as appropriate. Although many IRRs mirror one another, there is really no authoritative list of "primary" registries, and an AS operator may need to input their data to multiple IRRs to ensure their information is widely available. This still does not address how conflicting records should be resolved among IRRs. It would be useful to develop a specification of how/where to do this in the IRRs, what parts can be validated, and who and how would validate them.

AS operators can take these steps without delay, since they do *not* depend on learning or deploying any new technologies or systems (such as RPKI). It is important to note that the emergence of RPKI does *not* diminish the need for the IRRs, since (i) many AS operators use IRR data to construct route filters as part of today's best current practices and (ii) IRR data includes important information, such as organizational identity, address and point of contact information, that is *not* available in the RPKI but will continue to be required for interdomain routing security.

AS operators should use the Resource Public Key Infrastructure (RPKI) to certify their numbered resources: Looking forward, the RPKI offers a more rigorous and verifiable system for number resource certification. AS operators should begin to certify their number resources. In the case of North America, ARIN currently offers a hosted RPKI service, and intends to offer a fully delegated RPKI service at a later date. AS operators can use ARIN's RPKI service to generate (i) resource certificates that cryptographically verify that a resource has been allocated

or assigned to a specific entity and (ii) Route Origin Authorizations (ROAs) that specify which AS(es) may originate a specific IP address or range of addresses.

AS operators and the larger BGP security community should encourage a single global "root of trust" for the RPKI: The working group recognizes that, during the early stages of deployment, the RPKI may have multiple trust anchors ("roots"). However, for a global resource certification to be most effective, AS operators should stand with the Internet Architecture Board (IAB)[17] in encouraging the Regional Internet Registries (like ARIN) and the Internet Assigned Numbers Authority (IANA) to create a *single* root trust anchor for the RPKI, and ensure that the trust anchor is strictly aligned with the Internet number resource allocation hierarchy. Otherwise the onus for conflict resolution would fall to AS operators, who have no capability to resolve these conflicts. Furthermore, with no global root, any of the multiple trust anchors could assert, either intentionally, via error or through compromise, holdings of resources they have not been allocated. Having a global root would also help resolve issues that no one RIR can handle alone. An important example is the large amount of "legacy" IPv4 address space that was allocated before the RIRs were created. The RPKI root should provide RPKI services for the legacy IPv4 address holders.

4.2 Cautious, Staged Deployment of RPKI Origin Validation

Securing the global routing system using data from a resource certification system such as RPKI adds a new large and complex component to the list of factors on which the operational routing system must depend. Integration of this new infrastructure component to the global routing system unavoidably creates new failure modes, thus any deployment process must strive not only to increase security of the existing infrastructures, but also to reduce opportunities for creation of new vulnerabilities. It is important to note that even ASes that do not use the RPKI data would be affected by the way *other* ASes use this information to select and disseminate BGP routes.

The major dependencies exhibited by the currently proposed resource certification system (RPKI) are as follows:

1. Dependencies within the existing hierarchical system of business relationships for allocating IP addresses and AS numbers.

> Currently, the Internet Corporation for Assigned Names and Numbers (ICANN) and the Internet Assigned Numbers Authority (IANA) allocates large blocks of globally unique IP addresses to five Regional Internet Registries (RIRs), which cover all regions of the world. Subsequently, the RIRs delegate large portions of their address space to National Internet Registries (NIRs) or Local Internet Registries (LIRs) that include the ISPs, content providers, or large enterprises. This hierarchy for delegating IP address blocks was created with two simple goals: (i) to enable aggregation of IP prefixes to make the

[17] http://www.iab.org/documents/correspondence-reports-documents/docs2010/iab-statement-on-the-rpki/

routing system scalable and (ii) to provide an ability to adapt IP address allocation policies for particular regions, based on the needs of a very broad, multi-national, and, in some cases, multi-regional constituency of ISPs, content providers, and enterprises operating within those regions. In short, IANA sits at the root of an overall number allocation hierarchy with RIRs just below IANA, followed by NIRs and LIRs/ISPs at the lower branches of the allocation hierarchy tree.

2. Dependencies in the distributed RPKI repository systems, or similar systems.

The RPKI is a specialized version of the X.509 PKI certificate system designed to map CAs to the hierarchy of organizations allocating and assigning number resources, as a natural extension of their current practices. However, the repositories of published RPKI objects may either follow the same IP address allocation and assignment hierarchy or be outsourced as a hosted service in the case that the organization at a particular point in the allocation or assignment tree cannot or does not want to operate the requisite systems themselves.

3. Dependencies on the distributed system for transforming the distributed RPKI repository records into routing policy.

This set of dependencies is comprised of processes or protocols for fetching the RPKI data (e.g., certificates, ROAs, certificate revocation lists, manifests, and so on) from the distributed repositories to AS operators for use in their implementation of BGP route filtering. These repositories are potentially maintained at each branch in the number resource allocation or assignment tree. The data retrieved is used by network operating organizations and subsequently transformed into local instructions used to make routing policy decisions on learned routes in the operator's network, at each router in their network. Each ISP, content provider, or enterprise may choose its own policies for using (or not using) the RPKI information in routing policies.

We also observe that currently not enough is known about scaling properties and performance of synchronization of the distributed RPKI repositories (using rsync or similar techniques).

A concern opened up by any new critical database system that may control interdomain routing is its vulnerability to Denial of Service attacks. Protection is provided by significant redundancy, analogously to redundancy and use of BGP anycast for key DNS servers, such as root servers. This will significantly increase the cost of deploying RPKI and similar security solutions.

The risks exhibited by critical dependencies in the combined RPKI and BGP infrastructure bring two recommendations.

- **Independently of the RPKI infrastructure becoming available, network operators should not give up maintaining their local autonomy, and should remain solely**

responsible for setting their own routing policies, just as they are today (without RPKI).

- **Network operators should delay configuring their BGP routers for automatic rejection of routes deemed invalid in RPKI caches until they have significant trust in the accuracy, incorruptibility and reliability of RPKI infrastructure.**

In particular, it is always assumed that the action a BGP speaking router should take upon receiving a route that is invalid according to the RPKI is a local policy decision. Network operators may choose to apply the RPKI data in their network policies in whole or in part, and there should be no compulsory requirement for AS operators to use the RPKI data in routing decisions. It is expected that stages of RPKI adoption would follow the establishment of trust and confidence in the operation of the global RPKI system. The manner in which RPKI records are used by network operators is expected to evolve over time.

Importantly, a staged deployment not only allows network operators to gain experience with RPKI, but also substantially reduces the risk that naive or intentional corruption of records in RPKI would cause reachability problems.

4.3 Mitigating Risks Inherent in RPKI

Widespread participation in the RPKI could make the Internet routing system more secure by providing accurate information about which ASes are authorized to originate routes for each IP prefix. However, if the RPKI becomes a critical piece of the routing system infrastructure, inaccuracies in the RPKI data have the potential to impact the reachability of IP address blocks on the Internet, including those address blocks within which the RPKI repositories reside, as well as systems that enable access to the RPKI such as the DNS. In particular, the RPKI introduces new risks of (i) increased operational complexity and potential for misconfigurations and (ii) deliberate manipulations by parties that want to target specific Internet Service Providers and their customers. These risks could adversely affect the openness, robustness, and reliability of Internet communication in manners for which they are not vulnerable today. It is important to note, however, that misinformation in the RPKI will not *necessarily* harm reachability of IP address blocks, depending on how ASes use the RPKI data in their routing decisions. The working group has several recommendations to help mitigate these risks.

4.3.1 Operational Complexity and Misconfiguration

The current interdomain routing system is fairly complex. The RPKI introduces a new layer of complexity to the system that gives rise to the following two concerns:

1. **Building expertise:** Developing the expertise that is required to run the RPKI may well render the overall routing system more prone to failure and decrease connectivity overall, at least in short run, while operators gain experience with the new systems and operational procedures required to transition to RPKI. While this is true of any transition to a new network technology (e.g., IPv6, DNSSEC, etc) there is a particular concern that the deployment of the

RPKI requires operators to maintain a certificate authority for their resources (or rely on hosted models), a function that requires the development of new expertise, specialized systems, and operational procedures that are outside the typical area of expertise of most operators. Indeed, the complexity of running a certificate authority suggests that some organizations might move to a `hosted model' for RPKI where maintenance of their certificate authority is outsourced to a third party; the control placed these third parties is another subject for concern, as discussed in Section 4.3.2.

2. **Misconfigurations:** Like any system, the RPKI may be subject to misconfigurations and mistakes when data is input and interpreted. Thus, there is a longer-term concern that misconfigurations of the RPKI can impact IP address reachability. Two misconfigurations of particular concern include (1) failing to issue a new certificate before an existing one expires, and (2) issuing a new ROA for an IP address block that causes routes for its sub address blocks (that do not yet have ROAs) to become invalid. (The latter misconfiguration is quite common in early deployments of the RPKI.[18])

4.3.2 Deliberate Manipulations

IP address reachability can be impacted by deliberate manipulations of RPKI data. The RPKI operates in a hierarchical manner, enabling parties to digitally sign statements that they are delegating IP address blocks or authorizing ASes to originate IP address blocks into BGP. In both cases, the signing party must show a valid statement indicating that it was delegated (a superset of) this IP address block. It is possible for any *single* party above a target IP prefix in the hierarchy to *unilaterally* revoke that target IP prefix, potentially compromising its reachability. Moreover, the revoking party may not, necessarily, be in the same legal jurisdiction as the target IP prefix; this depends on the IP address allocation hierarchy, as well as the entity that controls the RPKI root keys. This is discussed at length in a recent study[19]; here we summarize the following: (i) the relationship between RPKI manipulations and existing techniques for delegating and blocking access to IP address blocks, and (ii) how manipulations of the RPKI might be performed, detected, and mitigated. As we discuss below, these manipulations could be the result (i) of normal business practices (e.g. reclaiming address space from a rogue customer), or (ii) of law-enforcement, government, or non-governmental pressure group wishing to block access to certain IP addresses, or even (iii) a party that abuses its position in the RPKI hierarchy.

It is recommended that a separate study be devoted to potential vulnerabilities of networked Critical Infrastructure and Key Resources (CIKR) to evaluate national security risks in this context.

[18] Matthias Wählisch, Olaf Maennel, Thomas C. Schmidt, Towards Detecting BGP Route Hijacking using the RPKI, http://conferences.sigcomm.org/sigcomm/2012/paper/sigcomm/p103.pdf

[19] Kyle Brogle, Danny Cooper, Sharon Goldberg, Leonid Reyzin, Impacting IP Prefix Reachability via RPKI Manipulations, 2013, http://www.cs.bu.edu/~goldbe/papers/RPKImanip.pdf

4.3.2.1 Comparison to the Status Quo

Business disputes: As transfers and leases of IP addresses become more common, business disputes may arise between the delegator of number resources and the recipient of the resource. For example, an ISP A may wish to take back number resources that it leased to ISP B because the lease has expired, or because B has failed to honor their contract with ISP A. Alternatively, ISP A may want to terminate the lease early in violation of its contract with ISP B, or because a customer of ISP B is hosting a website that is critical of ISP A. Note that ISP A does not even have to be a provider of ISP B, only the assigner of IP address resources. Today, if ISP A wants to take back its number resources, it would have to resort to non-technical means like contacting ISP B's new provider(s) and asking them not to route the prefix, or even direct negotiation with ISP B, and perhaps litigation as a last resort. On the other hand, we note that this this capability of the RPKI may appear attractive to ISPs frustrated with the ineffectiveness of current means for reclaiming address space from former customers or customers in violation of their agreement.

Blocking access to resources: In any jurisdiction in the world, a law-enforcement organization, governmental agency, or a non-governmental pressure group may desire to limit reachability for a particular resource on the Internet. There are numerous ways to effect this change. Perhaps the most common is intra-AS or intra-country censorship (e.g., traffic interception, TCP resets, or interception of DNS responses), but these approaches are intended to impact only the Internet users in the censoring AS or country. There have also been numerous high-profile cases of DNS takedowns - when records are removed from DNS name servers - that have a global impact, preventing a target domain name from being resolved to its IP address by users anywhere in the Internet (see for the example, the "bodog" incident[20]). However, during a DNS takedown, the target is still reachable by its IP address (at least by savvy users that know to specify the IP address explicitly). Moreover, the target may change its domain name in order to use a DNS server in another jurisdiction.

In contrast, routing-based methods for blocking access to Internet resources are arguably more drastic, because they can prevent an IP prefix from being reachable by large swaths of the Internet (see Section 4.3.3). Routing-based blocking can be accomplished today, and has been, through compelling a responsible ISP by a lawful order or otherwise to stop forwarding IP packets to the targeted network address. There is concern that the RPKI increases the risk and the impact of the routing-based blocking methods, for the following reasons:

1 Filtering at the ISP receiving the traffic or at some intermediate ISP along the way will leave the target accessible in other parts of the Internet and, therefore, reachable via Internet proxies.
2 A provider may be convinced to stop routing for its customer's IP address blocks, but a customer that has multiple providers, or that holds provider-independent IP address space will have other options.

[20] http://blog2.easydns.org/2012/02/29/verisign-seizes-com-domain-registered-via-foreign-registrar-on-behalf-of-us-authorities/

In contrast, the hierarchical nature of the RPKI means that a *single party* can impact the validity of any information lower in the RPKI hierarchy, and have *global* impact on all users that rely on the RPKI data. Moreover, that party may be in a different legal jurisdiction than the target.

4.3.2.2 Manipulations and Detectability

In this working group, we have found that a revoker can unilaterally impact that validity of any information below it in the RPKI hierarchy. Namely:

1. A *revoker* can unilaterally make any target IP prefix that it has previously delegated, *invalid* in the RPKI using small changes that are difficult to detect, and without explicitly employing certificate revocation lists.
2. The revoker can also unilaterally invalidate *any* target IP prefix below it in the RPKI allocation hierarchy. This invalidation can be "surgical" in the sense that it does not have to cause collateral damage to the validity of other address space, even if the target prefix is part of a bigger delegated address block. While the revoker's actions need not cause any other address space to become invalid, these actions may, in some cases, require significant refactoring of the RPKI, and therefore may be detectable by monitors.

Details of these manipulations are presented in Section 2 of a technical report.[21]
It is important to note, however, that even if a revocation and/or refactoring of the RPKI is detected, it will typically be difficult for a monitor to determine intent, e.g. whether the revocation was the result of normal business practices, or abuse, or a response to legal or extra-legal pressures to block access to a prefix.

4.3.3 Impact of Misconfigurations and Manipulations on IP Address Reachability

Revoking or misconfiguring a network IP address certificate or ROA in the RPKI does **not necessarily** compromise reachability for this address. The severity of the effects depends on how ISPs throughout the Internet *use* the RPKI data in their routing decisions. It is important to note that the RFCs do **not** specify exactly how parties should consume RPKI information (instead, they state that this is a matter of "local policy" at the ISPs that use the information). It is difficult to predict how ISPs will use RPKI information in their routing decisions. The following local policies seem plausible:

1. Using RPKI to allow a provider to decide whether or not to originate a prefix for its customer,
2. Using RPKI to assign a lower preference to BGP-learned routes classified as "invalid", or
3. Using RPKI to filter BGP-learned routes classified as invalid.

[21] Kyle Brogle et al., op. cit..

While RPKI invalidation can affect prefix reachability to different degrees in each of these three cases, the impact in all cases is typically quite nuanced and discussed in detail in the aforementioned technical report. The only exception is the case where routers use a strict ``drop invalid'' origin-authentication policy, so that invalidating a target prefix unequivocally prevents it from being reachable via BGP.

Given the early stage of the RPKI deployment, it is not possible to determine the policies that ASes will eventually use. For this reason, the impact of misinformation in the RPKI on prefix reachability is difficult to quantify, and warrants further study.

4.3.4 Recommendations for Reducing RPKI Risks

Mitigation of the risks of misinformation in the RPKI should be addressed by exploring means for monitoring and auditing modifications of the RPKI databases.

Transparency of RPKI Data: The RPKI purposefully does not standardize inclusion of data about the identity of resource holders in the resource certificates. In fact, certificate issuers may assign completely arbitrary and variable names to certificate subjects. As described to this group by people involved in RPKI design, the logic for this is twofold. First, duplicating the information already maintained in the RIR systems was thought to be redundant and subject to new potential inconsistencies. Second, inclusion of such information in RPKI certificates might lead to their misuse as identity certificates in contexts other than those intended. This latter issue was thought at the time to cause potential legal and liability concerns that could be undesirable for RIRs. Overall, such built-in lack of transparency could be a source of concern, if it is found that a tight synchronization between allocation data and RPKI data cannot be maintained. In order to build confidence in, debug and maintain this correspondence, it is important that human points of contact responsible for the operation of certificate authorities (CAs) in the distributed RPKI system are easy to identify. For this reason, we recommend that all CAs include RFC 6493 records that identify a point of contact for each CA (it is optional in the standards).

Detecting misconfigurations when entering RPKI data: AS operators can accidentally enter incorrect information into the RPKI, or neglect to renew certificates before they expire. We recommend that all RPKI CAs deploy tools that can help prevent RPKI misconfigurations. New information entered into the RPKI should be validated against the information already available in BGP routing tables, to detect when the addition of new RPKI information would make existing routes invalid. Tools that alert before a certificate expires will also play an important role in preventing misconfigurations. Other examples include alerts when the attestations are changed, whether for legitimate or spurious purposes. For example, RIPE already has a tool that alerts its customers of the potential impact of a requested ROA, to warn AS operators when a new ROA does not agree with the routing information seen in the operational BGP routing system.

Flagging suspicious changes in the RPKI data: Ongoing monitoring of the RPKI can help detect manipulations of certificates and ROAs. The monitoring should go beyond logging the revocation of certificates to include modifications and (re)issuing of ROAs. We recommend that creators of the RPKI software suites pay attention to the need for detailed logging and auditing, and that the operator and standards communities push for public sites that publish changes to the RPKI in real time, to aid in detecting and debugging problems. These sites can play a similar role for RPKI that BGP update collections (such as RouteViews and others) do for inspection and analysis of today's BGP routing tables - a way to gather data from multiple vantage points, compare the results against past history and other registries, and report the results. These sites can also play an important role in forensic analysis, through linking with other data sources (e.g. whois data). Correlation with external sources is particularly important given that it will be difficult for a monitoring and detection systems to determine the intent (e.g., normal business practice, abuse, takedown, etc) behind a refactoring or revocation in the RPKI.

Organizations managing the RPKI hierarchy should adopt policies that encourage the open dissemination of RPKI-related data: The RPKI data is valuable for improving BGP security, and for enabling analysis of the global routing system. Restrictions on sharing the data, or information derived from the data, could impede efforts to improve global routing security. The RIRs should offer free and open access to the RPKI data for researchers, AS operators, and others in the BGP community to perform analysis and disseminate the results without requiring formal approval.

4.4 Improving BGP Security Metrics and Measurements

Quantifying the security problems with BGP, and the effectiveness of proposed solutions, is important for informing decisions about which security solutions to deploy, and building confidence in the chosen solutions. The BGP security community - AS operators, researchers, RIRs and so on - should work to improve the quality of measurements and metrics available for evaluating the BGP routing system. Today, the absence of an accurate "ground truth" makes it difficult to detect and classify BGP security incidents accurately, and clearly identify the intent of the ASes involved. Still, the BGP security community can make progress on defining clearly-defined metrics for characterizing security incidents, and using these metrics to detect suspicious routes and their possible causes. The "ground truth" provided by RPKI can enable more accurate measurements. As more ASes participate in RPKI, the BGP security community can monitor the deployment of RPKI (and quantify any reductions in BGP security incidents), as well as suspected manipulations of the RPKI data.

The BGP security community should evaluate existing BGP security metrics, and extend them where necessary: Despite many years of experience with BGP, the community does not routinely use well-defined metrics for identifying routing security incidents, both accidental and malicious, and quantifying their scope and impact. Generally accepted metrics are crucial for calibrating the current levels of security problems, and evaluating the effectiveness of any proposed solution. Previous studies apply a wide range of methodologies, and typically study individual incidents or short periods of time. The community does not have broad consensus on

which security metrics and measurement methodologies are most effective. We recommend that the community evaluate and compare existing metrics and methodologies, and extend them where necessary, to better evaluate security problems and solutions. These studies should strive to understand the limits of our ability to infer the *intent* of an AS that originates incorrect information into the routing system, such as distinguishing an accidental misconfiguration from an intentional attack.

The BGP security community should perform continuous monitoring and analysis of BGP security incidents: The industry needs ongoing measurements of BGP security incidents. Since we cannot have a "flag day" for deploying any security solution, deployment will inevitably take place incrementally over a period of years. We recommend *continuous* measurement of security metrics, to record any changes in the number and severity of routing security incidents of various types with the growing number of networks deploying stronger routing security solutions, and periodic reporting of the measurement results. These studies can leverage long-established publicly available or commercial BGP update data warehouses collected from many routers worldwide, or develop additional techniques to augment current data sets.

The BGP security community should track the participation in RPKI: Confidence in the accuracy, completeness, and stability of the RPKI system is a prerequisite for any future coupling of RPKI data with real-time operation of the routing system. An important part of building confidence comes from monitoring the changes to the RPKI as more ASes participate and as numbered resources change hands. Throughout the deployment process, the RPKI data should be logged and audited to track deployment progress, verify the timeliness of maintaining ROAs, and compare the RPKI data with the BGP routes seen in the routing system. The BGP security community can also monitor any reduction in "unknown" or "invalid" routes for IP address blocks appearing in ROAs, to gauge whether and how the RPKI data is used to detect and prevent the spread of erroneous routing information.

Appendix 1: BGP Background

This section is intended for non-experts who have a need to understand the origins of BGP security problems. Participating in the global BGP routing infrastructure gives an organization some control over the path traffic traverses to and from its IP addresses (Internet destinations). To participate in the global BGP routing infrastructure, an organization needs:

- Assigned IP addresses, grouped into IP network addresses (aka prefixes) for routing.
- A unique integer identifier called an Autonomous System Number (ASN).
- A BGP router ready to connect to a neighbor BGP router on an Internet Service Provider's network (or another already connected AS) that is willing to establish a BGP session and exchange routing information and packet traffic with the joining organization.

Let's begin with the fundamental definition: A "route" in the present context means something very specific, namely a destination network address (IP prefix) and a collection of associated attributes that identify the sending router, and provide other information to intermediate routers to influence their decision how to direct the packet traffic to that destination. A BGP attribute frequently mentioned in this report is AS_PATH, which was originally meant to be a list of ASes transited by the route (for avoidance of routing loops), but now is often used for additional purposes. For router-to-router communication such "routes" are encapsulated in BGP messages sent over TCP connections.

The basic operation of BGP is remarkably simple – each BGP-speaking router can relay messages to its neighbors about routes to network addresses (prefixes) that it already knows, either because it "owns" these prefixes, or it already learned routes to them from another neighbor. As part of traveling from one border router to another, a BGP route announcement incrementally collects information about the ASes that the route "update" traversed in an attribute called AS_PATH. Therefore, every BGP route is constructed hop-by-hop according to local routing policies in each AS. This property of BGP is a source of its flexibility in serving diverse business needs, and also a source of vulnerabilities.

The operators of BGP routers can configure routing policy rules that determine which received routes will be rejected, which will be accepted, and which will be propagated further – possibly with modified attributes, and can specify which prefixes will be advertised as allocated to, or reachable through, the router's AS. In contrast to the simplicity of the basic operation of BGP, a routing policy installed in a BGP router can be very complex. A BGP router can have very extensive capabilities for manipulating and transforming routes to implement the policy, and such capabilities are not formally standardized, but instead, are largely dictated by AS interconnection and business relationships. Attributes of a route received from a neighbor can be transformed before a decision is made to accept or reject the route, and can be transformed again before the route is relayed to other neighbors; or, the route may not be disseminated at all.

All this works quite well most of the time – largely because of certain historically motivated trust and established communication channels among human operators of the global BGP routing system. This is the trust that a route received from a neighbor accurately describes a path to a prefix legitimately reachable through the neighbor ASes networks, and its attributes have not been tampered with. Notwithstanding the above, the "trust but verify" rule applies: Best Current Practices recommend filtering the routes received from neighbors. While this can be done correctly for well-known direct customers, currently there is no validated repository of the "ground truth" allowing for correct filtering of routes to all networks in the world. The Internet Routing Registries (IRRs) only partially play this role, because they are often incomplete and out-of-date.

Now observe that the BGP protocol itself provides a perfect mechanism for spreading malformed or maliciously constructed routes, unless the BGP players are vigilant in filtering them out from further propagation. However, adequate route filtering may not be in place, and from time to time a malicious or inadvertent router configuration change creates a BGP security incident: malformed or maliciously constructed routing messages will propagate from one AS to another simply by exploiting legitimate route propagation rules, and occasionally can spread to virtually all BGP routers in the world. Because some BGP-speaking routers advertisement policy advertise all local BGP routes to all external BGP peers by default, another example that commonly occurs involves a downstream of two or more upstream ASes advertising routes learned from one upstream ISP to another ISP – both the customer and the ISPs should put controls in place to scope the propagation of all routes to those explicitly allocated to the customer AS, but this is difficult given the lack of "ground truth" and controls automation. The resulting routing distortions can cause very severe Internet service disruptions, in particular effective disconnection of victim networks or third parties from parts or all of Internet, or forcing traffic through networks that shouldn't carry it, potentially opening higher-level Internet transactions up to packet snooping or man-in-the-middle attacks.

Despite all this, BGP is in fact a quite robust infrastructure. Any new routing security mechanism must not degrade the current state of BGP operation, although there may exist worst-case security violation scenarios that have not been observed "in the wild" so far. In practice, recovery from routing misconfigurations and security incidents is enabled by 1) technical feasibility of determining the source of the incident (in terms of identifying the offending ASN and/or prefix), and 2) Practical capability of identifying institutions (i.e., network operators) who can restore correct router configuration, or capability of identifying upstream Service Providers who can isolate the offending network. Transparency of ownership of IP address blocks and AS numbers allows verification of what is true, and repair what misplaced trust has let happen.

BGP Security Incidents and Vulnerabilities

In this section we classify the observed BGP security incidents, outline the known worst-case scenarios, and attempt to tie the incident to features of proposed solutions that could to prevent them. Many of the larger incidents are believed to have been the result of misconfigurations or mistakes rather than intentional malice or criminal intent. It has long been suspected that more

frequent, less visible incidents have been happening with less attention or visibility, as they were not recognized or topologically scoped as to not have been detected.

BGP security incidents usually originate in just one particular BGP router, or a group of related BGP routers in an AS, by means of changing the router's configuration leading to announcements of a peculiar route or routes that introduce new paths towards a given destination or trigger bugs or other misbehaviors in neighboring routers in the course of propagation.

There are no generally accepted criteria for labeling a routing incident as an "attack", and – as stressed in the recommendations – lack of broadly accepted routing security metrics that could automatically identify certain routing changes as "routing security violations". We recommend that research should be carried out to investigate this subject in greater depth, as provided in the main section of this report.

BGP security incidents that were observed to date can be classified as follows:

- **Route origin hijacking** (unauthorized announcements of routes to IP space not assigned to the announcer). Such routing integrity violations may happen under various scenarios: malicious activity, inadvertent misconfigurations ("fat fingers") that are likely most common causes of 'hijacking', or errors in traffic engineering. There are further sub-categories of such suspected security violations:
    - **Hijacking of unused IP space** such as repetitive hijacks of routes to prefixes within a large IP blocks assigned to an entity such as US government but normally not routed on the public Internet. Temporarily using these "unused" addresses enables criminal or antisocial purposes activities (spam, network attacks) while complicating efforts to detect and diagnose the perpetrators.
    - **Surgically targeted hijacks of specific routes and deaggregation attacks** on specific IP addresses. They may be hard to identify unless anomaly detection is unambiguous, or the victim is important enough to create a large commotion. Examples: Pakistan Hijacks YouTube[22] (advertisement of a more specific is globally accepted, and totally black-holes the traffic to the victim). There may be significantly more such attacks than publicly reported, as they may be difficult to distinguish from legitimate traffic engineering or network re-engineering activities.
    - **Unambiguous massive hijacks of many routes** where many distinct legitimate origin ASes are replaced by a new unauthorized origin AS advertising the hijacked routes. Significant recent incidents include a 2010 "China's 18-minute Mystery",[23] or a hijacking of a very large portion of the Internet for several hours by TTNet in 2004[24], or a 2006 ConEd incident.[25] Without knowing the

[22] http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml
[23] http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml
[24] Alin C. Popescu, Brian J. Premore, and Todd Underwood, "Anatomy of a Leak: AS9121," NANOG 34, May 16, 2005.
[25] http://www.renesys.com/blog/2006/01/coned-steals-the-net.shtml

motivations of the implicated router administrators it is difficult to determine if these and similar incidents were due to malicious intent, or to errors in implementations of routing policy changes.

- **Manipulation of AS_PATH attribute in transmitted BGP messages** executed by malicious, selfish, or erroneous policy configuration. The intention of such attacks is to exploit BGP routers' route selection algorithms dependent on AS_PATH properties, such as immediate rejection of a route with the router's own ASN in the AS_PATH (mandated to prevent routing loops), or AS_PATH length. Alternatively, such attacks may target software bugs in distinct BGP implementations (of which quite a few were triggered in recent years with global impact).

    - Unintentional leaks (announcement of invalid routes due to operator error) are the most common problem and often the ones with the largest impacts.[26] It may not be helpful to the victims to learn later that no malice was intended, and furthermore a large leak may be used to obscure a targeted attack.
    - A possibility of "man in the middle" (MITM) AS_PATH attacks detouring traffic via a chosen AS was publicly demonstrated at DEFCON in 2008[27]. Two other similar incidents were found in a 7-month period surrounding the DEFCON demo by mining of a BGP update repository conducted in 2009[28] but were not confirmed as malicious. This can occur either by accident as detailed above, or may be intentional. Additionally, such attacks may or may not attempt to obscure the presence of additional ASes in the AS path, should they exist. These are particularly problematic to identify as they require some knowledge of intent by the resource holder and intermediate ASes.
    - For routing incidents triggered by long AS_PATHs see House of Cards[29], AfNOG Takes Byte Out of Internet[30], Longer is Not Always Better[31] for actual examples.
    - AS_PATH poisoning – sometimes used by operators to prevent their traffic AS from reaching and/or transiting a selected AS, or steer the traffic away from certain paths. It is technically a violation of BGP protocol and could be used harmfully as well.
- Exploitations of router packet forwarding bugs, router performance degradation, bugs in BGP update processing
    - Example of a transient global meltdown caused by a router bug tickled by deaggregation[32] and several other cases cited there.

[26] http://tools.ietf.org/html/draft-grow-simple-leak-attack-bgpsec-no-help-00
[27] A. Pilosov and T. Kapela, "Stealing the Internet," DEFCON 16 August 10, 2008.
[28] C. Hepner and E. Zmijewski, "Defending against BGP Man-in-the-Middle attacks," Black Hat DC, February 2009.
[29] http://www.renesys.com/blog/2010/08/house-of-cards.shtml
[30] http://www.renesys.com/blog/2009/05/byte-me.shtml
[31] http://www.renesys.com/blog/2009/02/longer-is-not-better.shtml
[32] J. Cowie, The Curious Incident of 7 November 2011, NANOG 54, February 7, 2012

There are also BGP vulnerabilities that may have not been exploited in the wild so far, but that theoretically could do a lot of damage. The BGP protocol may exhibit various bizarre behaviors – such as persistent route oscillations[33] -- with certain routing policy configurations. The range of such unintended consequences is unknown with routing policies exceeding thousands of rules per router, and policy conflicts between different AS operators.

There have been several RFCs and papers addressing BGP vulnerabilities in the context of protocol standard specification and threat modeling, see the following Request For Comments (RFCs):

- RFC 4272 "BGP Security Vulnerabilities Analysis," S. Murphy, Jan 2006.
- RFC 4593 "Generic Threats to Routing Protocols", A. Barbir, S. Murphy, and Y. Yang, Oct 2006.
- Internet draft draft-grow-simple-leak-attack-bgpsec-no-help-00 "Route Leaks & MITM Attacks Against BGPSEC", D. McPherson, and S. Amante, Dec 2012.
- Internet draft draft-ietf-sidr-bgpsec-threats-01 "Threat Model for BGP Path Security", S. Kent and A. Chi, Feb 2012.
- K. Butler, T. Farley, P. McDaniel, and J. Rexford, "A survey of BGP security issues and solutions," in *Proceedings of the IEEE*, January 2010.

High-Level Requirements on Security Solutions

Any security enhancements to the inter-domain BGP routing system must not degrade any of the fundamentally desirable properties of routing that are already under stress: scalability, availability, support for multi-homing, and support for inter-domain traffic engineering.[34]

[33] http://tools.ietf.org/html/rfc3345
[34] RFC 6227 "Design Goals for Scalable Internet Routing", May 2011; RFC 4984 "IAB Workshop on Routing & Addressing", September 2007.

Appendix 2: Glossary

This appendix defines the acronyms used within this report.

Acronym	Expansion	Definition
ARIN	American Registry for Internet Numbers	One of 5 Regional Internet Registries (RIRs, see below). ARIN is the RIR for Canada, many Caribbean and North Atlantic islands, and the United States
APNIC	APNIC	Another of the five RIRs, covers the Asia Pacific region
AS	Autonomous System	Collection of connected networks represented by a set of IP prefixes subject to single, clearly defined routing policy
ASN	AS Number	The Autonomous System Number of an AS
AS_PATH	Autonomous System path	A sequence of ASNs between source and destination routers that form a directed route for packets to travel
BCP	Best Current Practices	Currently recommended operational wisdom, sometimes the subject of RFCs
BGP	Border Gateway Protocol	Interdomain routing protocol in operational use
BGPsec	BGP Security	A security extension to the Border Gateway Protocol
CA	Certificate Authority	Entity within a PKI that can issue and revoke certificates
DNSSEC	Domain Name Service Security Extensions	A suite of Internet Engineering Task Force (IETF) specifications for securing information provided by the Domain Name System (DNS)
INR	Internet Number Resources	IPv4 and IPv6 addresses and AS numbers.
IP	Internet Protocol	Network layer protocol operative on the Internet
IANA	Internet Addressing and Numbering Authority	A department of ICANN (see below) responsible for coordinating some of the key elements that keep the Internet running smoothly Specifically, IANA allocates and maintains unique codes and numbering systems that are used in the technical standards (“protocols”) that drive the Internet.

ICANN	Internet Corporation for Assigned Names and Numbers	A nonprofit private organization responsible for the coordination of the global Internet's systems of unique identifiers and, in particular, ensuring its stable and secure operation.
IRR	Internet Route Registry	Union of world-wide routing policy databases that use the Routing Policy Specification Language (RPSL)
LIR	Local Internet Registry	An organization that has been allocated a block of IP addresses by a regional Internet registry (RIR), and that assigns most parts of this block to its own customers. Most LIRs are Internet service providers, enterprises, or academic institutions. Membership in an RIR is required to become an LIR.
NIR	National Internet Registry	An organization under the umbrella of a Regional Internet Registry with the task of coordinating IP address allocations and other Internet resource management functions at a national level.
ROA	Route Origination Authorization	Authorizes a given AS to originate routes to a given set of prefixes
RIR	Regional Internet Registry	Organization that manages the allocation and registration of Internet number resources within a particular region of the world
RPKI	Resource Public Key Infrastructure	PKI defined by the IETF SIDR working group that contains certificates and objects which attest to rights of use for Internet Number Resources
RIPE NCC	Réseaux IP Européen Network Coordination Centre	RIR for Europe, Russia, the Middle East, and Central Asia
RPSL	Routing Policy Specification Language	Language commonly used by ISPs to describe their routing policies. The routing policies are stored at various “whois” databases including RIPE, RADB and APNIC. ISPs can use automated tools to generate router configuration files that match their policies.