SEPTEMBER 2012

WORKING GROUP 6
SECURE BGP DEPLOYMENT

Report

Table of Contents

1 Results in Brief

1.1 Mission Statement

The Border Gateway Protocol (BGP) controls inter-domain routing on the Internet. BGP relies on 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

This interim report is structured as an incremental continuation of the March 2012 Report on Secure BGP Deployment[1] by our Working Group 6. It is assumed that readers of this report are familiar with the deeper background information on current BGP vulnerabilities and security incidents contained in our earlier report, and with our initial recommendations. We intend to combine all interim reports in the final report to be released in March 2013.

In this contribution we focus on evaluation of risks associated with deployment and use of hierarchical resource certification system for Internet Number Resources[2] and their bindings. In order to establish a reference plane for such an evaluation we begin with currently used BGP security practices.

We direct attention to risks inherent in continuation of the Best Current Practices (BCP), if improvements such as those proposed in the Resource Public Key Infrastructure (RPKI) were not implemented. BCPs may appear robust enough, but are exclusively based on unverified and insecure databases of Internet Number Resources, and on informal personal networking among Internet engineers worldwide. They are no longer acceptable for today's business and social needs, with router misconfigurations and hijack related unreachability events lasting from tens of minutes to many hours

1 http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC-III-WG6-Final-Report.pdf

2 In this report, Number Resources refer to IP addresses and Autonomous System Numbers (ASNs).

in known incidents. BCPs will likely not be able to scale going forward with the growth of the Internet.

Given that the RPKI has not been deployed yet in any significant way, and that there has not been enough time to experiment and learn from early implementations, we look at it as follows:

We provide a high level analysis of risks of the RPKI that are potentially enabled within the standards as described in recently published RFCs (RFC 6480-6488, 6492, 6493), in particular the risks associated with abusive or punitive certificate manipulation enabled by RPKI's hierarchical structure.

The group came up with the following recommendations. They are described in greater depth, and with corresponding background information, in the main sections of this document.

1. **Desirability of Number Resource Certification**: All techniques for improving Internet routing security rely on authoritative, accurate, verifiable and tamper-proof information about which Autonomous Systems are authorized to announce and/or transit routes to which IP Address Blocks (i.e., IP Prefixes). Looking forward, the RPKI offers a more rigorous system for number resource certification. Network operators should experiment with certifying their number resources (i.e., allocated and assigned IP address blocks and AS numbers) as soon as production RPKI services are available in their region. This will allow operators to gain practical experience with the RPKI, identify any shortcomings of the current RPKI design, and aid in the safe, incremental deployment of BGP security solutions that rely on RPKI data.

2. **Need for Tools to Preview Impact of Submissions to RPKI Databases**: Network operators may make configuration mistakes in entering data into the RPKI. To reduce the likelihood of inadvertent errors, resource holders should have tools that compare new information against existing BGP routing tables before publishing the new data (e.g., the origin AS associated with an IP prefix) in the RPKI. This would allow network operators understand the impact of adding new information about their number resources ahead of commitment.

3. **Cautious, Staged Deployment Strategies:** ISPs should follow a cautiously phased deployment strategy for using RPKI-certified routing data. Different network operators may reasonably retrieve the RPKI data in different ways, and use the data in their routing policies in a customized fashion. The manner in which RPKI data is used by network operators is expected to evolve over time. We note that a phased deployment not only allows operators to gain experience with the RPKI, but also substantially reduces the risk that naive or intentional mistakes in the RPKI would cause reachability problems.

4. **Need for Transparency of RPKI Processes**: Transparency of administrative processes is the best preventive measure against erroneous or abusive actions and enables remediation. By design RPKI certificates need not disclose the organizational or personal identity of IP address and ASN holders. We recommend transparency in the processes, policies, and information bases that would allow anyone to correlate changes in the RPKI with changes in actual address allocations in a human understandable form. We recommend that additional informational records such as identification of human contact for certificate authorities be always employed (e.g., in routing registries) to assist in the operational management of RPKI data.

5. **Need for Monitoring Changes in the RPKI Data**: Confidence in the accuracy and stability of the RPKI system is a prerequisite for coupling RPKI data with real-time operation of the

routing system. An important part of building this confidence will come from monitoring the changes to the RPKI as more AS operators participate and number resources change hands. Throughout the deployment process, the RPKI data should be logged and audited to track deployment progress, detect configuration mistakes, and flag suspicious manipulations of RPKI data elements.

6. **Further Investigation of Risks Associated with the RPKI**: The design and deployment of the RPKI is still at a relatively early stage. As the community gains experience with publishing the RPKI data, and using the data to detect or avoid invalid routes, we recommend continued study of deployment scenarios and their associated risks. These studies should consider not only technical risks, but also policy issues such as manipulations of the RPKI data by powerful parties. These studies can aid in designing techniques for mitigating risks, including protocol improvements, enhanced operational practices, and new monitoring and auditing infrastructure.

While unanimity in recommendations was an objective from the outset, given the complexity of the subject each of the views expressed herein is not necessarily shared by all WG 6 members.

A glossary of key terminology is included at the end of the document.

2 Introduction

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:

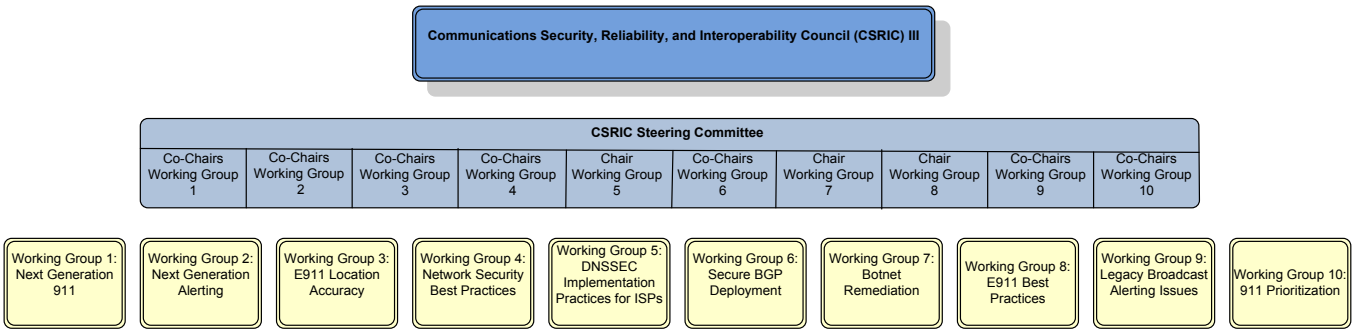

2.2 Working Group 6 Team Members

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
Mary Retka	Century Link
Isil Sebuktekin	Applied Communication Sciences
Ted Seely	Sprint
Greg Sharp	Internet Identity
Tony Tauber	Comcast
David Ward	Cisco
William Wells	TeleCommunication Systems

Table - List of Working Group Members

3 Objective, Scope, and Methodology

The security and robustness of the Internet routing system depends on availability of accurate information about the authorized holders of IP number resources. This information provides the basis upon which declarations of routing policies and constraints can be based. Examples include declarations that a given ASN is authorized to originate or transit a route to a given IP Address block.

In this section we discuss both existing and emerging information systems used to address these goals. We discuss the state of existing systems, i.e., RIR number resource "whois" databases, multiple Internet Routing Registries (IRRs), their use in current BCPs, and ways of improving their quality and trustworthiness. We discuss the emerging RPKI as way to certify number resources and authorize route origination. We identify further issues to study in the development of the RPKI system and describe an incremental approach to its deployment, evaluation and use.

These issues inform our recommendations for improving Internet routing security while minimizing risks in the operational use of existing and proposed supporting information systems.

3.1 The Need for Accurate Records about Number Resource Holders

In the Border Gateway Protocol (BGP), each Autonomous System (AS) border router receives routing UPDATE messages from its neighbors announcing reachability for IP address blocks (expressed as address prefixes) the neighbor manages or learns from other remote neighbor Autonomous Systems. Each BGP UPDATE message advertising a route contains an address prefix (e.g., 192.10.2.0/14) describing the block of routed IP addresses and various route attributes, most important of which is an AS_PATH attribute that lists the sequence of ASNs that have forwarded the UPDATE. These UPDATE messages are used by each receiving BGP router to select a "best path" to each prefix on the basis of its local routing policy. The BGP route selected as the best path is then forwarded to other neighbor ASes (as permitted by local policy) with the local ASN prepended to the AS_PATH attribute, and so on.

Erroneous BGP routing information originated by one border router can easily spread, hop by hop, to the rest of the Internet. To prevent the spread of invalid routes, border routers can be configured to perform route validation if the "ground truth" is known. The first target of validation is "origin validation" to detect when the origin AS (the very first ASN in the AS_PATH) in the received BGP UPDATE is not authorized to announce routes to the IP prefix. The results of validating individual routes can be used in various ways, including triggering alerts to operators, affecting the relative preference of routes in best path selection algorithm, or rejecting an UPDATE.

While by no means a complete solution for BGP security, origin validation is an important first step toward a more secure Internet routing system. Like all proposed BGP security solutions, origin validation relies on having accurate and timely information about number resources, such as the allocation or assignment of IP address blocks, the ASN(s) authorized to announce these address blocks, and the names of the organizations to which these IP address blocks and ASNs have been allocated or assigned.

In addition to origin validation, BGP routers may perform other more traditional forms of "filtering" of UPDATEs received from peers. In this subsection, we briefly review Best Current Practices (BCPs) for BGP filtering and the associated information systems used to support such filtering. Then we discuss the security benefits of the proposed Resource Public Key Infrastructure (RPKI) over current practices. We end by briefly discussing the interdependence of different sources of information about addressing and routing.

3.1.1 Best Current Practices for BGP Security and their Shortcomings

In today's Internet, interdomain routing security relies on AS operators configuring their BGP routers to filter invalid routes (e.g., routes with an invalid 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. Enforcement of the validity of routing information that originated several AS hops away, especially on links to AS neighbors who are not customers, is much more difficult and not typically implemented.

An ISP typically has different filtering policies for sessions with its customers (where a customer pays the provider for transit to and from the global Internet) than for other neighbors, such as peer ASes. For customer sessions, an ISP can prevent propagation of unauthorized routes by "whitelisting" allowed routes. This is done by maintaining a 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 Regional Internet Registry (RIR). These databases can be accessed either through the text-based, lightweight "whois" protocol or by a web interface to the same information.

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, to enable filtering on the AS-PATH information as well. For example, the customer's ASN should be the first hop in the path, followed by a known downstream customer. Some ISPs may employ only AS-PATH filters and not prefix filters. This practice is considered inadequate, since some common routing errors can allow for re-origination of often large numbers of illegitimate BGP updates which have the expected AS-PATH properties but are not authorized.

The success rate of these filtering techniques depends on the accuracy, truthfulness, completeness, and timeliness of the RIR IP allocation and assignment record keeping as well as on diligence of ISPs in constructing current and complete route filters from that information base. 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.

The filter contents 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, with varying degree of mirroring among themselves. Currently the IRRs are the only out-of-band network service matching prefixes to origin ASNs. The IRRs may be used as a way to check the list of prefixes a customer is asking the ISP to accept in their BGP peering. Alternately, the customer may indicate the proper set of objects to query in the IRR for the purpose of automating filter generation. It is typical in this scenario to indicate how to pull the relevant set of objects from the IRR. Examples include:

i. Origin AS (i.e., query the IRR repository database for prefixes with a given value in the "origin" field)
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 prefixes 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 such queries at regular intervals (e.g., some number of hours or days apart). The results of each run would compare the resulting list with the last returned list and would then update route filters on routers if necessary. Since this method has an automatic reliance on published data, the integrity of this data is important. The IRRs could use RPSL (Route Policy Specification Language) authentication and authorization model (see RFC2725). The RPSL authorization generally follows the address-assignment scheme but is rarely implemented. Most IRRs support one or more types of transactional authentication which can give a hint about who added a given record.

The remediation of BGP routing security failures, such as stopping of route hijacking (malicious or accidental – both inflict unreachability on the victim), or suppression of AS_PATH poisoning, 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 IRC, 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 does 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 “ground truth” about authorized holders of number resources and in particular who is authorized to announce which routes (which is nontrivial to determine in some cases even with perfect records due to origin AS changes for traffic engineering, or DDoS mitigation). A partial list of such shortcomings would include:

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

Moreover, route filtering is performed only on customer sessions, with little or no filtering on peering sessions and other non-customer sessions, which enables rapid global propagation of bogus routes. And, simply not every network operator in the world diligently follows the BCPs due to a lack of effort or other factors.

3.1.2 Resource Public Key Infrastructure (RPKI)

The Resource Public Key Infrastructure (RPKI) provides a way for the holders of Internet number

resources (that is, IP addresses and AS numbers) to formally demonstrate their ownership rights, 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 mapping directly 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 their certificates if they wish to do so. Certificates may also expire after a period of time if they are not explicitly renewed. The primary purpose of 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. In this, the RPKI is significantly different than typical PKI systems that are used to authenticate identity.

A ROA authorizes an AS to originate BGP UPDATES for one or more address blocks. Thus any BGP UPDATE 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 "origin validation" will result 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; or,
- ***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 as valid/invalid is missing from the system.

How an AS or BGP router reacts to specific origin validation results is always a matter of local routing policy. For invalid and/or unknown routes example local policies might include:

- generate an alarm to network operations staff;
- assign a lower preference to the route in best path selection; or,
- reject the route, such that it is not considered for best path.

These example policies demonstrate a range of responses from strictly advisory, to directly affecting the reachability of an address block along a specific path. 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.

Like the currently operating IRRs, the RPKI relies on AS participation to populate the repository databases. Yet, the RPKI has several advantages over IRRs that can help make the data more accurate and complete. First, today's IRRs do not have a global name space or 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 an authorized party 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. In contrast, the RPKI is not hampered by these legacy issues, and has a hierarchical structure where all certificate authorities follow the standards. 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 before the expiry period elapses when number resources are re-assigned or removed from allocation.

Based on the specifications developed in the IETF Secure Inter-Domain Routing (SIDR) working group, all five RIRs (AFRINIC, APNIC, ARIN, LACNIC and RIPE) have undertaken significant RPKI development and deployment activities and most are offering production RPKI services to their members. In North America, the American Registry for Internet Numbers (ARIN) started a pilot RPKI service in 2009, and plans to start offering an operational hosted service in the Fall of 2012, with a delegated service to follow.

The RPKI is intended as a certification of resources and the authorization of route origination. It is not intended to replace the whois databases or IRRs in their entirety. In particular, the identifying information (abuse contacts, etc) and information needed for conduct of business (organization addresses, etc), will continue to be needed and is not supplanted by the RPKI.

3.1.3 Interdependent Record Keeping Systems

Securing the global routing system using data from a resource certification system introduces at least one new large piece of global infrastructure to the operational routing system's dependency list. Integration of this new infrastructure to the global routing system will have new failure modes.

The dependencies that the current resource certification system (RPKI) has are:

1. **The existing hierarchical system of business relationships for allocating IP addresses and AS numbers.** Currently, ICANN/IANA allocates large blocks of globally unique IP addresses to five Regional Internet Registries (RIRs), which are located throughout 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. The hierarchy for delegating IP address blocks was created with two very simple goals: (i) preserve the aggregation of IP prefixes to make the routing system scale and (ii) provide an ability to create or 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. Effectively, 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.

2. **The distributed RPKI repository systems.** The RPKI is a specialized 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 does not want to operate the requisite systems themselves.

3. **The distributed system for transforming the distributed RPKI repository data into routing policy**. This is comprised of processes or protocols for fetching the RPKI data (Certificates, ROAs, and CRLs) from the distributed repositories. 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, subsequently transformed into data-structures 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 techniques for acquiring the RPKI data and using the information in routing policies.

In the RPKI system the network operators remain responsible for setting their own routing policies, just as they are today. In particular, what action a BGP speaking router should take upon receiving a route that is invalid according to the RPKI, remains the prerogative of the router operator. Network operators may choose to apply the RPKI data in their network policies in whole or in part, and there is no a requirement to use the system at all. It is expected that stages of RPKI adoption will follow the establishment of trust and confidence in the operation of the global RPKI system. Incentives to deploy and maintain the RPKI data should be used to encourage early adopters and serve as a catalyst to establishing the necessary trust in the system.

3.2 Risks Inherent in Using the RPKI

Widespread participation in the RPKI could make the Internet routing system more secure by providing timely and 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, the RPKI system's constituent parts could become attractive targets for parties that want to enforce policies, lawful orders, or inflict abuse. This could compromise the openness, robustness, and reliability of Internet communication.

3.2.1 Misconfigurations

Like any system, the RPKI may be subject to misconfigurations and mistakes when data is input. The authorization model of RPKI limits parties from accidentally entering information about IP prefixes held by others, but other errors (such as failing to issue a new certificate before an existing one expires) can certainly happen. These mistakes can impact prefix reachability in the network, depending on how RPKI information is consumed.

These mistakes are also well known from other technologies, such as DNSSEC. Certain of these failure modes have been recognized already and procedures were created to obviate them. 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.

3.2.2 Punitive Revocation

In any jurisdiction in the world a law-enforcement organization or other governmental agency or a non-governmental pressure group may desire to limit reachability for a particular resource on the Internet. One way of effecting these changes may be executed through changes to the DNS (i.e., removing the mapping of domain name to IP address) but a more drastic method is to block the reachability of the IP resource. This goal can be accomplished today, and has been, through compelling a responsible ISP by a lawful order or otherwise to stop routing IP packets to the targeted network address.

Because RIRs are subject to the lawful orders as well as extralegal pressures in the jurisdictions in which they operate, the local government may pressure or force an RIR, NIR, or LIR within its jurisdiction to compromise the reachability of IP resources which may be hosting objectionable content, notwithstanding the registry's contractual obligations with ICANN and its constituent members (in the common situation when a registry's members are located in jurisdictions other than the registry itself, the credibility of contractual obligations may be questionable). The reachability of IP resources in the routing system could also be compromised as a result of business disputes, or by the actions of abusive or compromised RIRs, NIRs, LIRs, or ISPs.

Finally, as transfers and leases of number resources 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 subtenant 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 number 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. To ISPs frustrated with the ineffectiveness of current means for reclaiming address space from former customers or customers in violation of their agreement, this capability of the RPKI may appear attractive.

It is interesting to note that if the legal or extralegal pressure could force revocation of a certificate, there's equal opportunity for the legal system or another pressure group to force re-issuance of a certificate. Thus legal actions can just as easily be to the benefit of ISPs and their customers as to their harm.

The working group has not been able to arrive at a consensus on the extent to which the RPKI increases the risk of actions that compromise prefix reachability. This topic is still the subject of debate in the broader Internet technical community as well.

3.2.2.1 Compromising prefix reachability today

We put this discussion in context by first discussing the extent to which lawful orders or punitive actions can compromise prefix reachability in the current interdomain routing system without the RPKI.

Today, if a third party wishes to block the reachability of a target IP prefix, it can pressure or force the ISPs that originate or provide transit for that prefix to stop doing so, rendering the target unreachable. The target, however, may still be able to find an alternate provider that is willing to route the prefix. The third party may also try to influence an ISP under its jurisdiction to filter the route or discard packets destined to that prefix, though the effect of this latter approach would be limited to transit customers of that ISP.

Alternatively, the third party could attempt to block the reachability more indirectly by requiring that the target prefix is removed from the IRRs, so that the target prefix is filtered by its providers, and rendered unreachable. However, compromising the reachability of an IP prefix in this manner is less effective because:

a. not all ISPs use IRRs to construct prefix filters; and,
b. IRR data is known to be occasionally stale or inaccurate.

ISPs that use IRR data to construct prefix whitelists typically make use of many different IRRs. The effectiveness of removing an entry from one IRR may not be deterministic. Effectiveness at any one ISP would depend on the intersection with the IRRs used by the ISP, the order in which multiple IRRs are searched, etc.

3.2.2.2 Revocation in the RPKI

The Certificate Authorities in the RPKI system have the ability to unilaterally overwrite or revoke their certificates issued to the CAs lower in the hierarchy. This triggers the invalidation of all resource certificates and/or ROAs issued by those CAs.

For this reason, there is a concern that a punitive revocation of a target prefix in the RPKI by various authorities or extralegal pressure groups could be used to block the target's reachability, and thus disconnect somebody from the Internet. The target may not be informed who and why caused their disconnection. However, the RPKI information is globally published so the target and the world know the RPKI action that was the cause. In short, law enforcement officers or other parties' requests to ISPs to block prefixes can be secret, but the RPKI has to be public to be effective.

3.2.2.3 Impact of revocation in RPKI on prefix reachability

It's important to note the revoking a network IP address certificate or ROA in the RPKI does not necessarily compromise reachability for this address. The severity of the effects of a punitive revocation depends on how ISPs throughout the Internet *use* the RPKI data in their routing decisions. An ISP may simply use the RPKI information as an offline source for reference about prefix to ASN mappings.

If an ISP uses the information about the mapping of IP address blocks to ASNs from the RPKI system in routers to affect BGP route selection, it is possible that the route to the IP address block will be not accepted or lowered in preference. The individual operators have the ability to decide the outcome of the route acceptance process based on policy that they define on each BGP speaking router under their control.

Using this capability, if an ISP merely assigns a lower preference to invalid routes the target prefix may still receive some or all traffic destined to the prefixes. A stricter policy of filtering invalid routes would cause the target of a punitive revocation to become unreachable.

It should be noted, however, that using a the less strict policy of depreferencing invalid routes (rather than dropping them) lowers the security level of the routing system; in particular, depreferencing invalid routes (instead of dropping them) means the routing system remains vulnerable to sub-prefix hijacks.

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 punitive revocations on prefix reachability is difficult to quantify, and warrants further study.

3.2.2.4 Comparison between the status quo and revocation in the RPKI

We compare the current practice to situations where a punitive revocation in the RPKI does result in limited reachability for the target prefix.

Suppose that, instead of requiring the target's ISPs to cease routing the target prefix, a third party instead demands a certificate revocation in the RPKI that can result in blocking reachability for the target prefix. In the former case, the target prefix may have multiple ISPs, requiring the third party to interface with multiple points of contact. Meanwhile, by electing to compromise reachability through the RPKI, the third party need only interface with a single point of contact (the delegator of the target address space), which could make the task simpler.

Next, consider revocation in the RPKI versus the deletion of an IRR route-object which provides mapping of a prefix to an origin ASN. In the case of the IRR, deletion requires ensuring that the record does not appear in any of the (possible multiple) IRR instances that an ISP uses to build its whitelists (Note: different ISPs may use different IRRs to construct their whitelists; in some cases, an ISP may use a single IRR to build its whitelist, while in other cases, multiple IRRs will be used for prefix whitelist construction). Since each IRR is run by a different entity that may be in a different jurisdiction, removing information from the RPKI could be easier, since it always involves only the single entity that issued the certificate for target prefix.

Finally, the RPKI changes the balance of power in the case of certain business disputes. Currently, an ISP who delegates address blocks to its customers has little effective means to reclaim address blocks from former customers or customers who have broken their contractual agreement. The ISP may need to resort to litigation to reclaim their legitimate resources. However, with the introduction of the RPKI, the delegator of the address space (ISP A) obtains unilateral power to easily and quickly revoke the space it delegated (to ISP B). No such power is available to ISP B: the recipient of the resource cannot do anything similar to counter the delegator's move; their only recourse may be a legal action.

3.2.2.5 Punitive revocation and collateral damage

A punitive revocation in the RPKI could cause "collateral damage" to reachability for other, non-objectionable, addresses that are adjacent to or covered by the prefixed targeted by the revocation. For example, if a small AS with a small /24 address block was targeted, but a third party required a RIR or NIR above it in the hierarchy to revoke the resource certificate for the covering /16 or larger prefix delegated to the small AS's provider, thousands or millions of other end users could be inadvertently affected. It is unclear today whether or not the possibility of collateral damage could prevent any parties from engaging in punitive revocations; some malicious parties may not be deterred by the prospect of collateral damage, while others may refuse to undertake a revocation that creates significant collateral damage.

However, in this working group we have found that it may be possible to avoid such serious collateral damage to reachability of other IP address during a revocation in the RPKI. While this is not the intended use of the RPKI, the RPKI specification in the RFCs makes it possible for anyone in the allocation hierarchy to revoke any IP prefix below it, and issue new certificates for more specifics regardless of the number of levels of the hierarchy there are between the revoker and the end-user of the prefix. With proper care, this revocation could be done with any desired granularity (in principle,

down to a single IP address), leaving all other IP prefixes and their corresponding originating ASNs valid within the RPKI. There is still potential for a less significant collateral damage, in that intermediate ISPs would be forced to announce many smaller prefixes, with consequent operational pain and routing table size increase.

It is unclear as of now whether a party's ability the discretely revoke a target prefix in the RPKI increases or decreases the risk of punitive revocations. It has also been noted that the same techniques could also be employed by points higher in the delegation hierarchy to remediate a punitive revocation.

It should be noted that today, in the absence of the RPKI, a third party can similarly order an ISP to stop routing a (possibly large) target prefix like a /16; in this case, the collateral damage would involve only the addresses covered by the target prefix that are originated by that ISP and its downstream customers. Limiting collateral damage through more fine-grained filtering, so that the ISP only stops routing a small number of addresses (e.g. a /24) is also possible. Note, however, as discussed in Section 3.2.2.4, the third party would need to order **all** of the target prefix's ISPs to stop routing the prefix to ensure the prefix is offline. With the RPKI, revocation involves only a single party -- the delegator of the address space -- which could make revocation simpler.

3.2.2.6 Mitigation through monitoring

The risk of punitive revocation may be mitigated by monitoring and auditing modifications of the RPKI databases. Whether a revoked certificate will stand out or not depends heavily on whether such modifications take place under normal conditions as well. The feasibility of monitoring-based mitigation warrants further study. The efficacy of such mitigations is also worth studying, if the number resources are no longer valid according to the RPKI, all participating parties will stop forwarding toward the affected resources, so monitoring merely informs that an action may have been taken, it cannot reverse that action. Nevertheless, it can enable remediation of a wrongful revocation.

3.2.3 Transparency of RPKI Process and Data

The RPKI purposefully did not standardize inclusion of data about the identity of resource holders in the issued resource certificates. In fact, certificate issuers may assign completely arbitrary and variable names to certificate subjects. 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.

Having said that, the RPKI is meant to be a resource certification infrastructure that accurately reflects the state of the underlying number resource allocation systems. In order to build confidence in, debug and maintain this correspondence, it is important that the process, policies, procedures and data sets associated with current resource allocation systems (e.g., RIR addressing registries and information systems) permit for easy, publicly accessible ways of monitoring, diagnosing and managing the state of RPKI system.

Another component of this issue is the ability to identify human points of contact (POC) responsible for the operation of certificate authorities (CAs) in the distributed RPKI system. For this reason, we

recommend that all CAs include RFC6493 records that identify a POC for each CA.

RIRs currently deploying production RPKI services are developing policies (i.e., certificate practice statements) and processes to insure that the RPKI provides an accurate reflection of current resource assignments in their regions. As ISPs and other entities in the address allocation hierarchy deploy RPKI systems, they will have to take similar steps to ensure that their RPKI data accurately reflects their current allocation state.

4 Recommendations

While current best common practices for protecting the global BGP routing system have "gotten us this far", it seems unlikely that the existing techniques can adequately scale to meet the challenges of the future Internet. Pressures on the addressing and routing system from IPv4 address exhaustion and other forces will only increase the complexity of resource certification and route authorization going forward. Long term improvements to BGP security requires the transition to widespread use of RPKI repositories of certified address ownership and routing data for route filtering. In order to achieve this goal RPKI deployment and adoption plans must minimize the potential of introducing new risks and failure modes associated with actual implementation and control of this new critical infrastructure. A staged deployment strategy for the RPKI system together with improvements to the publicly available RIR databases is necessary to build trust and confidence in RPKI services. Independent monitoring and analysis of RPKI services during this staged deployment will allow for early detection of bugs, adequate learning curve for the AS operators and development of new Best Practices. In this section we summarize our recommendations for incremental deployment of the RPKI in light of these issues.

1. Desirability of Number Resource Certification: All techniques for improving Internet routing security rely on authoritative, accurate and verifiable information about which Autonomous Systems are authorized to announce routes to which IP Address Blocks. To improve upon today's best common practices, network operators should begin by ensuring their Internet Routing Registry (IRR) records are public, complete, and up-to-date. ISPs that delegate portions of their address space to customers should register these allocations and assignments in RIR and IRR databases and maintain their accuracy. Looking forward, the RPKI offers a more rigorous and verifiable system for number resource certification. AS operators should begin certifying their numbered resources as soon as production RIR RPKI services are available in their region. In the case of North America, ARIN has announced that their RPKI pilot will transition to a production hosted-model system in Fall of 2012. A full delegated RPKI ARIN service should follow later. This will allow operators to gain practical experience with RPKI, identify benefits and shortcomings of the current RPKI design, and aid in the incremental deployment of BGP security solutions that rely on RPKI data.

2. Need for Tools to Preview Impact of Submissions to RPKI Databases: AS operators may make configuration mistakes in entering data into the RPKI. To reduce the likelihood of inadvertent errors, certificate authorities should have tools that compare new information against existing BGP routing tables (e.g., the origin ASes associated with each IP prefix) before publishing the new data in the RPKI. This would allow AS operators to understand the impact of adding new information about their number resources ahead of commit time. RIPE already has such a tool as part of its hosted service, and other parties involved in providing RPKI services should offer similar tools to reduce the likelihood of errors.

3. Cautious, Staged Deployment Strategies: ISPs should follow a phased deployment strategy for using certified routing data from RPKI. Different network operators may reasonably retrieve the RPKI data in different ways, and use the data in their routing policies in a customized fashion. Initially, the RPKI data could be helpful as part of a manual or automated vetting process at the time a customer BGP session is configured or reconfigured at customer request. Also, the RPKI data could be used for analysis of BGP-learned routes and detection of suspicious route announcements. After some period used to get acquainted with the RPKI systems and data, operators may consider having their routers use the RPKI data automatically in making routing decisions. For example, ISPs might use the RPKI data as one of several inputs in generating routing policies that de-preference or filter "known invalid" routes, or de-preference routes classified as "unknown." We note that a phased deployment not only allows operators to gain experience with the RPKI, but also substantially reduces the risk that naive or intentional mistakes in the RPKI would cause reachability problems.

4. Need for Transparency of RPKI Processes: Transparency of administrative processes is the best preventive measure against erroneous or abusive actions and enables remediation. However, the design of RPKI is such that certificates need not disclose the organizational or personal identity of IP address and ASN holders. We recommend that address allocation and registration policies be coordinated with RPKI certificate practice statements (CPSs) such that monitoring and human analysis of the state and any changes to the RPKI is easily achievable. We recommend that additional informational records (such as Ghostbusters Records in RFC6493 identifying a human contact for certificate maintenance) be readily available to assist in the operational management of the information systems that create (e.g, CAs) and store (e.g., publication points) RPKI data.

5. Need for Monitoring Changes in the RPKI data: Confidence in the accuracy and stability of the RPKI system is a prerequisite for coupling RPKI data with real-time operation of the routing system. An important part of building this confidence comes from monitoring the changes to the RPKI as more ASes participate and number resources change hands. Throughout the deployment process, the RPKI data should be logged and audited to track deployment progress, detect configuration mistakes, flag suspicious manipulations of ROAs, and so on. 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 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 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 or other registries, and report the results.

6. Further investigation of risks associated with the RPKI: The design and deployment of the RPKI are still at a relatively early stage. As the network operator community gains experience with publishing the RPKI data, and using the data to detect or avoid invalid routes, we recommend continued study of deployment scenarios and their associated risks. These studies should consider not only technical risks, but also policy issues such as manipulations of the RPKI data by powerful parties. These studies can aid in designing techniques for mitigating risks, including protocol improvements, enhanced operational practices, and new monitoring and auditing infrastructure.

Together, we believe these six recommendations can lead to an Internet routing system that is both more secure and more robust.

5 Conclusions

The Internet routing system is notoriously vulnerable to accidental misconfigurations and malicious attacks. We believe that following the steps outlined in this report will lead to a more secure routing system, while respecting the local autonomy of the many networks that make up the Internet.

Appendix: Glossary

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 5 RIRs, covers the entire Asia Pacific region
AS	Autonomous System	Collection of connected networks represented by a set of network prefixes subject to single and clearly defined routing policy
ASN	AS Number	The Autonomous System number of an AS
AS-Path		An AS-PATH is a sequence of intermediate ASes 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 within a country or economic unit.
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 (using automated tools) then generate router configuration files that match their business and technical policies.