[March, 2013]

WORKING GROUP 4

Network Security Best Practices

FINAL Report – BGP Security Best Practices

Table of Contents

Table of Contents

1 Results in Brief

1.1 Charter

This Working Group was convened to examine and make recommendations to the Council regarding best practices to secure the Domain Name System (DNS) and routing system of the Internet during the period leading up to some significant deployment of protocol extensions such as the Domain Name System Security Extensions (DNSSEC), Secure BGP (Border Gateway Protocol) and the like. The focus of the group is limited to what is possible using currently available and deployed hardware and software. Development and refinement of protocol extensions for both systems is ongoing, as is the deployment of such extensions, and is the subject of other FCC working groups. The scope of Working Group 4 is to focus on currently deployed and available feature-sets and processes and not future or non-widely deployed protocol extensions.

1.2 Executive Summary

Routing is what provides reachability between the various end-systems on the Internet be they servers hosting web or email applications; home user machines; VoIP (Voice over Internet Protocol) equipment; mobile devices; connected home monitoring or entertainment systems. Across the length and breadth of the global network it is inter-domain routing that allows for a given network to learn of the destinations available in a distant network. BGP (Border Gateway Protocol) has been used for inter-domain routing for over 20 years and has proven itself a dynamic, robust, and manageable solution to meet these goals.

BGP is configured within a network and between networks to exchange information about which IP address ranges are reachable from that network. Among its many features, BGP allows for a flexible and granular expression of policy between a given network and other networks that it exchanges routes with. Implicit in this system is required trust in information learned from distant entities. That trust has been the source of problems from time to time causing reachability and stability problems. These episodes have typically been short-lived but underscored the need for expanding the use of Best Current Practices (BCPs) for improving the security of BGP and the inter-domain routing system.

These mechanisms have been described in a variety of sources and this document does not seek to re-create the work done elsewhere but to provide an overview and gloss on the vulnerabilities and methods to address each. Additionally, the applicability of these BCPs can vary somewhat given different deployment scenarios such as the scale of a network's BGP deployment and the number of inter-domain neighbors. By tailoring advice for these various scenarios, recommendations that may seem confusing or contradictory can be clarified. Further, an appendix includes a table that indexes the risks and countermeasures according to different deployment scenarios.

Issues that the working group considered included:

- Session hijacking
- Denial of service (DoS) vulnerabilities
- Source-address filtering
- BGP injection and propagation vulnerabilities

- Hacking and unauthorized access to routing infrastructure
- Attacks against administrative controls of routing identifiers

Working Group 4 recommends that the FCC encourage adoption of numerous best practices for protecting ISPs' routing infrastructures and addressing risks related to routing that are continuously faced by ISPs.. Inter-domain routing via BGP is a fundamental requirement for ISPs and their customers to connect and interoperate with the Internet. As such, it is a critical service that ISPs must ensure is resilient to operational challenges and protect from abuse by miscreants.

SPECIAL NOTE: For brevity, and to address the remit of the CSRIC committee to make recommendations for ISPs, the term ISP is used throughout the paper. However, in most instances the reference or the recommendations are applicable to any BGP service components whether implemented by an ISP or by other organizations that peer to the Internet such as business enterprises, hosting providers, and cloud providers.

2 Introduction

CSRIC was established as a federal advisory committee designed to provide recommendations to the Commission regarding Best Practices and actions the Commission may take to ensure optimal operability, security, reliability, and resiliency of communications systems, including telecommunications, media, and public safety communications systems.

Due to the large scope of the CSRIC mandate, the committee then divided into a set of Working Groups, each of which was designed to address individual issue areas. In total, some 10 different Working Groups were created, including Working Group 4 on Network Security Best Practices. This Working Group will examine and make recommendations to the Council regarding best practices to secure the Domain Name System (DNS) and routing system of the Internet during the period leading up to a anticipated widespread implementation of protocol updates such as the Domain Name System Security Extensions (DNSSEC) and Secure Border Gateway Protocol (BGPsec) extensions.

The Working Group presented its report on DNS Security issues in September 2012.

This Final Report – BGP Best Practices documents the efforts undertaken by CSRIC Working Group 4 Network Security Best Practices with respect to securing the inter-domain routing infrastructure that is within the purview of ISPs, enterprises, and other BGP operators. Issues affecting the security of management systems that provide control and designation of routing and IP-space allocation records that BGP is based on were also considered.

Routing and BGP related services are necessary and fundamental components of all ISP operations, and there are many established practices and guidelines available for operators to consult. Thus most ISPs have mature BGP/routing management and infrastructures in-place. Still, there remain many issues and exposures that introduce major risk elements to ISPs, since the system itself is largely insecure and unauthenticated, yet provides the fundamental traffic control system of the Internet. This report enumerates the issues the group identified as most critical and/or that may need more attention.

2.1 CSRIC Structure

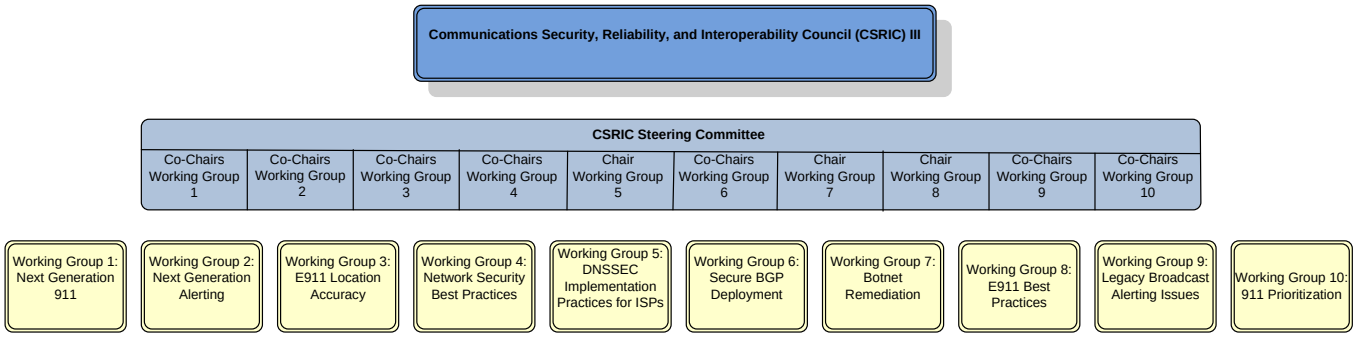

2.2 Working Group [#4] Team Members

Working Group [#4] consists of the members listed below for work on this report.

Name	Company
Rodney Joffe – Co-Chair	Neustar, Inc.
Rod Rasmussen – Co-Chair	Internet Identity
Mark Adams	ATIS (Works for Cox Communications)
Steve Bellovin	Columbia University
Donna Bethea-Murphy	Iridium
Rodney Buie	TeleCommunication Systems, Inc.
Kevin Cox	Cassidian Communications, an EADS NA Comp
John Crain	ICANN
Michael Currie	Intrado, Inc.
Dale Drew	Level 3 Communications
Chris Garner	CenturyLink
Joseph Gersch	Secure64 Software Corporation
Jose A. Gonzalez	Sprint Nextel Corporation
Kevin Graves	TeleCommunication Systems (TCS)
Tom Haynes	Verizon
Chris Joul	T-Mobile
Mazen Khaddam	Cox
Kathryn Martin	Access Partnership
Ron Mathis	Intrado, Inc.
Danny McPherson	Verisign
Doug Montgomery	NIST
Chris Oberg	ATIS (Works for Verizon Wireless)
Victor Oppleman	Packet Forensics
Elman Reyes	Internet Identity
Ron Roman	Applied Communication Sciences
Heather Schiller	Verizon
Jason Schiller	Google
Marvin Simpson	Southern Company Services, Inc.
Tony Tauber	Comcast
Paul Vixie	Internet Systems Consortium
Russ White	Verisign
Bob Wright	AT&T

Name	Company
Rodney Joffe – Co-Chair	Neustar, Inc.
Rod Rasmussen – Co-Chair	Internet Identity
Mark Adams	ATIS (Works for Cox Communications)
Steve Bellovin	Columbia University
Donna Bethea-Murphy	Iridium
Rodney Buie	TeleCommunication Systems, Inc.
Kevin Cox	Cassidian Communications, an EADS NA Comp
John Crain	ICANN
Michael Currie	Intrado, Inc.
Dale Drew	Level 3 Communications

Chris Garner	CenturyLink
Joseph Gersch	Secure64 Software Corporation
Jose A. Gonzalez	Sprint Nextel Corporation
Kevin Graves	TeleCommunication Systems (TCS)
Tom Haynes	Verizon
Chris Joul	T-Mobile
Mazen Khaddam	Cox
Kathryn Martin	Access Partnership
Ron Mathis	Intrado, Inc.
Danny McPherson	Verisign
Doug Montgomery	NIST
Chris Oberg	ATIS (Works for Verizon Wireless)
Victor Oppleman	Packet Forensics
Elman Reyes	Internet Identity
Ron Roman	Applied Communication Sciences
Heather Schiller	Verizon
Jason Schiller	Google
Marvin Simpson	Southern Company Services, Inc.
Tony Tauber	Comcast
Paul Vixie	Internet Systems Consortium
Russ White	Verisign
Bob Wright	AT&T

Table 1 - List of Working Group Members

3 Objective, Scope, and Methodology

3.1 Objective

This Working Group was convened to examine and make recommendations to the Council regarding best practices to secure the Domain Name System (DNS) and routing system of the Internet during the period leading up to what some anticipate might be widespread implementation of protocol updates such as the Domain Name System Security Extensions (DNSSEC) and Secure Border Gateway Protocol (BGPsec) extensions (though the latter outcome is not entirely uncontroversial).

DNS is the directory system that associates a domain name with an IP (Internet Protocol) address. In order to achieve this translation, the DNS infrastructure makes hierarchical inquiries to servers that contain this global directory. As DNS inquiries are made, their IP packets rely on routing protocols to reach their correct destination. BGP is the protocol utilized to identify the best available paths for packets to take between points on the Internet at any given moment. This foundational system was built upon a distributed unauthenticated trust model that has been mostly sufficient for over two decades but has some room for improvement.

These foundational systems are vulnerable to compromise through operator procedural mistakes as well as through malicious attacks that can suspend a domain name or IP address's availability, or compromise their information and integrity. While there are formal initiatives under way within the IETF (which has been chartered to develop Internet technical standards and protocols) that will improve this situation significantly, global adoption and implementation will take some time.

This Working Group will examine vulnerabilities within these areas and recommend best practices to better secure these critical functions of the Internet during the interval of time preceding deployment of more robust, secure protocol extensions.

This report covers the BGP portion of these overall group objectives.

3.2 Scope

Working Group 4's charter clearly delineates its scope to focus on two subsets of overall network security, DNS and routing. It further narrows that scope to exclude consideration of the implementation of DNSSEC (tasked to Working Group 5) and secure extensions of BGP (tasked to Working Group 6). While those groups deal with protocol modifications requiring new software and/or hardware deployments; WG4 is geared toward items that either don't require these extensions or are risks which are outside the scope of currently contemplated extensions.

For this report regarding BGP, the focus is on using known techniques within the Operator community. Some of these methods and the risks they seek to address are useful even in cases where protocol extensions are used in some future landscape.

3.3 Methodology

With the dual nature of the work facing Working Group 4, the group was divided into two sub-groups, one focused on issues in DNS security, another in routing security. Starting in December 2011, the entire Working Group met every two weeks via conference call(s) to review research and discuss issues, alternating between sub-groups. The group created a mailing list to correspond and launched a wiki to gather documents and to collectively collaborate on the issues. Additional subject matter experts were occasionally tapped to provide information to the working group via conference calls.

The deliverables schedule called for a series of reports starting in June 2012 that would first identify issues for both routing and DNS security, then enumerate potential solutions, and finally present recommendations. The initial deliverables schedule was updated in March in order to concentrate efforts in each particular area for separate reports. This first report on DNS security issues was presented in September 2012, and this, the second report on routing issues, is being published in March 2013.

Based on the discussions of the group, a list of BGP risks, potential solutions, and relevant BCP documents was created and refined over the course of the work. Subject matter experts in BGP then drove development of the initial documentation of issues and recommendations. These were then brought together into a full document for review and feedback. Text contributions, as completed, were reviewed, edited and approved by the full membership of Working Group 4.

4 Background

4.1 Deployment Scenarios

BGP is deployed in many different kinds of networks of different size and profiles. Many different recommendations exist to improve the security and resilience of the inter-domain routing system. Some of the advice can even appear somewhat contradictory and often the key decision can come down to understanding what is most important or appropriate for a given network considering its size, the number of external connections, number of BGP routers, size

and expertise of the staff and so forth.

We attempt to tailor the recommendations and highlight which are most significant for a given network operator's situation. Further background and information on routing operations can be found in the Appendix (Section 7) of this document for readers unfamiliar with this area of practice.

5 Analysis, Findings and Recommendations

The primary threats to routing include:

- Risks to the routers and exchange of routing information
- Routing information that is incorrect or propagates incorrectly
- General problems with network operations

5.1 BGP Session-Level Vulnerability

When two routers are connected and properly configured, they form a BGP peering session. Routing information is exchanged over this peering session, allowing the two peers to build a local routing table, which is then used to forward actual packets of information. The first BGP4 attack surface is the peering session between two individual routers, along with the routers themselves. Two classes of attacks are included here, *session hijacking* and *denial of service.*

5.1.1 Session Hijacking

The BGP session between two routers is based on the *Transport Control Protocol* (TCP), a session protocol also used to transfer web pages, naming information, files, and many other types of data across the Internet. Like all these other connection types, BGP sessions can be hijacked or disrupted, as shown in Figure 1.

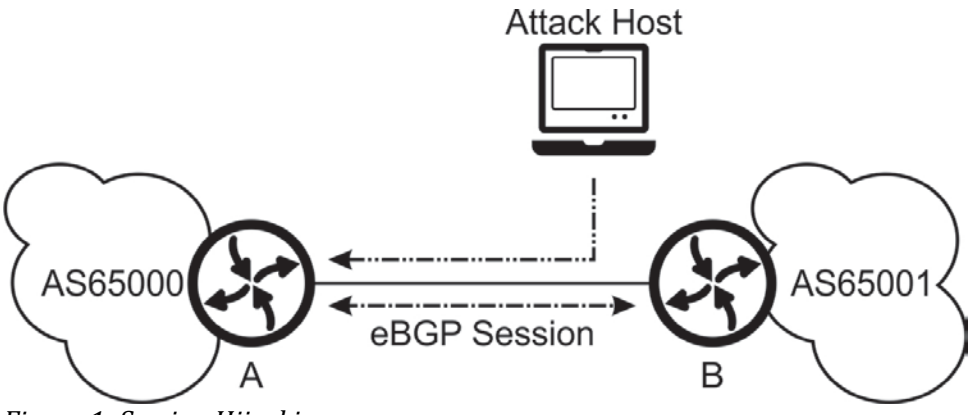

Figure 1: Session Hijacking

In this diagram, the attack host can either take over the existing session between Routers A and B, or build an unauthorized session with Router A. By injecting itself into the peering between Routers A and B, the attacker can inject new routing information, change the routing information being transmitted to Router A, or even act as a "man in the middle," modifying the routing information being exchanged between these two routers.

5.1.1.1 Session-Level Countermeasures

Current solutions to these types of attacks center on secure hash mechanisms, such as HMAC-MD5 (which has been deprecated) and HMAC-SHA. These mechanisms rely on the peering routers sharing a key (a *shared key* – essentially a password) that is used to calculate a cryptographic hash across each packet (or message) transmitted between the two routers, and included in the packet itself. The receiving router can use the same key to calculate the hash. If the hash in the packet matches the locally calculated hash, the packet could have only been transmitted by another router that knows the shared key.

This type of solution is subject to a number of limitations. First, the key must actually be shared among the routers building peering sessions. In this case, the routers involved in the peering session are in different administrative domains. Coming to some uniform agreement about how keys are generated and communicated (e.g. phone, email, etc.) with the often hundreds of partners and customers of an ISP is an impractical task.

There is the possibility that any key sharing mechanism deployed to ease this administrative burden could, itself, come under attack (although such attacks have never been seen in the wild).

Lastly, some concerns have been raised that burden of cryptographic calculations could itself become a vector for a Denial-of-Service (DoS) attack by a directed stream of packets with invalid hash components. One way to deny service is to make the processor that is responsible for processing routing updates and maintaining liveness too busy to reliably process these updates in a timely manner. In many routers the processor responsible for calculating the cryptographic hash is also responsible for processing new routing information learned, sending out new routing information, and even transmitting keep-alive messages to keep all existing sessions up. Since calculating the cryptographic hash is computationally expensive, a smaller flood of packets with an invalid hash can consume all the resources of the processor, thus making it easier to cause the processor to be to busy.

Other mechanisms, such as the Generalized TTL Security Mechanism (GTSM, described in RFC 5082), focus on reducing the scope of such attacks. This technique relies on a feature of the IP protocol that would prevent an attacker from effectively reaching the BGP process on a router with forged packets from some remote point on the Internet. Since most BGP sessions are built across point-to-point links (on which only two devices can communicate), this approach would prevent most attackers from interfering in the BGP session. Sessions built over a shared LAN, such as is the case in some Internet exchanges, will be protected from those outside the LAN, but will remain vulnerable to all parties that are connected to the LAN.

This solution is more complicated to implement when BGP speaking routers are not directly connected. It is possible to count the number of hops between routers and limit the TTL value to only that number of hops. This will provide some protection, limiting the scope of possible attackers to be within that many hops. If this approach is used consider failure scenarios of devices between the pair of BGP speaking routers, what impact those failures will have on the hop count between the routers, and if you want to expand the TTL value to allow the session to remain up for a failure that increases the hop count.

5.1.1.2 Session-Level Current Recommendations

For a network with a small (e.g., single-digit) number of eBGP neighbors, it is reasonable to follow the lead of what is specified by the upstream ISPs who may have a blanket policy of how they configure their eBGP sessions. A network with larger numbers of eBGP neighbors may be satisfied that they can manage the number of keys involved either through data-store or rubric. Note that a rubric may not always be feasible as you cannot ensure that your neighbors will always permit you to choose the key.

Managing the keys for a large number of routers involved in BGP sessions (a large organization may have hundreds or thousands of such routers) can be an administrative burden. Questions and issues can include:

- In what system should the keys be stored and who should have access?
- Should keys be unique per usage having only one key for internal usage and another key that is shared for all external BGP sessions?
- Should keys be unique per some geographical or geo-political boundary say separate keys per continent or per country or per router?
- Should keys be unique to each administrative domain, for example a separate key for each Autonomous System a network peers with?
- There is no easy way to roll over keys, as such changing a key is quite painful, as it disrupts the transmission of routing information, and requires simultaneous involvement from parties in both administrative domains. This makes questions of how to deal with the departure of an employee who had access to the keys, or what keys to use when peering in a hostile country more critical.

Another consideration is the operational cost of having a key. Some routing domains will depend on their peers to provide the key each time a new session is established, and not bother to make a record of the key. This avoids the problems of how to store the key, and ensure the key remains secure. However if a session needs to be recreated because configuration information is lost either due to accidental deletion of the configuration, or hardware replacement, then the key is no longer known. The session will remain down until the peer can be contacted and the key is re-shared. Often times this communication does not occur, and the peer may simply try to remove the key as a troubleshooting step, and note the session re-establishes. When this happens the peer will often prefer for the session to remain up, leaving the peering session unsecured until the peer can be contacted, and a maintenance window can be scheduled. For unresponsive peers, an unsecured peering session could persist, especially considering that the urgency to address the outage has now passed.

Despite these vulnerabilities having been widely known for a decade or more, they have not been implicated in any notable number of incidents. As a result some network operators have not found the cost/benefit trade-offs to warrant the operational cost of deploying such mechanisms while others have. Given these facts, the Working Group recommends that individual network operators continue to make their own determinations in using these counter-measures.

5.1.2 Denial of Service (DoS) Vulnerability

Because routers are specialized hosts, they are subject to the same sorts of Denial of Service (DoS) attacks any other host or server is subject to. These attacks fall into three types:

1. Attacks that seek to consume all available interface bandwidth making it difficult for enough legitimate traffic to get through such as UDP floods and reflective attacks
2. Attacks that seek to exhaust resources such as consume all available CPU cycles, memory, or ports so that the system is too busy to respond such as TCP SYN attacks
3. Attacks utilizing specially crafted packets in an attempt to cause the system to crash or operate in an unexpected way such as buffer overflow attacks, or malformed packet attacks that create an exception that is not properly dealt with

Bandwidth exhaustion attacks attempt to use so much bandwidth that there is not enough available bandwidth for services to operate properly. This type of an attack can cause routers to fail to receive routing protocol updates or keep-alive messages resulting in out-of-date routing information, routing loops, or interruption of routing altogether, such as happens when a BGP session goes down and the associated routing information is lost.

Resource exhaustion attacks target traffic to the router itself, and attempt to make the router exhaust its CPU cycles or memory. In the case of the former, the router's CPU becomes too busy to properly process routing keepalives and updates causing the adjacencies to go down. In the case of the latter, the attacker sends so much routing protocol information that the router has no available memory to store all of the required routing information.

Crafted packets attacks attempt to send a relatively small number of packets that the router does not deal with appropriately. When a router receives this type of packet it may fill up interface buffers and then not forward any traffic on that interface causing routing protocols to crash and restart, reboot, or hang. In some cases the router CPU may restart, reboot, or hang likely causing loss of all topological and routing state. One example was the "protocol 55" attack, where some router vendors simply did not code properly how to deal with this traffic type.

Some routers are specialized to forward high rates of traffic. These routers often implement their forwarding capabilities in hardware that is optimized for high throughput, and implement the less demanding routing functions in software. As such bandwidth exhaustion attacks are targeted at the routers interfaces or the backplane between those interfaces, or the hardware responsible for making forwarding decisions. The other types of attack target the software responsible for making the routing decisions.

Due to the separation between routing and forwarding, a fourth class of attacks are targeted at exhausting the bandwidth of the internal interconnection between the forwarding components and the routing components.

The section below on "Denial-of-Service Attacks on ISP Infrastructure" contains a discussion of disruptive attacks besides those targeting the exchange of BGP routing information.

5.1.2.1 Denial of Service Countermeasures

GTSM, described above, can be an effective counter to some forms of DoS attacks against routers, by preventing packets originating outside the direct connection between two BGP peers from being processed by the router under attack. GTSM cannot resolve simple buffer overflow problems, or DoS attacks that exploit weaknesses in packet processing prior to the TTL check, however.

Another mechanism currently used to prevent DoS attacks against routers is to simply make the interfaces on which the BGP session is running completely unreachable from outside the local network or the local segment. Using link-local addresses in IPv6, is one technique (with obviously limited applicability). Another approach is applying packet-filters on the relevant address ranges at the network edge. (This process is called *infrastructure filtering).*
Other well-known and widely deployed DoS mitigation techniques can be used to protect routers from attack just as they can be used to protect other hosts. For instance, *Control Plane Policing* can prevent the routing process on a router from being overwhelmed with high levels of traffic by limiting the amount of traffic accepted by the router directed at the routing processor itself.

5.1.2.2 Denial-of-Service Current Recommendations

Since routers are essentially specialized hosts, mechanisms that can be used to protect individual routers and peering sessions from attack are widely studied and well understood. What prevents these techniques from being deployed on a wide scale?

Two things: the perception that the problem space is not large enough to focus on, and the administrative burden of actually deploying such defenses. For instance, when GTSM is used with infrastructure filtering, cryptographic measures may appear to be an administrative burden without much increased security. Smaller operators, and end customers, often believe the administrative burden too great to configure and manage any of these techniques.

Despite these vulnerabilities having been widely known for a decade or more, they have not been implicated in any notable number of incidents. As a result some network operators have not found the cost/benefit trade-offs to warrant the operational cost of deploying such mechanisms while others have. Given these facts, the Working Group recommends that individual network operators continue to make their own determinations in using these counter-measures.

In dealing with vulnerabilities due to “crafted packets”, the vendor should provide notification to customers as the issues are discovered as well as providing fixed software in a timely manner.

Customers should make it a point to keep abreast of notifications from their vendors and from various security information clearing-houses.

5.1.2.2.1 Interface Exhaustion Attacks

Recommendations include:

1. Understanding the actual forwarding capabilities of your equipment in your desired configuration
2. Examining your queuing configuration
3. Carefully considering which types of traffic share a queue with your routing protocols, and if that traffic can be blocked, rate-limited or forced to another queue
4. Understanding packet filtering capabilities of your equipment, and under what scenarios it is safe to deploy packet filters
5. When it is safe to do so, tactically deploy packet filers upstream from a router that is being attacked

The first thing to consider with regard to attacks that attempt to consume all available bandwidth is to determine the actual throughput of the router. It is not safe to assume that a router with two or more 100 Gigabit Ethernet interfaces can receive 100G on one interface and transmit that same 100G out another interface. Some routers can forward at line rate, and some routers cannot. Performance may vary with packet size, for example forwarding traffic in software is more taxing on the CPU as the number of packets increases.

The next thing to consider is outbound queuing of routers upstream from the router that is being attacked. Routers typically place routing protocol traffic in a separate Network Control (NC) queue. Determine the characteristics of this queue such as the queue depth and frequency of it being serviced. These values may be tunable. Also consider what types of traffic are placed in this queue, and specifically what traffic an outside attacker can place in this queue. Consider preventing users of the network from being able to place traffic in this queue if they do not need to exchange routing information with your network. For direct customers running eBGP, limit traffic permitted into the NC queue to only traffic required to support their routing protocols. Consider rate-limiting this traffic so no one customer can fill up the queue. Note that rate limits will increase convergence time, so test a customer configuration that is advertising and receiving the largest set of routes, and measure how long it takes to re-learn the routing table after the BGP session is reset with and without the rate limits.

If the attack traffic has a particular profile, and all traffic matching that profile can be dropped without impacting legitimate traffic than a packet filter can be deployed upstream from the router that is under attack. Ensure that a deploying a packet filter will not impact the performance of your router by testing packets filters with various types of attacks and packet sizes on your equipment in a lab environment. Ensure that total throughput is not decreased, and that there is not a particular packet per second count that causes the router to crash, or become unresponsive, or stop forwarding traffic reliably, cause routing protocols to time out, etc.

If the upstream router belongs to a non-customer network, you will need to work with them to mitigate the attack. Additional bandwidth on the interconnect may allow you to move the bottleneck deeper into your network where you can deal with it.

Often the IP destination of these attacks is something downstream from the router. It is possible that some or all of the attack traffic may be destined to the router. In that case some of the mitigation techniques in the next section may also be helpful.

5.1.2.2.2 Resource Exhaustion Attacks

Recommendations include:

1. Consider deploying GTSM
2. Consider making router interfaces only reachable by directly connected network
3. Consider only permitting traffic sourced from configured neighbors
4. Consider deploying MD5
5. Deploy maximum prefixes

The first set of recommendations is to consider deploying mechanisms that restrict who can send routing protocol traffic to a router. The second set of recommendations restricts how much routing protocol state a neighbor can cause a router to hold.

GTSM, described above, can be an effective counter to some forms of DoS attacks against routers, by limiting who can send routing protocol traffic to the router by a configured hop-count radius. GTSM works by preventing packets originating outside the direct connection between two BGP peers from being processed by the router under attack. GTSM cannot resolve simple buffer overflow problems, or DoS attacks that exploit weaknesses in packet processing prior to the TTL check, however.

Another mechanism currently used to prevent DoS attacks against routers is to simply make the interfaces on which the BGP session is running completely unreachable from outside the local network or the local segment. Using link-local addresses in IPv6, is one technique (with obviously limited applicability). Another approach is applying packet-filters on the relevant address ranges at the network edge. (This process is called *infrastructure filtering).*

Some routers can dynamically generate packet filters from other portions of the router configuration. This enables one to create an interface packet filter that only allows traffic on the BGP ports from source IP addresses that belong to a configured neighbor. This means attempts to send packets to the BGP port by IP addresses that are not a configured neighbor will be dropped right at the interface.

The same session level protections discussed earlier, such as MD5 can also limit who can send routing information to only those routers or hosts that have the appropriate key. As such this is also an effective mechanism to limit who can send routing protocol traffic. While these packets will be processed by the router, and could possibly tax the CPU, they cannot cause the router to create additional routing state such as adding entried to the Routing Information Base (RIB) or Forwarding Information Base (FIB).

Lastly each neighbor should be configured to limit the number of prefixes they can send to a reasonable value. A single neighbor accidentally or intentionally de-aggregating all of the address space they are permitted to send could consume a large amount of RIB and FIB memory, especially with the large IPv6 allocations.

5.1.2.2.3 Crafted Packet Attacks

Crafted packet attacks typically occur when a router receives some exception traffic that the vendor did not plan for. In some cases it may be possible to mitigate these attacks by filtering the attack traffic if that traffic has a profile that can be matched on and all traffic matching that profile can be discarded. More often then not, this is not the case.

In all cases the vendor should provide new code that deals with the exception.

Recommendations include keeping current with all SIRT advisories from your vendors. When vulnerability is published move quickly to upgrade vulnerable versions of the code. This may require an upgrade to a newer version of the code or a patch to an existing version.

For larger organizations that have extensive and lengthy software certification programs, it is often more reasonable to ask the vendor to provide a patch for the specific version(s) of code that organization is running. If possible the vendor should provide the extent to which the code is modified to quantify how substantial the change is in order to help the provide plan what should be included in the abbreviated software certification tests.

For smaller organizations, or organizations that complete little or no software certification, the newer version of code with the fix in place should be deployed. Generally this is deployed cautiously at first to see if issues are raised with a limited field trail, followed by a more widespread deployment.

5.1.2.2.4 Internal Bandwidth Exhaustion Attacks

Other well-known and widely deployed DoS mitigation techniques can be used to protect routers from attack just as they can be used to protect other hosts. For instance, *Control Plane Policing* can prevent the routing process on a router from being overwhelmed with high levels of traffic by limiting the amount of traffic accepted by the router directed at the routing processor itself.

One should consider not only the impact on the router CPU, but also the impact on the bandwidth between the forward components and the router CPU. There may be some internal queuing in place on the interconnect between the forwarding components and the router CPU. It may be possible to influence which queue routing protocol traffic is placed in, or with queue traffic generated by customers is placed in, the depths and or servicing of these queues in order to separate and minimize the ability of non-routing traffic to impact routing traffic. If traffic generated by customers (for routing protocols or otherwise) can crowd out a network's internal routing protocol traffic, then operators may consider separately rate limiting this customer traffic.

5.1.3 Source-address filtering

Many Internet security exploits hinge on the ability of an attacker to send packets with spoofed source IP addresses. Masquerading in this way can give the attacker entre to unauthorized access at the device or application level and some BGP vulnerabilities are also in this category. The problem of source-spoofing has long been recognized and countermeasures available for filtering at the interface level.

5.1.3.1 Source-address spoofing example

Consider the diagram below which illustrates the legitimate bi-directional traffic flow between two hosts on the left-hand side. An attacker connected to another network can send IP packets with a source address field set to the address of one of the other machines, unless filtering is applied at the point that that attacker's host or network attaches to AS65002.

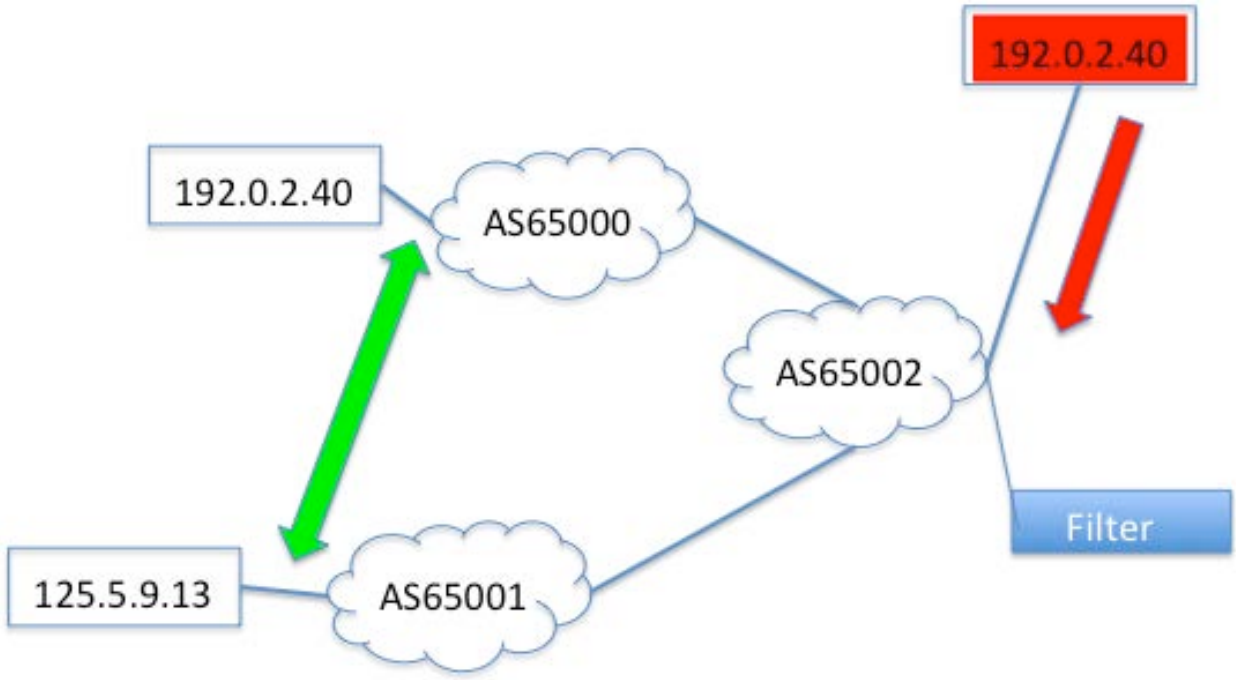

5.1.3.2 Source-address spoofing attacks

Though most IP transactions are bi-directional, attacks utilizing spoofed source IPs do not require bi-directional communication but instead exploit particular protocol or programmatic semantic weaknesses.
Exploits using this technique have covered many areas over the years including these types followed by some examples.

- Attacks against services which rely only on source IP of the incoming packet for authorization
- rsh, rlogin, NFS, Xwindows, etc.
- Attacks where the unreachability of the source can be exploited
- TCP SYN floods which exhaust resources on the server
- Attacks where the attacker masquerades as the “victim”
- Small DNS or SNMP requests resulting in highly asymmetric data flow back toward the victim
- Abusive traffic which result in the legitimate user getting blocked from the server or network

5.1.3.3 Source-address filtering challenges

The barriers to implementing these countermeasures have ranged from lack of vendor support to lack of solid motivation to implement them.

- Lack of proper vendor support: In older implementations of network devices, filtering based on the source address of a packet was performed in software, rather than hardware, and thus had a major impact on the rate of forwarding through the device. Most modern

network equipment can perform source filtering in the hardware switching path, eliminating this barrier to deployment.

- Lack of scalable deployment and configuration management: In older deployments, filters based on the source of traffic was configured and managed manually, adding a large expense to the entity running the network. This barrier has largely been resolved through remote triggered black hole, unicast Reverse Path Forwarding (uRPF), and loose uRPF options.
- Fear of interrupting legitimate traffic, for example in multi-homed situations: Vendors have created flexibility in uRPF filtering to reduce or eliminate this barrier. Future possible additions include "white lists," which would allow traffic to pass through a uRPF check even though it didn't meet the rules.
- Lack of business motivation: Unilateral application of these features does not benefit or protect a network or its customers directly; rather, it contributes to the overall security of the Internet. Network operators are realizing that objection to incurring this "cost" is being overcome by the realization that if everyone performs this type of filtering, then everyone benefits.

5.1.3.4 Source-address filtering recommendations

Filtering should be applied as close as possible to the entry point of traffic. Wherever one host, network, or subnet is attached, a feature such as packet filtering, uRPF, or source-address-validation should be used. Ensure adequate support from equipment vendors for subscriber-management systems (e.g. for Cable and DSL providers) or data-center or campus routers and switches.

Stub networks should also filter traffic at their borders to ensure IP ranges assigned to them *do not* appear in the source field of incoming packets and *only* those ranges appear in the source field of outgoing packets.

Transit networks should likewise use features such as uRPF. Strict mode should be used at a border with a topological stub network and loose mode between transit networks.

Transit networks that provide connectivity primarily stub networks, such as consumer ISPs, should consider uRPF strict mode on interfaces facing their customers. If these providers provide a home router to their customers they should consider making uRPF part of the default home router configuration.

Transit networks that provide connectivity to a mix of stub networks and multi-homed networks must consider the administrative burden of configuring uRPF strict mode only on stub customers and uRPF loose mode, or no uRPF on customers that are, or become multi-homed.

When using uRPF loose mode with the presence of a default route, one must special care to consider configuration options to include or exclude the default route.

The value of loose mode uRPF with networks in the default free zone is debatable. It will only prevent traffic with a source address of RFC-1918 space and dark IPs (IP addresses that are not routed on the Internet). Often these dark IP addresses are useful for backscatter techniques and

tracing the source(s) of a spoofed DoS attack. It is also important to consider if RFC-1918 addresses are used internal to the transit provider's network. This practice may become more common if ISPs implement Carrier Grade NAT.

It is also worth pointing out that some business customers depend on VPN software that is poorly implemented, and only changes the destination IP address when re-encapsulating a packet. If these customers are using non-routed IP addresses in their internal network then enabling uRPF will break these customers.

It is important to measure the impact of forwarding when enabling uRPF. Even when uRPF is implemented in hardware, the router must lookup the destination as well as the source. A double lookup will cause forwarding throughput to be reduced by half. This may have no in the forwarding rate if the throughput of the forwarding hardware is more than twice the rate of all the interfaces it supports.
Further, more detailed advice and treatment of this subject can be found in:

- IETF BCP38/RFC 2827 Network Ingress Filtering[1]
- BCP 84/RFC 3704 Ingress Filtering for Multihomed Networks[2]

5.2 ICANN SAC004 Securing the EdgeBGP Injection and Propagation Vulnerability

A second form of attack against the routing information provided by BGP4 is through injection of misleading routing information, as shown in figure 2.

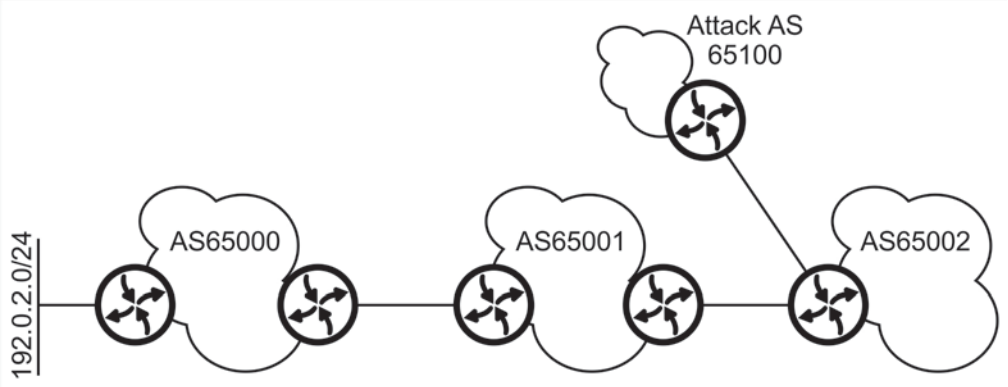

Figure 2: A Prefix Hijacking Attack

In this network, AS65000 has authority to originate 192.0.2.0/24. Originating a route, in this context, means that computers having addresses within the address space advertised are actually reachable within your network — that a computer with the address 192.0.2.1, for instance, is physically attached to your network.
Assume AS65100 would like to attract traffic normally destined to a computer within the 192.0.2.0/24 address range. Why would AS65100 want to do this? There are a number of possible motivations, including:

[1] http://tools.ietf.org/html/bcp38
[2] http://tools.ietf.org/html/bcp84

- A server with an address in this range accepts logins from customers or users, such as a financial web site, or a site that hosts other sensitive information, or information of value
- A server with an address in this range processes information crucial to the operation of a business the owner of AS65100 would like to damage in some way, such as a competitor, or a political entity under attack

AS65100, the attacker, can easily attract packets normally destined to 192.0.2.0/24 within AS65000 by simply advertising a competing route for 192.0.2.0/24 to AS65002. From within BGP itself, there's no way for the operators in AS65002 to know which of these two advertisements is correct (or whether both origins are valid – a configuration which does see occasional legitimate use). The impact of the bogus information may be limited to the directly neighboring AS(es) depending on the routing policy of the nearby ASes. The likelihood of the incorrect route being chosen can be improved by two attributes of the route:

- A shorter AS Path

A shorter AS Path has the semantic value of indicating a topologically "closer" network. In the example above, the normal propagation of the route would show AS65100 as "closer" to AS65001 thus, other factors being equal, more preferred than the legitimate path via AS65000.

- A longer prefix

Longer prefixes represent more-specific routing information, so a longer prefix is always preferred over a shorter one. For instance, in this case the attacker might advertise 192.0.2.0/25, rather than 192.0.2.0/24, to make the false route to the destination appear more desirable than the real one.

- A higher local-preference setting

Local-preference is the non-transitive BGP attribute that most network operators use to administratively influence their local routing. Typically, routes learned from a "customer" (i.e., paying) network are preferred over those where the neighboring network has a non-transit relationship or where the operator is paying for transit from the neighboring network. This attribute is more important in the decision algorithm for BGP than AS-path length so routes learned over such a session can draw traffic even without manipulation of the AS-path attribute.

This illustration can be used to help describe some related types of risks:

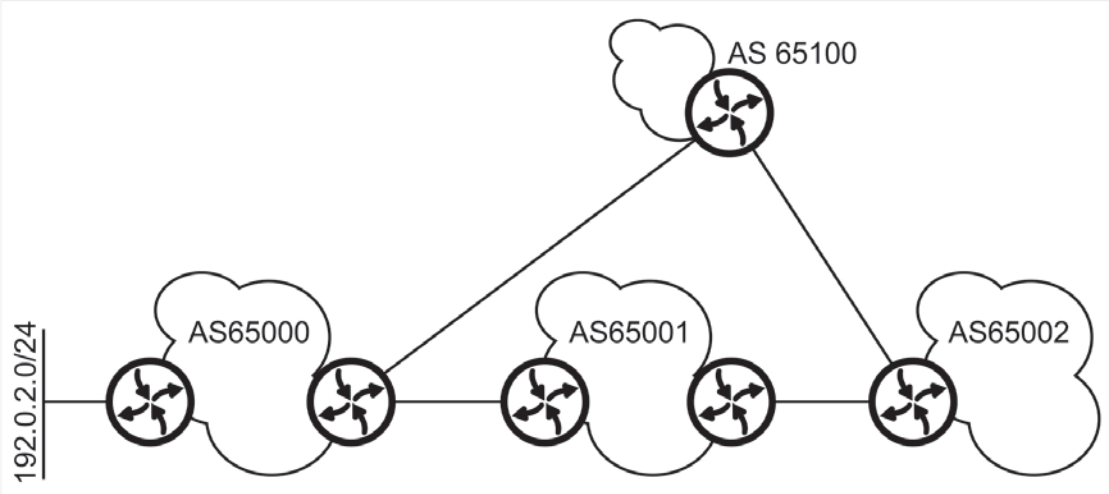

Figure 3: BGP Propagation Vulnerabilities

Route Leak: In this case, AS65100 is a transit customer of both AS65000 and AS65002. The operator of AS65100 accidentally leaks routing information advertised by AS65000 into its peering session with AS65002. This could possibly draw traffic passing from AS65002 towards a destination reachable through AS65000 through AS65100 when this path was not intended to provide transit between these two networks. Most often such happenings are the result of misconfigurations and can result in overloading the links between AS65100 and the other ASes. As the name suggests, this phenomenon has most often been the result of inadvertent mis-configuration.

Occasionally they can result in more malicious outcomes:

- **Man in the Middle:** In this case, all the autonomous systems shown have non-transit relationships. For policy reasons, AS65000 would prefer traffic destined to 192.0.2.0/24 pass through AS65100. To enforce this policy, AS65000 filters the route for 192.0.2.0/24 towards AS65100. In order to redirect traffic through itself (for instance, in order to snoop on the traffic stream), AS65100 generates a route advertisement that makes it appear as if AS65000 has actually advertised 192.0.2.0/24, advertising this route to AS65002, and thus drawing traffic destined to 192.0.2.0/24 through itself.
- **False Origination:** This attack is similar to the man in the middle explained above, however in this case there is no link between AS65000 and AS65100. Any traffic destined to 192.0.2.0/24 into AS65100, is discarded rather than being delivered.

Note that all three of these vulnerabilities are variations on a single theme: routing information that should not be propagated based on compliance with some specific policy nonetheless is.

5.2.1 BGP Injection and Propagation Countermeasures

5.2.1.1 Prefix Filtering

The key bulwark against entry and propagation of illegitimate routing announcements into the global routing system is prefix-level filtering; typically at the edge between the ISP and their customers. The usual method involves the customer communicating a list of prefixes and downstream ASes which they expect to be reachable through the connection to the ISP. The ISP will then craft a filter applied to the BGP session which explicitly enumerates this list of expected prefixes (a “prefix-list”), perhaps allowing for announcement of some more-specific

prefixes within the ranges such as might be needed by the customer to achieve some goals in adjusting the load of the customer's inbound traffic across their various connections. This configuration forms a "white-list", in security parlance, of possible downstream destinations but does not validate the overall semantic correctness of the resulting routing table.

5.2.1.1.1 Manual Prefix-Filter Limitations

The validity of the information in this list is obviously important and if the customer is either malicious or simply mistaken in the prefixes they communicate, the prefix filter could obviously still leave open a vulnerability to bogus route injection. Thus typically the information communicated by the customer is checked against registration records such as offered in the "whois" information available from Regional Internet Registries (RIRs) and/or others in the address assignment hierarchy. However, there is no information in the RIR information explicitly indicating a mapping between the address-assignment and the origin AS.

The reality is that the process of manually checking a filter that is a few thousand lines long, with hundreds of changes a week is tedious and time consuming. Many transit providers do not check at all. Others have a policy to always check, but support staff may grow complacent with updates from certain customers that have long filter lists or have filter lists that change weekly, especially if those customers have never had a questionable prefix update in the past. In these cases they may only spot-check, only check the first few changes and then give up, or grow fatigued and be less diligent the more lines they check.

Moreover, the entity name fields in the "whois" information are free-form and often can't be reliably matched to the entity name used by the ISP's customer records. Typically it is a fuzzy match between the customer name on record and the company name listed in the "whois" record. A truly malicious actor could order service with a name which is intentionally similar to a company name whose IP addresses they intend to use. Another possibility is that the company name on record is legitimate and an exact match to the company name on the "whois" record, but the customer is a branch office, and the legitimate holder of the IP addresses is the corporate office which has not authorized the branch office to use their space. It is also possible that the customer is the legitimate holder of the address space, but the individual who called in to the provider support team is not authorized to change the routing of the IP block in question. This problem is further complicated when a transit provider's customer has one or more downstream customers of its own. These relationships are typically hard or impossible to verify.

If every transit provider accurately filtered all of the prefixes their customers advertised, and each network that a transit provider peers with could be trusted to also accurately filter all of the prefixes of their customers, then route origination and propagation problems could be virtuailly eliminated. However, managing filters requires thousands of operators examining, devising, and adjusting the filters on millions of devices throughout the Internet. While there are processes and tools within any given network, such highly inconsistent processes, particularly when handling large amounts of data (a tedious process in and of itself), tends to produce an undesirable rate of errors. Each time an individual operator misjudges a particular piece of information, or simply makes a mistake in building a filter, the result is a set of servers (or services) that are unreachable until the mistake is found and corrected.

5.2.1.2 Internet Routing Registry (IRR)

The second source of information a provider can use as a basis for filtering received routing

information is a voluntary set of databases of routing policy and origination called Internet Routing Registries (IRRs). These IRRs allow providers to register information about their policies towards customers and other providers, and also allow network operators to register which address space they intend to originate. Some providers require their customers to register their address space in an IRR before accepting the customer's routes, oftentimes the provider will "proxy register" information on the customer's behalf since most customers are not versed in IRR details.

5.2.1.2.1 IRR Limitations

Because IRRs are voluntary, there is some question about the accuracy and timeliness of the information they contain (see *Research on Routing Consistency in the Internet Routing Registry* by Nagahashi and Esaki for a mostly negative view, and *How Complete and Accurate is the Internet Routing Registry* by Khan for a more positive view). Anecdotally, RIPE's IRR is in widespread use today, and some large providers actually build their filtering off this database, so the accuracy level is at least operationally acceptable for some number of operators. Some IRR repositories use an authorization model as well as authentication but none that primarily serve North America perform RPSL authorization using the scheme described in RFC2725 – *Routing Policy System Security*[3].

5.2.1.3 AS-Path filtering

Filters on the AS_PATH contents of incoming BGP announcements can also be part of a defensive strategy to guard against improper propagation of routing information. Some ISPs have used AS-path filters on customer-facing BGP sessions *instead of* prefix-filters. This approach is generally inadequate to protect against even the most naïve misconfigurations, much less a deliberate manipulation. Often a leak has involved either redistributing BGP routes inadvertently from one of a stub network's ISPs to the other. Another problem in the past has involved redistributing BGP routes into an internal routing protocol and back to BGP.

Where AS-path filters can be useful is to guard against an egregious leak. For instance an ISP would not expect ASNs belonging to known large ISPs to show up in the AS_PATH of updates from an enterprise-type customer network. Applying an AS-path filter to such a BGP session could act as a second line of defense to the specific prefix-list filter. Similarly, if there are networks which the ISP has non-transit relationships with, applying a similar AS-path filter to those sessions (which wouldn't be candidates for prefix filters) could help guard against a leak resulting in an unintended transit path.

5.2.1.3.1 AS-Path filtering limitations

Maintaining such a list of "known" networks which aren't expected to show up in transit adjacencies can be fairly manual, incomplete and error-prone. Again, applying a filter which validates the neighbor AS is in the path is useless since this state is the norm of what's expected.

5.2.1.4 Maximum-prefix cut-off threshold

Many router feature-sets include the ability to limit the number of prefixes that are accepted from a neighbor via BGP advertisements. When the overall limit is exceeded, the BGP session is torn down on the presumption that this situation is a dangerous error condition. Typically also a threshold can be set at which a warning notification (e.g. log message) to the Operations staff

[3] https://tools.ietf.org/html/rfc2725

is issued. This way a gradual increase in the number of advertisements will trigger a sensible manual raise in the cut-off threshold without causing an outage.

This tool can be used to guard against the most egregious leaks which can, if the numbers are large enough, exhaust the routing table memory on the recipient's routers and/or otherwise cause widespread network instability.

Typical deployments will set the threshold based on the current observed number of advertisements within different bands; for instance, 1-100, 100-1000, 1000-5000, 5000-10000, 10000-50000, 50,000-100,000, 100,000-150,000, 150,000-200,000, 200,000-250,000, 250,000-300,000.

5.2.1.4.1 Maximum-prefix limitations

When the threshold is exceeded, the session is shut down and manual intervention is required to bring it back up. In the case where a network has multiple interconnection points to another network (thus multiple BGP neighbors), all sessions will typically go down at the same time assuming all are announcing the same number of prefixes. In this case, it may be the case that all connectivity between the two networks is lost during this period. Obviously this measure is an attempt to balance two different un-desirable outcomes so must be weighed judiciously.

Above 10000 or perhaps 50000 (e.g., a full Internet routing table from a transit provider), applying maximum-prefix thresholds provide limited protection. A small number of neighbors each advertising a unique set of 300,000 routes would fill the memory of the receiving router anyway. However if these neighbors are all advertising a large portion of the Internet routes, with many routes overlapping, then the limit offers some protection.

5.2.1.5 Monitoring

Aside from a proactive filtering approach, a network operator can use various vantage points external to their own network (e.g, "route servers" or "looking glasses") to monitor the prefixes for which they have authority to monitor for competing announcements which may have entered the BGP system. Some tools such as BGPmon have been devised to automate such monitoring.

5.2.1.5.1 Monitoring Limitations

Obviously, this approach is reactive rather than proactive and steps would then need to be taken to contact the offending AS and/or intermediate AS(es) to stop the advertisement and/or propagation of the misinformation. Also, the number of such vantage points is limited so a locally impacting bogus route may or may not be detected with this method.

5.2.2 BGP Injection and Propagation Recommendations

The most common router software implementations of BGP do not perform filtering of route advertisements, either inbound or outbound, by default. While this situation eases the burden of configuration on network operators (the customers of the router vendors), it has also caused the majority of unintentional inter-domain routing problems to date. Thus it is recommended that network operators of all sizes take extra care in configuration of BGP sessions to keep unintentional routes from being injected and propagated.

Stub network operators should configure their outbound sessions to only explicitly allow the

prefixes which they expect to be advertising over a particular session.

ISPs should explicitly filter their inbound sessions at the boundary with their "customer edge". The inter-provider connections between large ISPs are impractical locations for filtering given the requirement for significant dynamism in BGP routing and traffic-engineering across the global Internet. However, the cumulative gains accrued when each ISP filters at this "customer edge" are significant enough to lessen the residual risk of not filtering on these "non-customer" BGP sessions.

ISPs (and even stub networks) should also consider using AS-path filters and maximum-prefix limits on sessions as a second line of defense to guard against leaks or other pathological conditions.

5.3 Other Attacks and Vulnerabilities of Routing Infrastructure

There are many vulnerabilities and attack vectors that can be used to disrupt the routing infrastructure of an ISP outside of the BGP protocol and routing-specific operations. These are just as important to address as issues the working group has identified within the routing space itself.

The largest attack surface for routing infrastructure likely lies within the standard operational security paradigm that applies to any critical networked asset. Therefore the working group looked at including BCPs relating to network and operational security as part of addressing these issues, and ISPs should be aware that they are likely to see attacks against their routing infrastructure based on these "traditional" methods of computer and network intrusion.

5.3.1 Hacking and unauthorized 3rd party access to routing infrastructure

ISPs and all organizations with an Internet presence face the ever-present risk of hacking and other unauthorized access attempts on their infrastructure from various actors, both on and off network. This was already identified as a key risk for ISPs, and CSRIC 2A – Cyber Security Best Practices was published in March 2011 to provide advice to address these types of attacks and other risks for any ISP infrastructure elements, including routing infrastructure. The current CSRIC III has added a new Working Group 11 that will report out an update to prior CSRIC work in light of recent advancements in cybersecurity practices and a desire of several US government agencies to adopt consensus guidelines to protect government and critical infrastructure computers and networks.

A recent SANS publication, *Twenty Critical Security Controls for Effective Cyber Defense: Consensus Audit Guidelines (CAG)*[4] lays out these principals and maps them out versus prior work, including another relevant document, NIST SP-800-53 *Recommended Security Controls for Federal Information Systems and Organizations.*[5] The SANS publication appears to be a primary driver for Working Group 11's work. The entire document is available for review, and we have included the 20 topic areas here for reference:

[4] http://www.sans.org/critical-security-controls/
[5] http://csrc.nist.gov/publications/PubsSPs.html

Critical Control 1: Inventory of Authorized and Unauthorized Devices
Critical Control 2: Inventory of Authorized and Unauthorized Software
Critical Control 3: Secure Configurations for Hardware and Software on Laptops, Workstations, and Servers
Critical Control 4: Continuous Vulnerability Assessment and Remediation
Critical Control 5: Malware Defenses
Critical Control 6: Application Software Security
Critical Control 7: Wireless Device Control
Critical Control 8: Data Recovery Capability
Critical Control 9: Security Skills Assessment and Appropriate Training to Fill Gaps
Critical Control 10: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches
Critical Control 11: Limitation and Control of Network Ports, Protocols, and Services
Critical Control 12: Controlled Use of Administrative Privileges
Critical Control 13: Boundary Defense
Critical Control 14: Maintenance, Monitoring, and Analysis of Audit Logs
Critical Control 15: Controlled Access Based on the Need to Know
Critical Control 16: Account Monitoring and Control
Critical Control 17: Data Loss Prevention
Critical Control 18: Incident Response Capability
Critical Control 19: Secure Network Engineering
Critical Control 20: Penetration Tests and Red Team Exercises

Because this work is being analyzed directly by Working Group 11 to address the generic risk to ISPs of various hacking and unauthorized access issues, Working Group 4 will not be commenting in-depth in this area, and refers readers to reports from Working Group 11 for comprehensive, and updated coverage of these risks when they issue their report. We will comment upon current BCPs for ISPs to look to adopt in the interim, and provide further background around risks unique to running BGP servers/routers in this area.

An ISP's routing infrastructure is an important asset to protect, as gaining control of it can lead to a wide variety of harms to ISP customers. Further, an ISP's staff computers, servers, and networking infrastructure also rely upon their own routers to correctly direct traffic to its intended destinations. The ISP's own sensitive data and processes could be compromised via hacked routers/servers. Thus routers should be included on the list of network assets that are assigned the highest level of priority for protection under any type of ISP security program.

There are many industry standard publications pertaining to overall cybersecurity best practices available for adoption by ISPs or any organization at risk of attack, including prior CSRIC reports. It is incumbent upon ISPs to maintain their overall security posture and be up-to-date on the latest industry BCPs and adopt the practices applicable to their organization. Of particular note is the IETF's RFC 4778 - *Current Operational Security Practices in Internet Service Provider Environments*[6] which offers a comprehensive survey of ISP security practices. An older IETF publication, but still active BCP, that still applies to ISP environments can be found with BCP 46, aka RFC 3013 *Recommended Internet Service Provider Security Services and Procedures*[7]. NIST also puts out highly applicable advice and BCPs for running

[6] http://www.ietf.org/rfc/rfc4778.txt
[7] http://www.apps.ietf.org/rfc/rfc3013.txt

government networks, with the most currently relevant special report, NIST SP-800-53.

The ultimate goal of someone attempting unauthorized access to routing infrastructure would be to either deny customer use of those servers or, more likely, insert false entries within the router to misdirect the users of those routers. This is a functional equivalent to route injection and propagation attacks as already described in section 5.2. So the analysis and recommendations presented in section 5.2.1.5 with respect to monitoring for and reacting to route injection and propagation attacks apply in the scenario where an attacker has breached a router to add incorrect entries.

5.3.1.1 Recommendations

1) ISPs should refer to and implement the practices found in CSRIC 2A – Cyber Security Best Practices that apply to securing servers and ensure that routing infrastructure is protected.
2) ISPs should adopt applicable BCPs found in other relevant network security industry approved/adopted publications. Monitor for applicable documents and update. Three documents were identified that currently apply to protecting ISP networks: IETF RFC 4778 and BCP 46 (RFC 3013); NIST special publications series: NIST SP-800-53
3) ISPs should ensure that methods exist within the ISP's operations to respond to detected or reported successful route injection and propagation attacks, so that such entries can be rapidly remediated.
4) ISPs should consider implementing routing-specific monitoring regimes to assess the integrity of data being reported by the ISP's routers that meet the particular operational and infrastructure environments of the ISP.

5.3.2 ISP insiders inserting false entries into routers

While insider threats can be considered a subset of the more general security threat of unauthorized access and hacking, they deserve special attention in the realm of routing security. ISP insiders have unparalleled access to any systems run by an ISP, and in the case of routers, the ability to modify entries is both trivially easy and potentially difficult to detect. Since routers don't typically have company-sensitive information, are accessed by thousands of machines continuously, and are not usually hardened or monitored like other critical servers, it is relatively easy for an insider to alter a router's configuration in a way that adversely affects routing.

The analysis and recommendations for this particular threat do not differ significantly from those presented in Section 5.3.1 of this report - Hacking and unauthorized 3rd party access to routing infrastructure. However, it is worth paying special attention to this particular exposure given the liabilities an ISP may be exposed to from such difficult-to-detect activities of its own employees.

5.3.2.1 Recommendations

1) Refer to section 5.3.1.1 for generic hacking threats.

5.3.3 Denial-of-Service Attacks against ISP Infrastructure

Denial-of-Service (DoS) and Distributed Denial-of-Service (DDoS) are some of the oldest and most prolific attacks that ISPs have faced over the years and continue to defend against today.

Typically, an external actor who is targeting some Internet presence or infrastructure to make it unusable is behind such attacks. However, DoS/DDoS attacks come in many flavors that can be broadly lumped into two primary categories: logic attacks and resource exhaustion/flooding attacks.[8] Logic attacks exploit vulnerabilities to cause a server or service to crash or reduce performance below usable thresholds. Resource exhaustion or flooding attacks cause server or network resources to be consumed to the point where the targeted service no longer responds or service is reduced to the point it is operationally unacceptable. We will examine the latter type of attack in this section of analysis, as resource exhaustion . Logic attacks are largely directed to break services/servers and can be largely addressed with the analysis and recommendations described above with respect to BGP specific issues and also put forward in section 5.3.1 that cover protecting networked assets from various hacking and other attacks.

There is a large variety of flooding attacks that an ISP could face in daily operations. These can be targeted at networks or any server, machine, router, or even user of an ISP's network. From the perspective of routing operations, it is helpful to differentiate between "generic" DoS attacks that could affect any server, and those that exploit some characteristic of BGP that can be utilized to affect routers in particular, which have already been covered.

Due to the long history, huge potential impact, and widespread use of various DoS and DDoS attacks, there is an abundance of materials, services, techniques and BCPs available for dealing with these attacks. ISPs will likely have some practices in place for dealing with attacks both originating from their networks and that are being directed at their networks and impacting their services. The IETF's RFC 4732 *Internet Denial-of-Service Considerations*[9] provides an ISP with a thorough overview of DoS/DDoS attacks and mitigation strategies and provides a solid foundational document. The SANS Institute has published a useful document for ISPs that is another reference document of BCPs against DoS/DDoS attacks entitled *A Summary of DoS/DDoS Prevention, Monitoring and Mitigation Techniques in a Service Provider Environment*[10].

As mentioned in section 5.3.1, there are several documents that cover general ISP security concerns, and those typically include prescriptive advice for protecting a network against DoS/DDoS attacks. Such advice can be found in previously cited documents including prior CSRIC reports: CSRIC 2A – Cyber Security Best Practices[11], the IETF's RFC 4778 - *Current Operational Security Practices in Internet Service Provider Environments*[12], BCP 46, RFC 3013 *Recommended Internet Service Provider Security Services and Procedures*[13] and NIST's special report, NIST SP-800-53.

For the most part, an ISP's routers for interdomain routing must be publicly available in order for the networks they serve to be reachable across the Internet. Thus measures to restrict access

[8] http://static.usenix.org/publications/library/proceedings/sec01/moore/moore.pdf
[9] http://tools.ietf.org/rfc/rfc4732.txt
[10] http://www.sans.org/reading_room/whitepapers/intrusion/summary-dos-ddos-prevention-monitoring-mitigation-techniques-service-provider-enviro_1212
[11] http://www.fcc.gov/pshs/docs/csric/WG2A-Cyber-Security-Best-Practices-Final-Report.pdf
[12] http://www.ietf.org/rfc/rfc4778.txt
[13] http://www.apps.ietf.org/rfc/rfc3013.txt

that can be implemented for an ISP's internal infrastructure are unavailable as options for these connecting routers. This leaves an ISP with limited choices for DDoS protection, including the traditional approaches of overprovisioning of equipment and bandwidth and various DoS/DDoS protection services and techniques.

5.3.3.1 Recommendations

1) ISPs should implement BCPs and recommendations for securing an ISP's infrastructure against DoS/DDoS attacks that are enumerated in the IETF's RFC 4732 *Internet Denial-of-Service Considerations* and consider implementing BCPs enumerated in the SANS Institute reference document of BCPs against DoS/DDoS attacks entitled *A Summary of DoS/DDoS Prevention, Monitoring and Mitigation Techniques in a Service Provider Environment*.
2) ISPs should refer to and implement the BCPs related to DoS/DDoS protection found in CSRIC 2A – Cyber Security Best Practices that apply to protecting servers from DoS/DDoS attacks.
3) ISPs should consider adopting BCPs found in other relevant network security industry approved/adopted publications that pertain to DoS/DDoS issues, and monitor for applicable documents and updates. Four that currently apply to protecting ISP networks from DoS/DDoS threats are IETF RFC 4778 and BCP 46 (RFC 3013); NIST special publications series: NIST SP-800-53; and ISOC Publication *Towards Improving DNS Security, Stability, and Resiliency*.
4) ISPs should review and apply BCPs for protecting network assets against DoS/DDoS attacks carefully to ensure they are appropriate to protect routing infrastructure.

5.3.4 Attacks against administrative controls of routing identifiers

Blocks of IP space and Autonomous Systems Numbers (ASNs) are allocated by various registries around the world. Each of these Regional Internet Registries (RIR's) is provided IP space and ASN allocation blocks by IANA, to manage under their own rules and practices. In-turn, several of these registries allow for country or other region/use specific registries to sub-allocate IP space based on their own rules, processes and systems. Each RIR maintains a centralized "whois" database that designates the "owner" of IP spaces or ASN's within their remit. Access to the databases that control these designations, and thus "rights" to use a particular space or ASN is provided and managed by the RIR's and sub RIR registries depending upon the region. Processes for authentication and management of these identifier resources are not standardized, and until recently, were relatively unsophisticated and insecure. This presents an administrative attack vector allowing a miscreant to use a variety of account attack methods, from hacking to password guessing to social engineering and more that could allow them to assume control over an ASN or IP space allocation. In other fields, such attacks would be considered "hijacking" or "account take-over" attacks, but the use of the word "hijacking" in the BGP space to include various injection and origin announcements complicates the common taxonomy. Thus for this section, we will refer to account "hijacking" as "account take-over".

The primary concern for most ASN and prefix block owners and the ISPs that service them in such scenarios is the take-over of active space they are using. A miscreant could literally "take

over" IP space being routed and used by the victim, much like an origin attack as described in section 5.3. In this case, it would be equivalent to a full take-over, with the majority of the global routing system recognizing the miscreant's announcement as being the new "legitimate" one, with all the inherent risks previously described. The real owner will have to prove their legitimacy and actual legal ownership/control of the resource that has been taken over. Depending upon the authentication scheme the registry uses, this can prove difficult, especially for legacy space and older ASN registrations.

A corollary of this attack scenario, is a miscreant taking over "dormant" IP space or an unused ASN, and thus "squatting" in unused territory[14]. While not impactful on existing Internet presence, squatting on IP space can lead to many forms of abuse, including the announcement of bogus peering arrangements, if the squatted resource is an ASN.

In a take-over scenario, a miscreant typically impersonates or compromises the registrant of the ASN and/or IP space in order to gain access to the management account for that ASN or CIDR block. Until recently, nearly all RIRs and registries used an e-mail authentication scheme to manage registrant change requests. Thus, if the registrant's e-mail address uses an available domain name, the miscreant can register the domain name, recreate the administration email address, and authenticates himself with the registry. If the domain isn't available, the criminal could still try to hijack the domain name registration account to gain control of that same domain. If a registry or RIR requires more verification for registrant account management, the criminal use various social engineering tricks against the registry staff to get into the management account.

Once a criminal has control of the registration account, they can update the information there to allow them to move to a new peering ISP, create new announcements from their "new" space, or launch any sort of BGP-type attack as listed above. Even more basically, the criminal can simply utilize their new control of the ASN/prefix to have their own abusive infrastructure announced on the Internet for whatever process they would like. This includes direct abuse against the Internet in general (e.g. hosting malware controllers, phishing, on-line scams, etc.), but also the ability to impersonate the original holder of the space they have taken over. Of course they can also intercept traffic originally destined for the legitimate holder of the space as previously described for various route-hijacking scenarios.

The end result of an administrative account take-over is likely to be similar to other injection attacks against routing infrastructure as covered in section 5.2. Thus, ISPs will want to consult BCPs covering techniques for monitoring and reacting to those types of attacks. These BCPs cover the general effects of a BGP origin attacks – dealing with service interruptions and the worldwide impacts. ISPs typically do not have direct access or control to RIR or other registry account information that has been compromised in most hijacking attacks. The ISP is dependent upon the affected registry to restore control of the ISP's management account, or in the case of a serious breach, the registrar/registry's own services. Once control is re-established, the original, correct information needs to be re-entered and published again. This will usually mean updating

[14] For a full description of the taxonomy of hijacking, squatting, and spoofing attacks in routing space, see *Internet Address Hijacking, Spoofing and Squatting Attacks - http://securityskeptic.typepad.com/the-security-skeptic/2011/06/internet-address-hijacking-spoofing-and-squatting-attacks.html*

routing table entries/BGP announcements, and fixing any account information that has been modified.

The industry has largely been slow to adopt security measures to protect account access for controlling ASN and CIDR block management that are found in other online services like financial services, e-commerce, or even some ISP management systems. The industry also has many participants with a wide variety of geographical regions, with few standards and requirements for the security of registration systems, and very limited oversight. This means it is often difficult to find support for typical online security tools like multi-factor authentication, multi-channel authentication, and verification of high-value transactions. This has been changing in recent years at the RIR level, with ARIN, RIPE, and several other IP registries at various levels of authority implementing new controls and auditing account information. Despite this, gaps exist, especially with "legacy" data entered many years ago before current management systems and authentication processes were implemented.

While there is scant guidance on this topic area for ASN/IP block management, ICANN's Security and Stability Advisory Committee (SSAC) has released two documents to address these issues in the domain name space which is quite analogous to the provisioning of IP space. These documents provide BCPs for avoiding and mitigating many of these issues. SAC 40 *Measures to Protect Domain Registration Services Against Exploitation or Misuse*[15], addresses issues faced by domain name registrars and offers numerous BCPs and recommendations for securing a registrar against the techniques being used by domain name hijackers. Many of the BCP's presented there would be applicable to RIR's and other IP address provisioning authorities, including ISPs managing their own customers. SSAC 44, *A Registrant's Guide to Protecting Domain Name Registration Accounts*[16], provides advice to domain name registrants to put in place to better protect their domains from hijacking. Similar techniques could be used by operators to protect their own IP space allocations. Given the limited choices and practices followed by various IP space allocators, ISPs need to carefully evaluate their security posture and the practices of their RIR's or other IP space allocators with these BCPs in mind.

5.3.4.1 Recommendations

1) ISPs and their customers should refer to the BCPs and recommendations found in SSAC 44 *A Registrant's Guide to Protecting Domain Name Registration Accounts* with respect to managing their ASN's and IP spaces they register and use to provide services.
2) ISPs should review the BCPs and recommendations found in SAC 40 *Measures to Protect Domain Registration Services Against Exploitation or Misuse* to provide similar protections for IP space they allocate to their own customers.

6 Conclusions

Working Group 4 has recommended the adoption of numerous best practices for protecting the inter-domain BGP routing system. As a distributed infrastructure requiring several actors to both enable and protect it, network operators face challenges outside of their direct control in tackling many of the issues identified. The more widely Best Current Practices are utilized, the more robust the whole system will be to both bad actors and simple mistakes. See Appendix 7.3 for a

[15] http://www.icann.org/en/groups/ssac/documents/sac-040-en.pdf
[16] http://www.icann.org/en/committees/security/sac044.pdf

tabular display of risks indexed with the appropriate countermeasures as discussed in the body text of the document.

7 Appendix

7.1 Background

Note that in order to remain consistent with other CSRIC III reports, considerable portions of "Salient Features of BGP Operation" have been taken verbatim from the Appendix section of the CSRIC III, Working Group 6 Interim Report published March 8, 2012. Other parts have been taken from the NIST (National Institute of Standards and Technologies) report entitled Border Gateway Protocol Security.

7.1.1 Salient Features of BGP Operation

This section is intended for non-experts who have a need to understand the origins of BGP security problems.

Although unknown to most users, the Border Gateway Protocol (BGP) is critical to keeping the Internet running. BGP is a routing protocol, which means that it is used to update routing information between major systems. BGP is in fact the primary inter-domain routing protocol, and has been in use since the commercialization of the Internet. Because systems connected to the Internet change constantly, the most efficient paths between systems must be updated on a regular basis. Otherwise, communications would quickly slow down or stop. Without BGP, email, Web page transmissions, and other Internet communications would not reach their intended destinations. Securing BGP against attacks by intruders is thus critical to keeping the Internet running smoothly.

Many organizations do not need to operate BGP routers because they use Internet service providers (ISP) that take care of these management functions. But larger organizations with large networks have routers that run BGP and other routing protocols. The collection of routers, computers, and other components within a single administrative domain is known as an autonomous system (AS). An ISP typically represents a single AS. In some cases, corporate networks tied to the ISP may also be part of the ISP's AS, even though some aspects of their administration are not under the control of the ISP.

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.

7.1.2 Review of Router Operations

In a small local area network (LAN), data packets are sent across the wire, typically using Ethernet hardware, and all hosts on the network see the transmitted packets. Packets addressed

to a host are received and processed, while all others are ignored. Once networks grow beyond a few hosts, though, communication must occur in a more organized manner. Routers perform the task of communicating packets among individual LANs or larger networks of hosts.

To make internetworking possible, routers must accomplish these primary functions:

- Parsing address information in received packets
- Forwarding packets to other parts of the network, sometimes filtering out packets that should not be forwarded
- Maintaining tables of address information for routing packets.

BGP is used in updating routing tables, which are essential in assuring the correct operation of networks. BGP is a dynamic routing scheme—it updates routing information based on packets that are continually exchanged between BGP routers on the Internet. Routing information received from other BGP routers (often called "BGP speakers") is accumulated in a *routing table*. The routing process uses this routing information, plus local policy rules, to determine routes to various network destinations. These routes are then installed in the router's *forwarding table*. The forwarding table is actually used in determining how to forward packets, although the term routing table is often used to describe this function (particularly in documentation for home networking routers).

7.2 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 de-aggregation attacks on specific IP addresses. They may be hard to identify unless anomaly detection is unambiguous, or the victim is important enough to create a large commotion. Examples: Pakistan Hijacks YouTube[17] (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"**[18],** or a hijacking of a very large portion of the Internet for several hours by TTNet in 2004[19], or a 2006 ConEd incident[20]. 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[21], AfNOG Takes Byte Out of Internet[22], Longer is Not Always Better[23] 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

[17] http://www.renesys.com/blog/2008/02/pakistan-hijacks-youtube-1.shtml
[18] http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml
[19] Alin C. Popescu, Brian J. Premore, and Todd Underwood, Anatomy of a Leak: AS9121. NANOG 34, May 16, 2005.
[20] http://www.renesys.com/blog/2006/01/coned-steals-the-net.shtml
[21] http://www.renesys.com/blog/2010/08/house-of-cards.shtml
[22] http://www.renesys.com/blog/2009/05/byte-me.shtml
[23] http://www.renesys.com/blog/2009/02/longer-is-not-better.shtml

2008[24]. Two other similar incidents were found in a 7-month period surrounding the DEFCON demo by mining of a BGP update repository conducted in 2009[25] 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 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 de-aggregation[26] 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.

[24] A. Pilosov and T. Kapela, Stealing the internet, DEFCON 16 August 10, 2008
[25] C. Hepner and E. Zmijewski, Defending against BGP Man-in-the-Middle attacks, Black Hat DC February 2009
[26] J. Cowie, The Curious Incident of 7 November 2011, NANOG 54, February 7, 2012

7.3 BGP Risks Matrix

BGP Routing Security Risks Examined by WG 4			
Network Operator Role	**Risks**	**Report Sect.**	**Recommendations**
Stub network (e.g. Enterprise, Data Center)	Session-level threats	5.1.1	• Consider MD5 or GTSM if neighbor recommends it
	DoS (routers and routing info)	5.1.2	• Control-Plane Policing (rate-limiting) • Keep up-to-date router software
	Spoofed Source IP Addresses	5.1.3	• Use uRPF (unicast Reverse Path Forwarding) in strict mode or other similar features at access edge of network (e.g. data-center or campus). • Filter source IP address on packets at network edge to ISPs
	Incorrect route injection and propagation	5.2.1	• Keep current information in "whois" and IRR (Internet Routing Registry) databases • Outbound prefix filtering • Use monitoring services to check for incorrect routing announcements and/or propagation
	Other Attacks (e.g., hacking, insider, social engineering)	5.3	• Consider many recommendations about operational security processes
Internet Service Provider Network	Session-level threats	5.1.1	• Consider a plan to use MD5 or GTSM including flexibility to adjust to different deployment scenario specifics
	DoS (routers and routing info)	5.1.2	• Control-Plane Policing (rate-limiting)

			• Keep up-to-date router software
	Spoofed Source IP Addresses	5.1.3	• Use uRPF (unicast Reverse Path Forwarding) in strict or loose mode as appropriate (e.g. strict mode at network ingress such as data-center or subscriber edge, loose mode at inter-provider border)
	Incorrect route injection and propagation	5.2.1	• Keep current information in "whois" and IRR (Internet Routing Registry) databases • Consult current information in "whois" and IRR (Internet Routing Registry) databases when provisioning or updating customer routing • Implement inbound prefix filtering from customers • Consider AS-path filters and maximum-prefix limits as second line of defense • Use monitoring services to check for incorrect routing announcements and/or propagation
	Other Attacks (e.g., hacking, insider, social engineering)	5.3	• Consider many recommendations about operational security processes

7.4 BGP BCP Document References

Network Protection Documents

NIST Special Publication 800-54 *Border Gateway Protocol (BGP) Security Recommendations*

WG2A - Cyber Security Best Practices

SANS: *Twenty Critical Security Controls for Effective Cyber Defense: Consensus Audit Guidelines (CAG)*

NIST Special Publication 800-53 *Recommended Security Controls for Federal Information Systems and Organizations*

IETF RFC 4778 - *Current Operational Security Practices in Internet Service Provider Environments*

IETF RFC 3013 *Recommended Internet Service Provider Security Services and Procedures*

Source Address verification/filtering

IETF BCP38/RFC 2827 *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*

BCP 84/RFC 3704 *Ingress Filtering for Multihomed Networks*

BCP 140/RFC 5358 *Preventing Use of Recursive Nameservers in Reflector Attacks*

ICANN SAC004 *Securing the Edge*

DoS/DDoS Considerations

IETF RFC 4732 *Internet Denial-of-Service Considerations*

SANS *A Summary of DoS/DDoS Prevention, Monitoring and Mitigation Techniques in a Service Provider Environment*