



MARCH 2012

WORKING GROUP 6
SECURE BGP DEPLOYMENT

Report

Table of Contents

1. Results in Brief
 2. Introduction
 3. Objectives, Scope, and Methodology
 4. Recommendations
 5. Conclusions
- Appendix: Background on BGP Security

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

Although the working group's mission statement addresses ISPs, all network operators participating in inter-domain routing on the Internet should be concerned about BGP security. As such, the group's recommendations apply not only to ISPs, but also to content providers, enterprise networks, and other stakeholders in the global Internet routing system. In addition, as BGP is the glue that holds the disparate parts of the Internet together, global adoption of a security solution must be weighed against the primary goal of ensuring a robust and reliable system. Since the routing system has no central authority, and the many constituent networks have different objectives and business concerns, any viable security solution must preserve the local autonomy of these networks.

During the first few months of meetings, the working group discussed and compared proposed solutions for improving BGP security, and came to the following three recommendations, discussed in greater depth in the remainder of this report:

- **Establishing ground truth through resource registration and certification:** BGP security hinges on having a stronger notion of ground truth, through registration and certification of Internet number resources and associated routing system bindings, such

as which Autonomous Systems (ASes) may originate BGP reachability information (*routes*) for particular IP address blocks (in the form of *prefixes*). To that end, we recommend that AS operators: (i) ensure their Internet routing registry (IRR) records are public, complete, and up-to-date, (ii) encourage ARIN to deliver their hosted Resource Public Key Infrastructure (RPKI) service, and (iii) encourage a single global “root of trust” for the RPKI.

- **Phased deployment of techniques that detect and prevent route hijacking:** Each AS can apply local policies for storing, disseminating, and using information about certified number resources to detect and prevent route hijacking. To this end, we recommend that AS operators (i) track the ongoing developments in the secure BGP community and (ii) consider phased deployment strategies for using certified routing data in ways that are consistent with their own internal policies. In addition, the BGP security community should investigate the new risks introduced by resource certification.
- **Metrics and measurements for evaluating security problems and solutions:** Better security metrics, and continuous monitoring to compute these metrics, can help quantify the frequency and scope of routing security incidents, and evaluate the effectiveness of proposed security improvements. We recommend that the BGP security community (i) evaluate existing BGP security metrics, and extend them where necessary, and (ii) perform continuous monitoring and analysis of BGP security incidents.

While the principal objective of the working group was development of recommendations related to “Secure BGP Deployment”, there was consensus among the working group members that Internet number resource allocation, certification, operational procedures, and an array of other externalities have considerable implications on the ability to better secure the global Internet routing system.

While unanimity in recommendations was an objective from the outset, each of the views expressed herein is not necessarily shared by all of the working group members. In addition, we note that the working group is strictly *advisory* in nature, putting forth recommendations that encourage the market-led adoption of security technologies rather than advocating any regulatory policy or inventing any new security solutions.

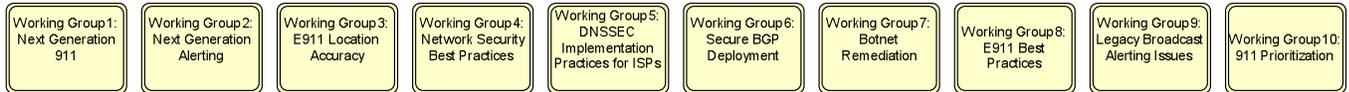
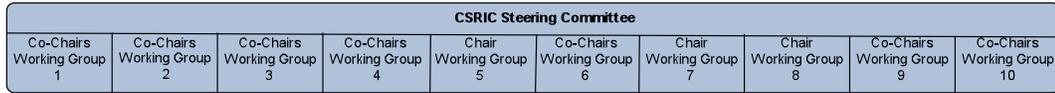
2 Introduction

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:

Communications Security Reliability and Interoperability Council (CSRIC) III



2.2 Working Group 6 Team Members

Working Group 6 consists of the members listed in the following table.

Name	Organization
Andy Ogielski, Co-Chair	Reneseys
Jennifer Rexford, Co-Chair	Princeton University
Shane Amante	Level3
Daniel Awduche	Verizon
Ron Bonica	Juniper
Jay Borkenhagen	AT&T
Martin Dolly	ATIS/AT&T
Andy Ellis	Akamai
Sharon Goldberg	Boston University
Adam Golodner	Cisco
Kyle Hambright	Las Vegas Metro Police
Lars Harvey	Internet Identity
Michael Kelsen	Time Warner Cable
Ed Kern	Cisco
Eric Lent	Comcast
Danny McPherson	Verisign
Doug Maughan	DHS S&T
Doug Montgomery	NIST
Christopher Morrow	Google
Sandra Murphy	SPARTA
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 1 - List of Working Group Members

3 Objective, Scope, and Methodology

3.1 Objective

This working group has been chartered to address improvements to the security of the Internet's inter-domain routing system, including both the IPv4 and IPv6 versions of the Internet protocols. The Internet routing system is a large, loosely interconnected global system whose architecture, operation, and topology reflect diverse business relationships among network operators and enterprises, on national and international levels. The Internet consists of tens of thousands of separately administered networks known as Autonomous Systems (ASes).

The current implementation of inter-domain routing on the Internet relies on the Border Gateway Protocol version 4 (BGPv4), which allows each AS to originate route announcements for IP address blocks (or *prefixes*) for which it provides reachability. In BGP, each AS learns routes to remote destination prefixes from neighboring ASes, and chooses to utilize, discard, and/or propagate those routes to other neighbors according to local policies. Routing information that does not accurately reflect the intent of the authorized originator or some intermediate party can easily propagate from one AS to another, leading to serious global consequences such as “traffic black-holing” (where traffic never reaches its destination) or “traffic detouring” (where traffic is re-routed through an intermediate network that may observe or analyze the traffic).

For more detail on the security vulnerabilities of BGP, see the recent surveys on the topic¹, as well as the Appendix below.

To protect their networks and their customers, many network operators follow Best Current Practices (BCPs) such as configuring their routers to perform defensive filtering of route announcements, or to limit the number of routes they accept from each BGP neighbor. However, ASes remain vulnerable to incorrect information that originates several hops away. Further improvements to BGP security rely on having a common notion of “ground truth” about which AS(es) can originate routes for each IP prefix, and ASes using this information to detect or prevent the spread of invalid routing information.

3.2 Scope and Methodology

During the first several working group meetings, the working group members compared a range of candidate security solutions across a variety of dimensions, including technical maturity, deployment 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.

This discussion sets the stage for many of the recommendations that follow in the next section to encourage the incremental deployment of solutions based on Internet number resource certification and origin authorization.

3.2.1 BEST CURRENT PRACTICES

¹ See RFC 4272 on “BGP Security Vulnerabilities Analysis,” and “A Survey of BGP security issues and solutions” from *Proceedings of the IEEE*, January 2010.

Today's Best Current Practices rely on individual ASes to configure their routers to block invalid routes, based mainly on local information. For example, an AS should filter BGP announcements from customers to explicitly permit only prefixes each customer is authorized to advertise, preventing them from effectively advertising routes for IP address blocks they do not "own," or wrongly exporting routes learned from one peer or provider to another peer or provider. This is particularly important since some router implementations, by default, announce *all* routes in the local BGP table to neighboring ASes. An AS may apply other filtering policies, such as filtering routes for private addresses (e.g., RFC 1918 private use addresses, which should not be announced in the global Internet), for very small address blocks (which could easily overwhelm the BGP routing tables), or limiting the number of prefixes accepted from each neighbor. In fact, even if more sophisticated security technologies are deployed, each AS should still perform such "defensive filtering" to protect itself (and the rest of the Internet) from accepting and potentially propagating erroneous routing information.

If every AS applied BCPs, many of the security problems with BGP today would disappear. Unfortunately, some network administrators do not apply these practices, leaving an AS vulnerable to erroneous routing information propagated by ASes multiple hops away. Thus, while BCPs can be incrementally deployed, achieving sufficient participation has proven elusive for an array of reasons, some of which we will discuss later, and some of which are the topic of CSRIC III Working Group 4, *Network Security Best Practices*.

In terms of cost, the BCPs introduce relatively little capital expense, since today's commercial routers can support the necessary filtering policies. Network operators do incur operational cost to configure BGP sessions with neighbors, although automation reduces these costs and the complexity of implementation. These operational practices are mature and well understood.

Internet Route Registries (IRRs) have been proposed as a source of information about who is authorized to announce what IP address blocks. However, given the current state of the IRRs, ASes cannot create filtering policies with sufficient "reach" and efficacy. An AS is typically not in a good situation to judge the veracity of BGP routing information sent by far-away ASes, or even the information published in IRRs. As such, AS operators are understandably conservative, if not extremely reluctant, in specifying route filters derived directly and/or solely from IRRs, to avoid unintentionally filtering valid information that may reduce Internet reliability and performance for their customers. In addition, an AS cannot easily constrain how its neighbors select and announce routes it announces.

3.2.2 ORIGIN CERTIFICATION

To protect against accidental prefix hijacking from an unauthorized origin AS, organizations participating in BGP routing could conceivably establish and maintain ground truth of all legitimate ASes that are authorized to originate each address block (*prefix*) in the global routing system. Organizations could acquire certificates that prove they have been assigned a particular Internet number resource, such as an IP address block or Autonomous System number. The use of a Resource Public Key Infrastructure (RPKI) currently has momentum in routing security circles, 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 validated by third parties. During early deployment, a system for number resource certification such as RPKI may even be used to inform and fortify other provisioning and configuration systems such as the prefix origin binding information

commonly contained in the IRRs.

Based on the work in the SIDR (Secure Inter Domain Routing) working group at the IETF (Internet Engineering Task Force), different regional Internet registries (RIRs) have started offering 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 initial operational service by the end of the first half of 2012.

In addition to certifying their resources, ASes have the option of using the RPKI data to detect, de-preference, or filter “known invalid” BGP routes received from neighbors. Network administrators can use the RPKI information when configuring their existing routers, or apply software upgrades that allow local caches of RPKI data to update individual routers automatically with the necessary data. This technique is effectuated through recently standardized router-to-RPKI protocols and populates soft-state memory in the routers to inform local BGP policy functions of authorized origin AS/prefix bindings. This information can augment existing BGP policies to protect against unauthorized ASes originating a given address prefix in the routing system. In terms of cost, many newer generations of routers have the capacity necessary to check the origin ASes of BGP routes, as no cryptographic operations are necessary. That said, some older routers do not have the capacity to apply all of these checks, or to load software that enables importing information from an RPKI cache that contains the binding information.

These new RPKI cache-to-router protocols, if used, may introduce new attack surfaces that need further analysis beyond the current provisioning and configuration systems operators utilize. In addition, legitimate changes in the origin ASes may take time to propagate through the RPKI system and be uploaded to the local routers, particularly if some ASes only download new RPKI information periodically (e.g., daily). Origin certification is amenable to incremental deployment, since the technology does not require any changes to the BGP protocol itself, and an AS can start using the RPKI data locally to influence its routing decisions, even before other ASes begin utilizing such a system. While origin certification handles the most prevalent security incidents (e.g., misconfigured origin ASes or naïve prefix-hijacking attacks), more sophisticated attacks that exploit BGP’s hop-by-hop dissemination of routes or manipulate the sequence of ASes in the BGP AS_PATH attribute are still effective, as long as the attacker ensures that the origin AS matches an authorized origin for the prefix in question.

3.2.3 PATH VALIDATION

To protect against attacks that involve the AS_PATH attribute in a BGP route announcement, the IETF SIDR working group is designing and standardizing techniques for cryptographic protections of the path information. In contrast to policy derived from origin certification, cryptographic protection of the BGP AS_PATH attribute, and subsequent path validation by BGP routers, changes the BGP protocol to have ASes cryptographically sign and verify the hops in the BGP AS_PATH attribute. After selecting a best route for a destination prefix, the AS signs the update messages sent to each neighbor; these neighbors, in turn, verify the signatures along the path before accepting the announcement. Checking these signatures prevents an attacker from artificially shortening, lengthening, poisoning, or otherwise modifying the sequence of ASes in the AS_PATH. As such, under path validation, the BGP AS_PATH attribute is an accurate reflection of the sequence of ASes that propagated the BGP route announcement.

The standardization process for path validation is not yet complete, and commercial implementations of the technology are not yet available. That said, the technology is under active design and specification at the IETF. In terms of cost, path validation will require software upgrades to the routers to perform the necessary cryptographic operations. New router hardware may be necessary as well, since the signing and validating of routes requires additional processing and memory.

With cryptographic protections in BGP AS_PATH validation, it is critical to understand that ASes that propagated the route (i.e., advertised “transit” reachability to the destination through their local AS) may very well have done so unbeknownst to or without the authorization of the resource holder, for purposes benign or malicious, even if the origin in the path is an authorized origin – this is simply an artifact of distributed database protocols and autonomous local routing operations by routing system participants. However, if this were to occur, the semantic integrity protections BGP AS_PATH validation affords would require the intermediate AS to accurately record and reflect the local AS(es) in the AS_PATH attribute. For this reason, route advertisement implications related to *policy and intent* of the resource holder or other intermediate ASes is outside of the scope of the current BGP path validation work occurring in the IETF SIDR working group. The areas of routing *policy* and *resource holder intent* are a critical area of study for routing security, but expressly outside the scope of this focused path validation effort.

4 Recommendations

Our recommendations fall in three main categories.

I. Establishing ground truth through resource registration and certification

Improving routing security depends on having more effective ways to identify invalid routing information received from neighboring ASes, in particular which AS(es) are authorized to originate an IP prefix, in a manner separate from the global routing system. Establishing ground truth through accurate and complete routing registries, and certification of number resources, are important steps for many security solutions. As such, the working group explored ways all Internet routing system operators -- such as enterprises, content providers, and ISPs -- can improve the accuracy of routing registries and participate in resource certification infrastructure.

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 invalid or suspicious routing information is a prerequisite for most BGP security solutions. Yet, maintaining accurate registries will likely become more challenging in the years ahead, as depletion of IPv4 address space leads to more “trading” of address blocks between different institutions and potentially weaker incentives for updating the registries with truthful information. We recommend that every AS operator ensure its routing registry records are public, complete, and up-to-date. ISPs that delegate portions of their address space to customers should register and maintain accuracy of these address allocations and assignments in IRRs or whois systems as appropriate. 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). We stress

that the use of routing registries provides one reliable method of validating IP address and AS number assignments outside of the routing protocol, and into the future, may be informed and contain formally verifiable information through global resource certification infrastructure (e.g., RPKI).

AS operators should encourage ARIN to deliver a hosted RPKI service: Security enhancements to BGP hinge on the availability of certified resource data that indicate which entities are allowed to announce routes for IP addresses. We recommend ARIN members encourage ARIN to deliver their hosted RPKI service during the first half of 2012, as scheduled². When the ARIN service becomes available, resource holders should begin making use of that service to certify their resources and gain operational experience with Internet number resource certification. In the meantime, AS operators can experiment with ARIN's pilot RPKI service to start gaining operational experience, ahead of the hosted service.

AS operators should encourage a single global "root of trust" for the RPKI: The working group recognizes that, during the early stages of RPKI deployment, the global routing system may have multiple trust anchors ('roots'). However, for a global resource certification to be most effective, Autonomous System operators should stand with the Internet Architecture Board (IAB)³ in encouraging the RIRs 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 will fall to AS operators, who have little to 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.

We note that meaningful improvements in BGP security will require worldwide cooperation for certifying number resources among all Regional Internet Registries and countries. We recommend that the FCC work with its international partners to encourage wider participation in registering and certifying routing resources.

II. Phased deployment of techniques that detect and prevent route hijacking

The registration and certification of number resources can, over time, enable AS operators to detect erroneous routing information and prevent it from propagating through the routing system.

Any successful BGP security solution must balance the desire for better global security against the need for each AS operator to make independent local decisions, with minimal reliance on centralized control and without introducing undue complexity, additional latency, or brittleness to the local system or global infrastructure. As such, while we recommend that AS operators follow ongoing developments in the BGP security community, we intentionally stop short of advocating specific technical choices or route-selection policies. Each network operator should understand the implications of these systems and mechanisms, and exert due consideration when

²<http://www.nanog.org/meetings/nanog54/abstracts.php?pt=MTkxNCZuYW5vZzU0&nm=nanog54>

³<http://www.ietf.org/mail-archive/web/ietf-announce/current/msg07028.html>

effectuating these mechanisms in their networks. Instead, we suggest a phased approach to storing, disseminating, and using certified routing data in operational networks.

AS operators should track the developments in the secure BGP community: Given the importance of BGP routing security, we recommend that ISPs, content providers, and anyone else configuring BGP-speaking routers to pay close attention to the ongoing developments in the global routing operations and BGP security communities. Technical experts operating BGP-speaking networks are encouraged to actively participate in the IETF (Internet Engineering Task Force) working group on SIDR (Secure Inter Domain Routing), as well as in their respective RIR and network operations group (NOG) communities.

AS operators should consider phased deployment strategies for using certified routing data in ways that are consistent with their own internal policies: As BGP security technology and supporting infrastructure matures, ASes can move toward storing, distributing, and using certified routing information in their networks. Different AS operators may reasonably retrieve the information in different ways, and apply varying mechanisms controlling whether and how they use this information to detect, de-preference, or filter known invalid routes. In the early stages, an AS may perform offline analysis of observed BGP routes against the certified origin ASes, or use the RPKI data as one of several inputs in constructing route filters. At this time, we recommend that any direct, *automatic* feedback into the routing system be undertaken with due caution, so as to maintain the stability of the current system. In the meantime, ASes could perform lab testing to explore the engineering and operational implications of possible real-time, automated use of this resource certification data in the future.

The BGP security community should investigate the new risks introduced by resource certification: Currently, a small handful of Regional and National Internet Registries throughout the world provide assurance as to the global uniqueness of the Internet numbers (i.e., IP addresses and AS numbers) they delegate to downstream organizations. The RIRs and NIRs specifically do not provide any assurance of the *global reachability* of Internet numbers in the routing system. In the future, with a resource certification system for Internet numbers in place, RIRs and NIRs would provide an even more critical function of determining whether, or not, all destinations on the Internet are not only globally unique, but also globally reachable. While helping prevent illicit route hijacking, this functionality also raises significant legal, policy, and economic questions. More specifically, RIRs and NIRs could become attractive targets to various outside entities/organizations that wish to enforce various types of policies or lawful orders that could dramatically affect the openness, robustness, and reliability of communications on the Internet. We recommend that experts in the legal and policy realm, particularly as it relates to Internet governance, study these concerns very carefully in order to evaluate not only the benefits, but also the risks associated with deployment and use of a hierarchical resource certification system for Internet numbers.

In the coming months, the working group will explore phased-deployment strategies that balance the goal of better BGP security with the needs for a resilient inter-domain routing system and the local autonomy of ASes, and analyze risks introduced by resource certification.

III. Metrics and measurements for evaluating security problems and solutions

Quantifying the security problems with BGP, and the effectiveness of proposed solutions, is

important for informing decisions about which security solutions to deploy, and timing dependencies associated with this deployment. As such, the working group discussed the need for better metrics and measurements of the routing system, and will continue working on these issues in the months ahead. At this time, we have the following initial recommendations:

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 operationally defined metrics for identifying routing security incidents, both accidental and malicious, and quantifying their scope and impact⁴. 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, and the community does not have broad consensus on which security metrics and measurement methodologies are most effective. We recommend that the community compare existing metrics and methodologies, and extend them where necessary, to better evaluate security problems and solutions.

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 datasets.

It is beyond the scope of the current report to recommend any particular set of metrics or a continuous monitoring methodology, or any specific action by AS operators, at this time. If desired, these issues will be addressed in future reports.

5 Conclusions

This report recommends a stronger notion of ground truth for which ASes can originate routes for each IP address block, and a phased deployment of techniques that detect and prevent some route hijacking attacks while preserving the local autonomy of ASes, and better security metrics and continuous measurement. In the coming months, the working group will review these initial recommendations and continue to explore more concrete recommendations that can facilitate improvements in the security of the Internet’s inter-domain routing system. The working group will also attempt to capture the subtle distinction between cryptographically verifiable AS_PATH validation functions, and the intersection of *policy* and *intent* of resource holders.

Appendix: Background on BGP Security

Salient Features of BGP Operation

⁴ For comparison, consider the annual data breach reports published by Verizon, or quarterly reports published by Messaging Anti-Abuse Working Group (MAAWG). http://www.maawg.org/email_metrics_report

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.

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 standardized, but instead, are largely dictated by AS interconnection and business relationships. 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.

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 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”. 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 incidents to features of proposed solutions that could 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.

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”.

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”), 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 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 YouTube⁵ (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”⁶, or a hijacking of a very large portion of the Internet for several hours by TTNNet in 2004⁷, or a 2006 ConEd incident⁸. Without knowing the 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).
 - For routing incidents triggered by long AS_PATHs see [House of Cards](#)⁹, [AfNOG Takes Byte Out of Internet](#)¹⁰, [Longer is Not Always Better](#)¹¹ for actual examples.
 - Route leaks - A possibility of “man in the middle” (MITM) AS_PATH attacks detouring traffic via a chosen AS was publicly demonstrated at DEFCON in 2008¹². Two other similar incidents were found in a 7-month period surrounding the DEFCON demo by mining of a BGP update repository conducted in 2009¹³ but were not confirmed as malicious. This can occur either by accident as detailed above, and is sometimes referred to as route “leaks”, 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.
 - 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

⁵ <http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml>

⁶ <http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml>

⁷ Alin C. Popescu, Brian J. Premore, and Todd Underwood, Anatomy of a Leak: AS9121. NANOG 34, May 16, 2005.

⁸ <http://www.renesys.com/blog/2006/01/coned-steals-the-net.shtml>

⁹ <http://www.renesys.com/blog/2010/08/house-of-cards.shtml>

¹⁰ <http://www.renesys.com/blog/2009/05/byte-me.shtml>

¹¹ <http://www.renesys.com/blog/2009/02/longer-is-not-better.shtml>

¹² A. Pilosov and T. Kapela, Stealing the internet, DEFCON 16 August 10, 2008

¹³ C. Hepner and E. Zmijewski, Defending against BGP Man-in-the-Middle attacks, Black Hat DC February 2009

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 deaggregation¹⁴ and several other cases cited there.

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 does not have solid mathematical foundations, and certain bizarre behaviors – such as persistent route oscillations – are quite possible.

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-foo-sidr-simple-leak-attack-bgpsec-no-help-01 “Route Leak Attacks Against BGPSEC”, D. McPherson and S. Amante, Nov 2011.
- Internet draft draft-ietf-sidr-bgpsec-threats-01 “Threat Model for BGP Path Security”, S. Kent and A. Chi, Feb 2012.

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, support for multi-homing, and support for inter-domain traffic engineering¹⁵.

From the perspective of network operators, network availability is more fundamental than security, and especially so in national emergencies and disaster recovery situations. New security solutions should not impede network availability and restoration.

¹⁴ J. Cowie, The Curious Incident of 7 November 2011, NANOG 54, February 7, 2012

¹⁵ RFC 6227 “Design Goals for Scalable Internet Routing”, May 2011; RFC 4984 “IAB Workshop on Routing & Addressing”, September 2007.