



September, 2014

WORKING GROUP 6
Long-Term Core Internet Protocol Improvements
Final Report

Table of Contents

1	Results in Brief	3
2	Introduction.....	4
2.1	CSRIC Structure.....	4
2.2	Working Group 6 Team Members.....	4
3	Open Recursive DNS Server Subgroup.....	5
3.1	Objective, Scope, and Methodology.....	5
3.2	Background.....	6
3.3	Analysis, Findings and Recommendations.....	10
4	BGP Subgroup	13
4.1	Objective, Scope and Methodology.....	13
4.2	Background.....	14
4.3	Analysis and Findings	18
5	Recommendations.....	22
5.1	Working Group 6 DNS & Routing Recommendation(s)	22
5.2	DNS Subgroup recommendations	22
5.3	Routing Subgroup recommendations	24
6	Conclusions.....	25
	Appendix A: Getting Started with RPKI	26

1 Results in Brief

Working Group 6 was chartered to look at long-term improvements to the core Internet protocols – Domain Name System (DNS) and Border Gateway Protocol (BGP). Earlier CSRIC working groups had looked extensively at DNS and BGP:

- CSRIC III
 - Working Group 4 – BGP Best Practices¹
 - Working Group 4 – DNS Best Practices²
 - Working Group 5 – Measurements of DNSSEC Deployments³
 - Working Group 6 – Secure BGP Deployment⁴
- CSRIC II
 - Working Group 2a – Cyber Security Best Practices⁵
 - Working Group 8 – Internet Service Provider Network Protection Practices⁶

Given the thorough review of DNS and BGP by the earlier CSRIC working groups, Working Group 6 believed it was best to focus on two topic areas of DNS and BGP respectively, that deserved a focused study, rather than doing a broad survey and refresh of the best practices that had been put forth by the 6 previous CSRIC working groups.

For DNS, the group focused on Open Recursive DNS Servers as they have been and continue to be a prime target for malicious actors to use for attacks as DNS uses the User Datagram Protocol (UDP), and anything UDP-based is trivially susceptible to source IP address spoofing. As part of doing the analysis, it became evidently clear that there was no quick-fix for dealing with the millions of open recursive DNS servers^{7,8} and this is further hampered by the fact that the majority these are home routers⁹ that are outside the control of the network operators.

In the realm of routing security, the working group sought to provide through its collective expertise illumination on some of the less well-understood areas. A taxonomy was laid out to

¹ “WORKING GROUP 4 Network Security Best Practices FINAL Report – BGP Security Best Practices.” March 1, 2013. Accessed August 19, 2014. http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG4_Report_March_2013.pdf.

² “WORKING GROUP 4 Network Security Best Practices FINAL Report – DNS Best Practices.” March 1, 2013. Accessed August 19, 2014. http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRICIII_9-12-12_WG4-FINAL-Report-DNS-Best-Practices.pdf.

³ “Working Group 5 DNSSEC Implementation Practices for ISPs Final Report on Measurement of DNSSEC Deployment.” February 22, 2013. Accessed August 19, 2014. http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG5_Report_March_2013.pdf.

⁴ “WORKING GROUP 6 Secure BGP Deployment Final Report.” March 1, 2013. Accessed August 19, 2014. http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG6_Report_March_2013.pdf.

⁵ “Working Group 2A Cyber Security Best Practices Page 1 of 24 Final Report.” March 1, 2011. Accessed August 19, 2014. <http://transition.fcc.gov/pshs/docs/csric/WG2A-Cyber-Security-Best-Practices-Final-Report.pdf>.

⁶ “FINAL REPORT Internet Service Provider (ISP) Network Protection Practices.” December 1, 2010. Accessed August 19, 2014. http://transition.fcc.gov/pshs/docs/csric/CSRIC_WG8_FINAL_REPORT_ISP_NETWORK_PROTECTION_2010213.pdf.

⁷ “Open Resolver Project.” Open Resolver Project. Accessed August 26, 2014. <http://openresolverproject.org>.

⁸ “The Shadowserver Foundation: DNS Scanning Project.” The Shadowserver Foundation: DNS Scanning Project. Accessed August 26, 2014. <https://dnsscan.shadowserver.org>.

⁹ “24 Million Home Routers Expose ISPs to Massive DNS-based DDoS Attacks - Nominum.” Nominum. April 2, 2014. Accessed August 11, 2014. <http://nominum.com/news-post/24m-home-routers-expose-ddos/>.

give precision to the discussion of routing security. A discussion of recent events reported in the news or trade press put these incidents in context of the history and understanding of Internet anomalies. By reviewing a number of existing and past measurement projects, the working group demonstrated the broad range of these projects, their goals, and their fitness for operational use or lack thereof. Lastly, an appendix seeks to describe in narrative fashion how to interact with the RPKI infrastructure, which may add robustness to the routing infrastructure and information flow.

The working group did an extensive analysis of the problem spaces for open recursive DNS servers and inter-domain routing to identify recommendations to put forth. As the working group studied the problem space, it became evident that the problems associated with open recursive DNS servers and inter-domain routing are not the sole responsibility of any one industry sector or segment, but rather a shared responsibility of the broader Internet community. As such the recommendations put forth by this group reflect that shared responsibility. The group identified ten (10) recommendations that are listed in section 5 of this report.

2 Introduction

2.1 CSRIC Structure

Communications Security, Reliability, and Interoperability Council (CSRIC) IV									
CSRIC Steering Committee									
Chair or Co-Chairs: Working Group 1	Chair or Co-Chairs: Working Group 2	Chair or Co-Chairs: Working Group 3	Chair or Co-Chairs: Working Group 4	Chair or Co-Chairs: Working Group 5	Chair or Co-Chairs: Working Group 6	Chair or Co-Chairs: Working Group 7	Chair or Co-Chairs: Working Group 8	Chair or Co-Chairs: Working Group 9	Chair or Co-Chairs: Working Group 10
Working Group 1: Next Generation 911	Working Group 2: Wireless Emergency Alerts	Working Group 3: EAS	Working Group 4: Cybersecurity Best Practices Working	Working Group 5: Server-Based DDoS Attacks	Working Group 6: Long-Term Core Internet Protocol Improvements	Working Group 7: Legacy Best Practice Updates	Working Group 8: Submarine Cable Landing Sites	Working Group 9: Infrastructure Sharing During Emergencies	Working Group 10: CPE Powering

Table 1 - Working Group Structure

2.2 Working Group 6 Team Members

Working Group 6 consists of the members listed below.

Name	Company
William Check (Chair)	National Cable and Telecommunications Association (NCTA)
Tony Tauber (Subgroup chair)	Comcast
Matt Tooley (Subgroup chair)	National Cable and Telecommunications Association (NCTA)
Daniel Awduche	Verizon
Luke Berndt	Department of Homeland Security
D. Blumenthal	PIR.org
Matthew Bretan	
Roy Cormeir	Nsight.com
Rob Fleishman	Xercole
Chris Garner	CenturyLink
Mike Geller	Cisco
Sharon Goldberg	Boston University
Joseph Lorenzo Hall	Center for Democracy & Technology
Kurin Jacob (FCC Liaison)	Federal Communications Commission
Sun Joseph	CableLabs

Merike Kaeo	Internet Identity
Michael Kelsen	Time-Warner Cable
Mazen Khaddam	Cox Communications
Padm Krishnaswamy	Federal Communications Commission
Warren Kumari	Google
David LaBianca	
Eric Lent	Comcast
Jason Livingood	Comcast
John Marnho	CTIA – The Wireless Association
Ron Mathis	Intrado
Christopher Mikkelson	CenturyLink
Doug Montgomery	National Institute for Standards and Technology
Chris Morrow	Google
Andy Agielski	Renesys
Eric Osterweil	Verisign
Richard Perlotto	Shadow Server
Jennifer Rexford	Princeton University
Brian Rexroad	AT&T
Chris Richardson	Internet Identity
Ronald Ritchey	Bank of America
Chris Roosenraad	Time-Warner Cable
Joe St. Sauver	University of Oregon
Ted Seely	Sprint
Tom Soroka	US Telecom Association
Todd Szymanski	Sprint
Sounil Yu	Bank of America
Eric Ziegast	Farsight Security

Table 2 - List of Working Group Members

Working group 6 was divided into two sub-groups – Open Recursive DNS Servers and Routing.

3 Open Recursive DNS Server Subgroup

3.1 Objective, Scope, and Methodology

3.1.1 Objective & Scope

The DNS has been leveraged as an attack platform because it primarily uses the UDP, and anything UDP-based is trivially susceptible to source IP address spoofing (for example, TCP traffic spoofing requires both modifying the source IP address and some header fields as session number, which requires ongoing network monitoring). Recursive DNS servers will answer all end-user queries unless they are otherwise configured with restrictions. As a result of misconfiguration or non-restriction, unmanaged (i.e. unmanaged here means a server that is either poorly managed or was not intended to be operated as a recursive DNS server), recursive DNS servers may answer DNS queries from hosts outside of their administrative domain or otherwise off of the networks on which they reside.

Thus, unless properly configured, a recursive DNS server in a U.S. network may answer spoofed queries, potentially sent from overseas networks and vice versa. DNS servers configured and responding in this manner are commonly referred to as “Open Resolvers” or “Open DNS Resolvers”. Open recursive DNS servers have been and continue to be a prime target for malicious actors to use as part of reflection/amplification attacks as described in section 3.2.3 via spoofed source IP addresses. These attacks pose a significant threat to the U.S. and the global Internet infrastructure.

This DNS sub-team of Working Group 6, Long Term Protocol Improvements, is chartered to examine and make recommendations to the Council regarding best practices to address the issue of open recursive DNS servers. The issue of open recursive DNS servers is not isolated to one subset of the Internet ecosystem. The problem cuts across a large portion of the Internet ecosystem and requires the cooperative, collaborative efforts of all members of the ecosystem to mitigate the problem posed by open recursive DNS servers. The sub-team will examine and recommend best practices to address this problem.

It should be noted that some other protocols have come into use as part of reflection/amplification attacks since the problem with open recursive DNS servers was identified¹⁰. Network Time Protocol (NTP) and Character Generator (chargen) have shown to be similar or greater attack vectors. Underlying many of these attacks is packet-level source address forgery or spoofing, a well-known vulnerability where an attacker generates and transmits UDP packets purporting to be from the victim's IP address.

3.1.2 Methodology

The DNS subgroup of Working Group 6 began by doing a review of the problem and working to identify well-known best practices for the problem. The DNS subgroup conducted bi-weekly conference calls to review the identified well-known best practices as part of a gap-analysis of the solution space. The group then worked to identify new best practices for addressing the problem with open recursive DNS servers. This was followed by reviewing and assessing the various Internet community measurement projects to determine what and if any recommendations could be made to the FCC with regards to how to best leverage these efforts.

3.2 Background

3.2.1 Brief Overview of DNS

The DNS was created to provide a scalable, flexible, dynamic, robust and resilient network service to help people, the applications they use, and their computing devices reference services and computers using names instead of IP addresses. Users or IT staff would either manually configure a name server IP address for their computer and applications to use to lookup names, or their computers would be dynamically configured to use one or more name server IP addresses provided by their Internet service provider (ISP).

The caching resolving name server (“resolver”) would receive client requests from the user's computer and recursively make DNS requests in the hierarchical structure of authoritative name servers related to the domain name. Upon receiving an answer or giving up, the result would be returned to the client.

¹⁰ Rossow, Christian. “Amplification Hell: Revisiting Network Protocols for DDoS Abuse.” Internet Society. February 23, 2014. Accessed August 11, 2014. http://www.internetsociety.org/sites/default/files/01_5.pdf.

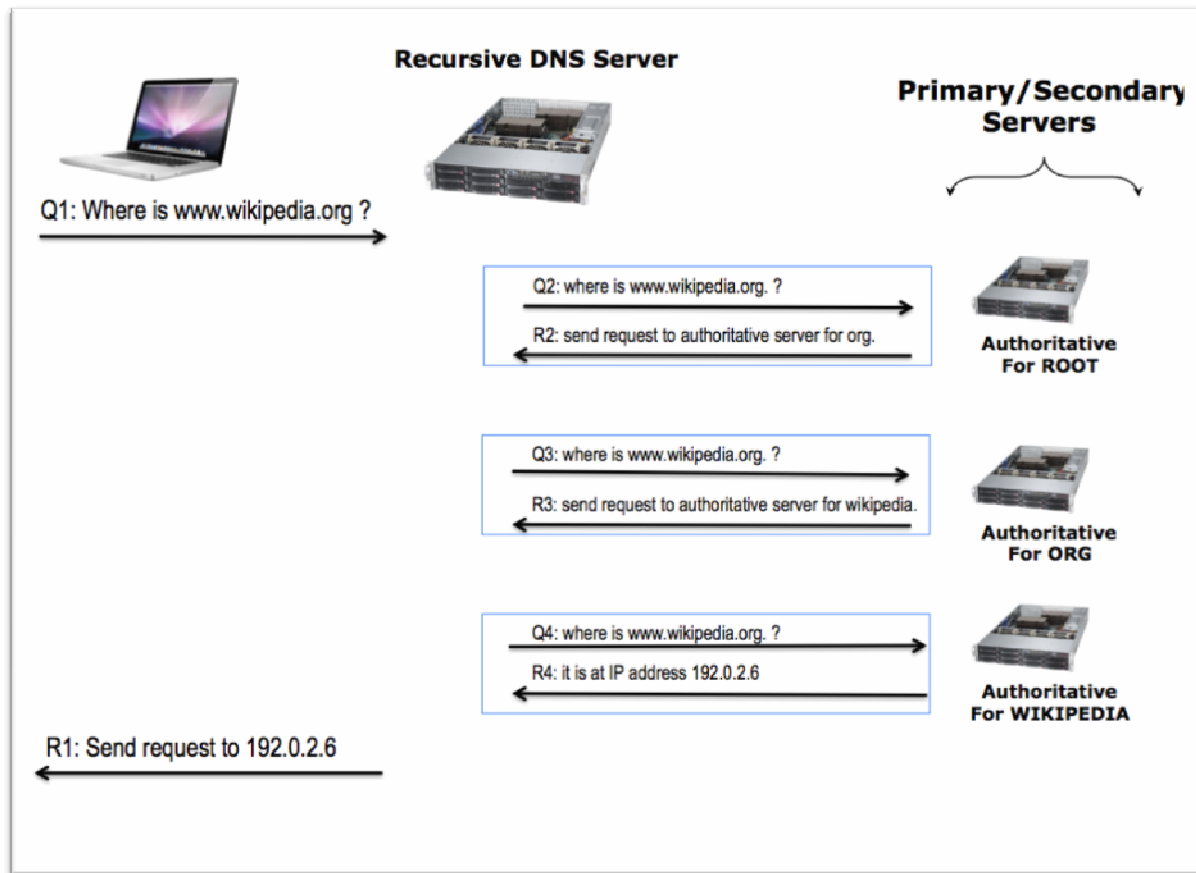


Figure 1 Example of a Hierarchical DNS Query

To increase efficiency of the overall DNS system, the information would be cached in memory for a specified period of time so that if the same client or another client using the same resolver needed to look up the information again, the server would not have to retrieve the same information again from authoritative servers.

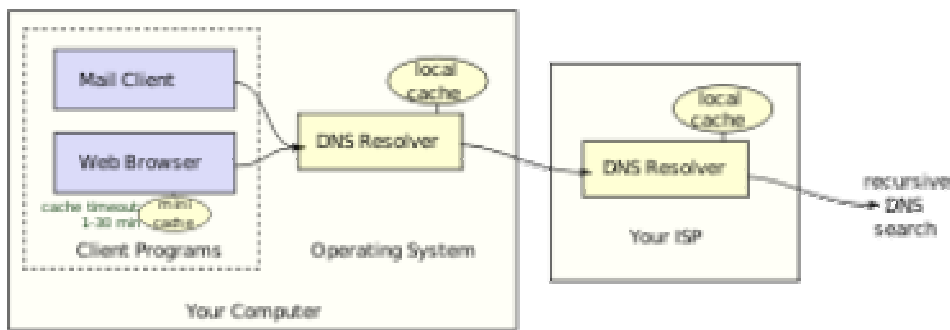


Figure 2 Example of DNS resolution with a local cache

The client-server protocol used a lightweight UDP query and response for most requests. If a UDP response was not received, clients and resolvers would try a TCP connection to receive the data. For efficiency purposes, most network administrators architected their use of DNS to use UDP as much as possible, so that both query and response data could fit within a 512-byte UDP packet. Later enhancements via EDNS0 (Extension Mechanism for DNS) allow DNS responses to extend significantly beyond 512 bytes through packet fragmentation.

3.2.2 Actors and Roles in the DNS System

A potentially large population of users could efficiently be served by a single recursive caching resolver. At the time that the DNS protocols were developed and deployed in the 1980's and 1990's, the most popular reference implementation software was the Berkeley Internet Name Domain (BIND). The software provided both the client to caching resolver service for clients, and the authoritative name service for domain names.

While the Internet had been mostly used for collaborative commercial, government and academic research, a new class of consumer and office computers started directly accessing the Internet with the release of Windows 95. Due to this increasing use of the Internet, the inclusion of resolvers in operating systems became widespread. Microsoft also developed and supported their own name server software, as an alternative to BIND to run on Windows NT 4.0 in 1996.

Absent the network threats seen today, network administrators typically made resolvers openly available on the Internet for their customers and Internet users in general. As DHCP and home gateway routers became more widely used from the mid-1990's to the early 2000's, more users automatically configured their computers using resolver IP addresses provided by their ISPs. The distribution of available resolvers and authoritative servers across the Internet seemed beneficial.

3.2.3 Security Issues with Open Recursive DNS Servers

Security concepts related to DNS evolved in response to the attempts to leverage the name server software as a method of gaining unauthorized access to server and network infrastructure. The concept of utilizing firewalls and restricting access to name service evolved in response to attacks, particularly in commercial or enterprise networks attached to the Internet.

Attempts by individuals or organizations to alter the information served by resolvers to their clients ("cache poisoning")¹¹ brought improvements and hardening to the BIND software. Other name server software was independently developed that separated resolver functionality from authoritative server functionality. Software became specialized for serving the role of either a resolver or an authoritative server.

As denial of service attacks evolved in the 1990's, a general vulnerability in UDP and Internet Control Message Protocol (ICMP), and the loose trust relationships in network hardware and routing architecture, allowed an attacker to forge source addresses. While a solution was proposed¹² by the IETF in 1998, it was not widely adopted during the rapid growth of Internet infrastructure. Attacks were limited to trivial Internet protocols like "echo" or ICMP "ping". Attacks that could flood network connections were first seen, and spoofed sources started to allow attackers to not be directly attributed. With reflected attacks, a normally benign server could now be used by an attacker to direct traffic to a victim. One of the first examples of such an attack, which also employed amplification, was the Smurf Attack.¹³

¹¹ "US-CERT CA-1997-22." BIND - the Berkeley Internet Name Daemon. May 26, 1996. Accessed August 11, 2014. <http://www.cert.org/historical/advisories/CA-1997-22.cfm>.

¹² Ferguson, P., and D. Senie. "RFC 2267 - Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing." RFC 2267 - Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing. January 1, 1998. Accessed August 18, 2014.

¹³ "Smurf Attack." Wikipedia. July 23, 2014. Accessed August 11, 2014.

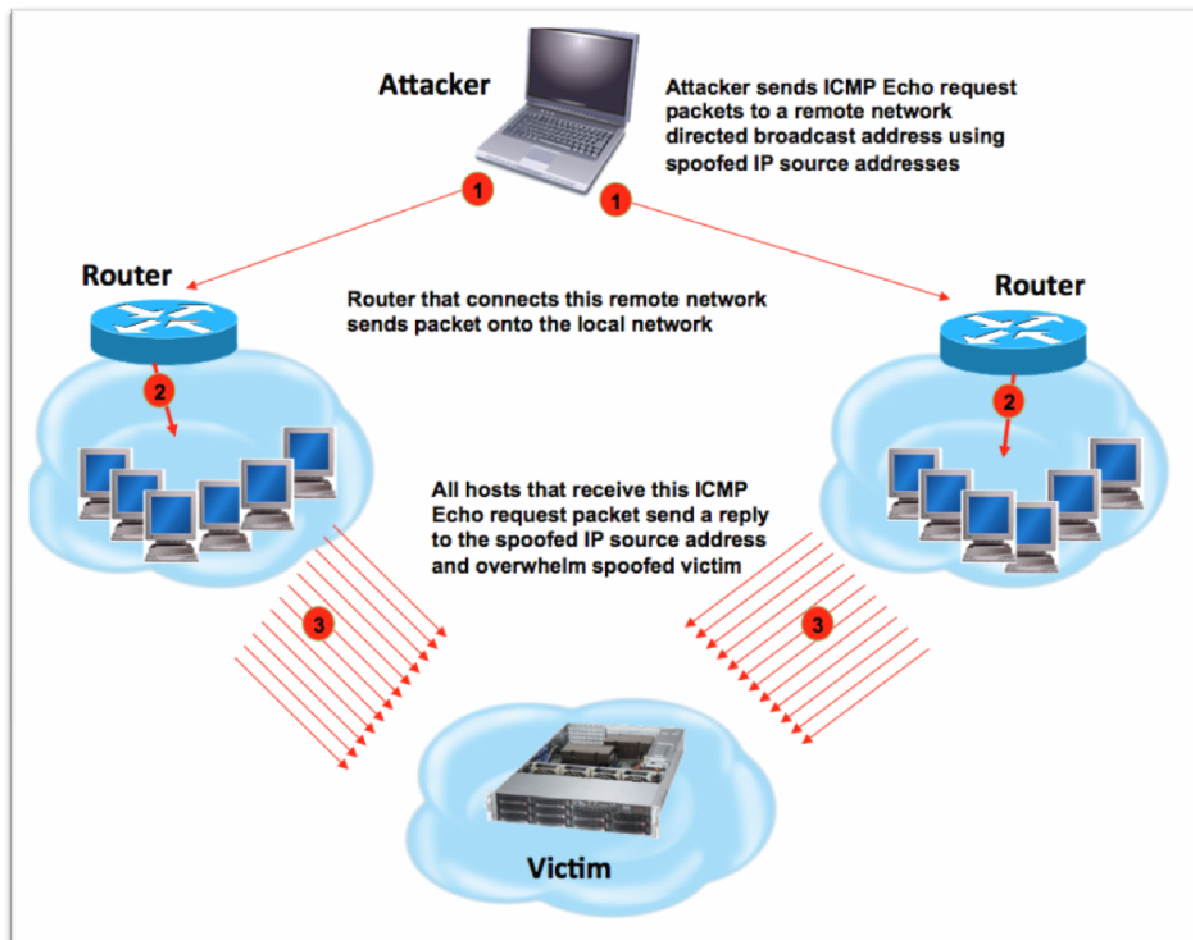


Figure 3 Example Smurf Attack

As the need arose in the late 1990's to efficiently serve larger DNS records, extensions were made to the DNS protocol to enable larger sized UDP responses to queries to avoid having to fall back to TCP [RFC 2671]. In the early 2000's it was discovered that one could use a small DNS UDP query from a forged source, and large UDP response from an authoritative server, to create an amplification effect on attack traffic¹⁴. If one controlled a computer that allowed forged UDP packets into the Internet, and attacker can use a small amount of traffic to generate a large amount of response traffic from a well-connected authoritative server. The same attack could also be performed between clients and caching resolvers. Not only were attacks effective, they were also successful in hiding the source origin of the true attacker, perhaps being launched from even another continent out of reach of the victim's law enforcement. By 2002, the first "botnet" was used to successfully implement a DNS Amplification Attack as part of its capabilities to affect critical Internet infrastructure.¹⁵

With the increasing sophistication of attack leveraging the DNS system, the Internet community

http://en.wikipedia.org/w/index.php?title=Smurf_attack&oldid=618101041.

¹⁴ Vaugh, Randal, and Gadi Evron. "DNS Amplification Attacks." .

http://www.bandwidthco.com/whitepapers/netforensics/dns/spoof-poison-hijack/DNS_Amplification_Attacks.pdf (accessed June 23, 2014).

¹⁵ "DNS Backbone DDoS Attacks." Wikipedia. March 7, 2014. Accessed August 11, 2014.

http://en.wikipedia.org/w/index.php?title=Distributed_denial_of_service_attacks_on_root_nameservers&oldid=615442441.

worked to make the DNS system more resilient. While load balancers strengthened servers' ability to process more traffic, it also had the adverse effect of increasing their impact on a victim, as much more spoofed traffic could be directed at the victim. Servers were placed behind protective layers, such as firewalls, to provide a line of defense.

Further efforts were made to inhibit malicious traffic from reaching DNS servers by implementing source IP validation techniques such as unicast Reverse Path Forwarding (uRPF)¹⁶. This would cause packets to be dropped based on the capability of the router to reach the source IP. With the release of BIND 9 in 2000, Response Rate Limiting (RRL) was another method available to limit the attack amplification provided by recursive servers¹⁷. Also included in BIND 9 was the option to disable recursion on DNS servers. It wasn't until BIND 9.4 was released in 2007 that recursion was closed-by-default.

Even with all these efforts DNS amplification attacks continue to be an issue today. These attacks continue to grow in size and complexity, such as the 2013 attack on Spamhaus recorded at over 300 GB/s. There is a continued effort to combat DDOS at all levels. One of the most recent efforts by the Internet community and industry was the US Anti-Bot Code of Conduct (ABC), which includes steps to help protect customers from being used in a botnet¹⁸.

3.3 Analysis, Findings and Recommendations

3.3.1 Analysis and Findings

The working group identified many different security issues due to Open recursive DNS servers. The issue of Open recursive DNS servers is not a new issue and the issue itself is well understood and documented. The inherent vulnerabilities of DNS and its underlying transport protocol, UDP, are documented in numerous reports from a number of organizations. The two vulnerabilities in combination with open recursive DNS servers have been and continue to be exploited by bad actors to create what is known as a Reflective DNS Amplification DDoS attack. CSRIC III's Working Group 4, DNS Best Practices, as part of their report¹⁹ provided a set of recommendations for protecting the DNS infrastructure from being exploited as summarized below:

- ISPs should refer to CSRIC III Working Group 5 for a discussion of DNSSEC
- ISPs should review DNS infrastructure to ensure it is consistent with RFC5452
- ISPs should ensure methods exist to respond to detected DNS cache poisonings
- ISPs should consider implementing DNS-specific monitoring regimes to assess the integrity of data being reported by the ISP's recursive servers

¹⁶ Cisco, "Understanding Unicast Reverse Path Forward." Accessed June 20, 2014. <http://www.cisco.com/web/about/security/intelligence/unicast-rpf.html>

¹⁷ US-CERT, "Alert (TA13-088A) DNS Amplification Attacks." Accessed June 20, 2014. <http://www.us-cert.gov/ncas/alerts/TA13-088A>.

¹⁸ Federal Communications Commission – Communications, Security, Reliability and Interoperability Council III, "FCC." Accessed June 20, 2014. http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG7_Report_March_2013.pdf

¹⁹ Federal Communications Commission - Communications, Security, Reliability and Interoperability Council III, "FCC." Accessed March 4, 2014. http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRICIII_9-12-12_WG4-FINAL-Report-DNS-Best-Practices.pdf.

- ISPs should refer to and implement practices found in CSRIC 2a – Cybersecurity Best Practices that apply to servers and ensure that recursive name server infrastructure is protected
- ISPs should adopt applicable BCPs found in other relevant network security industry approved/adopted publications including IETF RFC 2182, 2845, 3013, 3645, 5358, 4732, & 4778, NIST SP-800-53/81, and ISO publication: Towards Improving DNS Security, Stability, and Resiliency.
- ISPs should refer to and implement the Best Common Practices found in SSAC 40 & 44.

These attacks have continued and have been growing in size since the release of CSRIC III's report in 2012²⁰. In addition to CSRIC, the Internet Engineering Task Force (IETF)²¹, the Internet Corporation for Assigned Names and Numbers (ICANN) Security and Stability Advisory Committee (SSAC)²², the United States Computer Emergency Readiness Team (US-CERT)²³, Public Safety Canada²⁴, and Cisco²⁵ are just some of the organizations that have published reports that describe best practices for managing open recursive DNS servers due to the vulnerabilities with DNS and UDP.

3.3.2 CSRIC III Recommendations

CSRIC III was chartered from March 19, 2011 to March 18, 2013. CSRIC III had two working groups that looked at DNS. Working Group 4 looked at DNS Best Practices²⁶ and Working Group 5 that looked the Measurement of DNSSEC Deployment.²⁷

Working Group 4 focused on best practices to secure DNS and the routing system of the Internet during the period leading up to the successful global implementation of the Domain Name System Security Extensions (DNSSEC) and Secure BGP extensions. The working group did an exhaustive analysis of the DNS system and vulnerabilities as part of identifying recommendations for securing it.

Working Group 5 focused on best practices for implementing DNSSEC and measuring validation of DNSSEC implementations by ISPs.

²⁰ Wikipedia, "Distributed Denial of Service Attacks on Root Nameservers." Accessed March 4, 2014. http://en.wikipedia.org/w/index.php?title=Distributed_denial_of_service_attacks_on_root_nameservers&oldid=615442441.

²¹ Internet Engineering Task Force, "RFC 5358/BCP 140 - Preventing Use of Recursive Nameservers in Reflector Attacks." Accessed March 4, 2014. <http://www.ietf.org/rfc/rfc5358.txt>.

²² ICANN- Security and Stability Advisory Committee (SSAC), "SSAC Reports and Advisories." Accessed March 4, 2014. <https://www.icann.org/en/groups/ssac/documents>.

²³ US-CERT, "Alert (TA13-088A) DNS Amplification Attacks." Accessed March 4, 2014. <http://www.us-cert.gov/ncas/alerts/TA13-088A>.

²⁴ Public Safety Canada, "DNS Open Resolvers Best Practices." Accessed March 4, 2014.

<http://www.publicsafety.gc.ca/cnt/rsrscs/cybr-ctr/2013/tr13-002-eng.aspx>.

²⁵ Cisco, "DNS Best Practices, Network Protections, and Attack Identification." Accessed March 4, 2014.

<http://www.cisco.com/web/about/security/intelligence/dns-bcp.html>.

²⁶ Federal Communications Commission - Communications, Security, Reliability and Interoperability Council III, "FCC." Accessed March 4, 2014. http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRICIII_9-12-12_WG4-FINAL-Report-DNS-Best-Practices.pdf.

²⁷ Federal Communications Commission - Communications, Security, Reliability and Interoperability Council III, "FCC." Accessed March 4, 2014.

http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG5_Report_March_2013.pdf

3.3.3 Internet Standards Bodies Reports

The IETF has a number of RFCs and BCPs that it has published with respect to securing DNS. The RFCs include:

- RFC 2182 – Selection and Operation of Secondary DNS Servers
- RFC 2845 – Secret Key Transaction Authentication for DNS (TSIG)
- RFC 3013 – Recommended Internet Service Provider Security Services and Procedures
- RFC 3645 – Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG)
- RFC 4732 – Internet Denial of Service Considerations
- RFC 4778 – Current Operational Security Practices in Internet Service Provider Environments
- RFC 5358 – Preventing Use of Recursive Nameservers in Reflector Attacks

3.3.4 Group Findings

The Internet community has undertaken a number of projects to comprehensively and consistently measure the prevalence of open recursive DNS servers. Many of these efforts lack scientific due-diligence, are under resourced, and are continually undergoing improvements. Two of the larger Internet community projects – Shadowserver.org and OpenResolverProject.org – periodically scan the Internet looking for DNS servers that respond to recursive requests. Both projects generate reports with the IP addresses that respond to requests. Scanning for open recursive DNS servers is not without its challenges. Even though a server may respond when scanned, it does NOT necessarily mean that it can be used effectively in an attack. The method used by some of the Internet community projects can lead to over reporting the number open recursive DNS servers that are susceptible to attack. The other challenge is that although it is possible to identify susceptible DNS resolvers by IP address, it can be difficult to identify and contact the rightful operator of the DNS resolver.

The working group prioritized the issues it identified and surveyed these existing documents for the most appropriate way to address the identified risks. Those documents and relevant portions thereof, are then referenced both in the analysis and the group’s specific requirements.

Based upon the analysis performed, the group findings were:

- IP address spoofing continues to be one of the primary threat vectors for DDoS attacks leveraging Open recursive DNS servers.
- Misconfigured DNS servers continue to lead to DNS resolvers being configured as “open”.
- Some low-cost network-connected devices²⁸ expose DNS and other services to the Internet but have no fundamental reason to do so. These open recursive DNS servers reside on home networks outside the control of the network operator and are effectively unmanaged. The end-user of the unmanaged open recursive DNS server is unaware of the fact they are operating an open recursive DNS server and the consequence of it.

²⁸ “24 Million Home Routers Expose ISPs to Massive DNS-based DDoS Attacks - Nominum.” Nominum. April 2, 2014. Accessed August 11, 2014. <http://nominum.com/news-post/24m-home-routers-expose-ddos/>.

- DNS software updates are not consistently available nor applied²⁹
- Internet-wide community efforts to comprehensively and consistently measure the prevalence of open recursive DNS servers lacks scientific due-diligence, resulting in making it hard to quantify the impact of the previously published open recursive DNS server best practices.³⁰

4 BGP Subgroup

4.1 Objective, Scope and Methodology

4.1.1 Objective & Scope

The FCC CSRIC IV has commissioned Working Group 6 to work on improvements to Core Internet Protocols. One of the areas of attention is the routing protocols and procedures that allow for functional and proper connectivity between Internet Service Providers (ISPs). BGP (Border Gateway Protocol) is the only protocol in use today for this purpose. Over the past two decades this method for exchanging routing information between administrative domains has shown itself to be largely robust and scalable to meet the growing and changing needs of the Internet, its users and applications.

However, over that period there have also been a number of cases where the system has been interrupted due to structural weakness in the protocol or the overall routing system. For the most part, these problems have been understood as naïve configuration errors which often were able to have broad impact due to lack of adherence to best practices by some parties. In some cases, apparent intentional actions have been employed in the service of some higher-level malicious goal such as traffic disruption, injection, eavesdropping, censorship³¹, or analysis.

The previous FCC CSRIC III body also studied the topic of routing security in two separate working groups. Working Group 4 described best current practices in its report focusing primarily on various forms of route filtering at the border between ISPs and their customers.³²

The statement from Working Group 6 described the motivation behind the afore-mentioned protocol extensions, as well as some recommendations for a careful path to realize some benefits from those extensions³³.

CSRIC IV commissioned Working Group 6 to follow up on the recommendations brought forth

²⁹ “SAC065 - SSAC Advisory on DDoS Attacks Leveraging DNS Infrastructure.” *ICANN Security and Stability Advisory Committee*. (2014): 8. <https://www.icann.org/en/groups/ssac/documents/sac-065-en.pdf> (accessed March 4, 2014).

³⁰ “SAC065 - SSAC Advisory on DDoS Attacks Leveraging DNS Infrastructure.” *ICANN Security and Stability Advisory Committee*. (2014): 7. <https://www.icann.org/en/groups/ssac/documents/sac-065-en.pdf> (accessed March 4, 2014).

³¹ McCullagh, Declan. “How Pakistan Knocked YouTube Offline (and How to Make Sure It Never Happens Again) - CNET.” CNET. February 25, 2008. Accessed August 18, 2014.

³² Federal Communications Commission – Communications, Security, Reliability and Interoperability Council III, “FCC.” Accessed August 19, 2014.

http://www.fcc.gov/bureaus/pshs/advisory/csr3/CSRIC_III_WG4_Report_March_2013.pdf

³³ Federal Communications Commission - Communications, Security, Reliability and Interoperability Council III, “FCC.” Accessed July 9, 2014.

http://transition.fcc.gov/bureaus/pshs/advisory/csr3/CSRIC_III_WG6_Report_March_2013.pdf

from CSRIC III and also carry the work further by reflecting on events and developments over the intervening time. CSRIC III Working Group 6 was chartered to examine BGP and to recommend a framework for the industry regarding an incremental adoption of secure routing procedures and protocols based upon existing work in the industry and research in order to create incentives for wider scale, incremental deployment of secure BGP protocols and practices by ISPs in a market-driven, cost effective manner.

4.1.2 Methodology

The BGP subgroup of Working Group 6 began by doing a review of the recommendations and findings of the previous CSRIC BGP Working Groups. The BGP subgroup conducted bi-weekly conference calls to review recent events that have occurred involving routing systems. The group then worked to identify current efforts and developing practices to secure BGP. Based on this, the group created a list of tools and measurement projects relevant to BGP and the global routing system as well as a guide to Resource Public Key Infrastructure (RPKI).

The subgroup developed a number of main thrusts:

- Taxonomy
 - The ability to discuss a complicated technical topic coherently benefits greatly from definition and consistent use of some key terms of art.
- Recent Events
 - The interest in the topic of routing security by CSRIC IV was motivated in part by some recent events that gained some level of publicity. Some of these are reviewed and characterized in the context of the working group's efforts.
- Tools and measurement
 - The ability to measure various aspects of the global routing system are key to the understanding both of security incidents and anomalies and of the progress of changes and enhancements meant to address these types of incidents and anomalies.

4.2 Background

4.2.1 Review of CSRIC III

CSRIC III had two different working groups that touched on the topic of Routing System security. Working Group 4 focused on capabilities available in existing deployed hardware and software. Working Group 6 took up the question of enhancements to routing protocols; effectively meaning BGP.

4.2.1.1 Existing Technologies

The recommendations for improving routing security and stability with existing technologies were categorized by CSRIC III by threat vector. The recommendations of CSRIC III are summarized below by the threat vector categories used by CSRIC III:

- Session-level threats
 - Consider a plan to use MD5 or GTSM including flexibility to adjust to different deployment scenario specifics.
- DoS (Denial of Service) on routers and routing info
 - Control-Plane Policing (rate-limiting)

- Spoofed Source IP Addresses
 - 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
 - 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
 - Relying solely on AS-path filters is inadequate because it will allow mistaken re-origination or leaks of routing information to propagate
 - 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)
 - Consider many recommendations about operational security processes

4.2.1.2 Protocol Enhancement and Extensions

Working Group 6 highlighted the following high-level needs related to security of the routing system at the global level:

- Accurate records about Internet number resource holders
- Cautious, staged deployment of RPKI origin validation
- Mitigating risks inherent in the RPKI
- Improving BGP security metrics and measurements

The executive summary of that report contains a short narrative for each of these points and more supporting information in the body of the report.

4.2.2 Taxonomy

In order to arrive at a common understanding of the problem space and challenges of the improving the security of the routing system, the working group undertook to produce taxonomy of routing components and pathologies.

Overload - The possible result of traffic misdirection. Resources that commonly are overloaded in these situations may be network links or processors.

Black-hole - As a noun or more often a verb used to describe the situation where traffic is discarded by an intermediate point between the source and destination of the intended flow.

Mis-origination - A BGP routing announcement with an origin AS different from the proper and authorized origin AS. This does not necessarily connote a malicious act.

Hijack - A mis-origination or other malicious act which is intended to disrupt reachability to the legitimate destination or enable traffic disruption, injection, eavesdropping, censorship, or analysis.

Leak - Propagation of routing announcement(s) beyond the intended scope. The result can be redirection of traffic through an unintended path which may enable eavesdropping or traffic analysis and may or may not result in an overload or black-hole. Most often leaks have been accidental.

A recent IETF Internet Draft submission [draft-sriram-route-leak-protection](#)³⁴ contains a more refined classification of route leaks provides further illustration of route leaks based on observed events on the Internet..

4.2.3 Recent Events

Some events during the 2013-2014 timeframe (between CSRIC III and CSRIC IV) have helped to motivate the investigation and consideration of these topics.

4.2.3.1 Belarus and Iceland Hijack Report

One recent report³⁵ which garnered some attention from the press and the ISP community was made by Renesys (note: Renesys was acquired by Dyn in May 2014³⁶), a commercial company whose business is monitoring the global routing system and performing analysis for the purposes of security and business intelligence. In summary, this report described observations by Renesys of events within the routing system of certain events that they interpreted as incorrect routing announcements specifically crafted with malicious intent.

Specifically, Renesys proposed that the goal of the putative attackers was to re-direct traffic through an intermediate network before being delivered to its intended destination. Several types of mischief might be visited in such conditions including traffic monitoring, modification, and/or interception. Part of the analysis behind the conclusion that these incidents weren't simply accidental had to do with the particular propagation of routing updates which didn't simply result in misdirection and dropping of traffic but allowed for an apparently viable path from the intermediate network to the legitimate destination. Further, the particular networks that appeared to have been targeted were those assigned to financial and other enterprises that might have traffic of significant financial value to an attacker.

While experts may debate the interpretation of the events, it is also worth considering what weaknesses in the routing protocol and practices may have been behind the events that were

³⁴ Sriram, K., and D. Montgomery. "Draft-sriram-route-leak-protection-00 - Enhancement to BGPSEC for Protection against Route Leaks." Draft-sriram-route-leak-protection-00 - Enhancement to BGPSEC for Protection against Route Leaks. July 4, 2014. Accessed August 11, 2014.

³⁵ Cowie, Jim. "The New Threat: Targeted Internet Traffic Misdirection." The New Threat: Targeted Internet Traffic Misdirection - Renesys. November 19, 2013. Accessed August 11, 2014. <http://www.renesys.com/2013/11/mitm-internet-hijacking/>.

³⁶ "Dyn Acquires Renesys, The Global Authority On Internet Intelligence." Managed DNS. May 21, 2014. Accessed August 18, 2014. <http://dyn.com/blog/dyn-acquires-renesys-the-global-authority-on-internet-intelligence-2/>.

evident. A more recent example attack is the Canadian Bitcoin Hijack³⁷. For instance, if there was per-prefix filtering on the BGP session between the originator of the bogus routing information and the upstream ISPs, likely the problem could have been avoided. However, some speculation suggested that one or more upstream ISPs might have been collaborators in the attack. AS-path filters are insufficient to prevent a rogue AS from injecting routing information for prefixes for which they are not authorized.

If route origin AS verification based on RPKI were in place, such an attack might well still have been available to a determined attacker as shown by past research.³⁸ Full protection of the BGP AS_PATH as envisioned under the BGPSEC³⁹ might have thwarted the attack.

In the article, Renesys proposed increased monitoring of the routing system and Internet topology through use of both passive analysis of routing information and active tests of topology and reachability. Such approaches have been tested and employed in various ways over several years.

4.2.3.2 Chinese Traffic Incident

Sometimes Internet anomalies are reported in the press as “routing” incidents but aren’t actually caused by errors in the routing system, as the technical community understands it. One recent example of this phenomenon took place on January 22nd, 2014. One narrative indicated, “...many of China’s 500 million internet users were mysteriously rerouted...”⁴⁰ A significantly more technical account described how the problem appears to have been caused by the “Great Firewall” of China and related to an improper DNS to IP address mapping.⁴¹ The outcome in this case was that a remote network was flooded with unwanted traffic, and a large number of users were unable to reach the sites they wanted, but the cause was not an error in the routing system.

It should be noted that some firewalls manipulate DNS⁴² in unusual ways as part of their intended function so even methods such as DNSSEC would not have helped prevent the problem. For instance oftentimes clients inside a firewall see a private view of some DNS information. In the case of some types of application proxies, DNS clients receive modified information. DNSSEC validation is performed by recursive servers, not stub clients (e.g. user machines).

4.2.3.3 Indosat (Thailand) Incident

On April 2nd, 2014, an ISP in Thailand named Indosat (AS4761) appeared to re-originate over

³⁷ Toonk, Andree. “The Canadian Bitcoin Hijack.” BGPmon. August 12, 2014. Accessed August 18, 2014. <http://www.bgpmon.net/the-canadian-bitcoin-hijack/>.

³⁸ Pilosov, Alex, and Tony Kapela. “Stealing The Internet An Internet-Scale Man In The Middle Attack.” August 10, 2008. Accessed August 11, 2014. <https://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosov-kapela.pdf>.

³⁹ Lepinski, M., and S. Turner. “An Overview of BGPsec Draft-ietf-sidr-bgpsec-overview-05.” July 4, 2014. Accessed August 11, 2014. <http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-overview>.

⁴⁰ Perlroth, Nicole. “Big Web Crash in China: Experts Suspect Great Firewall.” Bits Big Web Crash in China Experts Suspect Great Firewall Comments. January 22, 2014. Accessed August 11, 2014.

⁴¹ Percy. “Internet Outage in China on Jan 21.” Internet Outage in China on Jan 21. January 21, 2014. Accessed August 11, 2014. <https://zh.greatfire.org/blog/2014/jan/internet-outage-china-jan-21>.

⁴² “Preparing for DNSSEC: Best Practices, Recommendations, and Tips for Successful Implementation.” Cisco. Accessed August 11, 2014. <http://www.cisco.com/web/about/security/intelligence/dnssec.html>.

400,000 prefixes (nearly the full routing table). The same AS had mis-originated a couple thousand prefixes some years earlier (2011).⁴³ In both cases if their upstream provider had applied some best practices (e.g. explicit prefix-filters, max-prefix limits) to Indosat's BGP session, the impact of this mistake could have been lessened or avoided. The incident appears to have been fairly short-lived and the impact not very widespread.

4.3 Analysis and Findings

4.3.1 Best Practices

The industry has arranged a set of best practices with regard to routing hygiene. In summary, these involve filtering routing information at ingress and egress from a given Autonomous System or other administrative domain. Optimal implementation and upkeep of these practices can require some level of diligence in dynamic ISP environments and their utilization is far from ubiquitous. This situation is what has left the door open to the naïve incidents that have been seen to date. Some methods employed by determined attackers could pierce such a level of defense.

The Internet Society (ISOC) has acted to convene a recent project for “Collective Responsibility and Collaboration for Routing Security” to spread the adoption of best practices. As of this writing, that effort has produced a draft document⁴⁴ and is circulating it in the Network Operations community for feedback, support, and awareness. These norms can be summarized as prefix-level route filters, anti-spoofing or other source-address validation, and published and reliable contact information for incident response.

4.3.2 Protocol Extensions and Enhancements

In an effort to address some of the structural weaknesses of BGP, some extensions to the basic protocol have been proposed and developed in the SIDR (Secure Inter-Domain Routing) Working Group⁴⁵ inside the IETF (Internet Engineering Task Force).⁴⁶ These proposed enhancements are still at an early stage of development and implementation. They have properties that are as yet unproven and a number of concerns have been raised among the ISP and Network Operator community about their fitness to meet the objectives for which they were designed. When using any external information source(s) in an automated fashion to effect (filter, set preference, etc.) BGP routing one must gain confidence and experience with the properties of the information sources (e.g., completeness, correctness, security / authorization model of its creation, availability / robustness of its repositories, responsiveness to change/revoked information, etc.) and the process and tool-chains that digest such information to affect running BGP processes. The introduction of new sources of information, or new process/tool-chains to consume their information must be taken with care and in an incremental fashion, until such confidence can be established. In all cases, operational plans should address

⁴³ Zmiewski, Earl. “Indonesia Hijacks the World - Renesys.” Renesys Indonesia Hijacks the World Comments. April 3, 2014. Accessed August 11, 2014. <http://www.renesys.com/2014/04/indonesia-hijacks-world/>.

⁴⁴ “Collective Responsibility and Collaboration for Routing Resilience and Security.” Collective Responsibility and Collaboration for Routing Resilience and Security. June 1, 2014. Accessed August 11, 2014. <https://www.routingmanifesto.org/manifesto/>.

⁴⁵ “Secure Inter-Domain Routing (sidr).” Secure Inter-Domain Routing (sidr) - Charter. Accessed August 11, 2014. <http://datatracker.ietf.org/wg/sidr/charter/>.

⁴⁶ “Internet Engineering Task Force (IETF).” Internet Engineering Task Force (IETF). Accessed August 11, 2014. <http://www.ietf.org/>.

the possibility that these information sources could become corrupted, or unavailable for some amount of time and require “roll back” procedures to mitigate that situation.

These concerns are true for all forms of routing policy information (I.e., IRRs, PRKI, and other sources). It is important to understand the operational policies and processes of the routing policy information repositories and the underlying properties of their information models (e.g., properties / consistency checks that can be verified by external parties). In the case of Internet Routing Registries, little is standardized in this area, and one must consult the properties of individual instances of IRRs (see: <http://www.irr.net/docs/list.html>). In the case of RPKI, the properties of the information models, the operations of repositories and the algorithms and protocols for accessing RPKI data, distilling it, and conveying it to BGP routing processes is addressed by the RFCs from the IETF SIDR WG (<http://datatracker.ietf.org/wg/sidr/documents/>). Having said that, RPKI based filtering is a relatively new technology. While production RPKI repositories are operational in all RIR regions and commercial routers ship with support for consumption and use of this information for BGP origin validation, it will take time to establish operational confidence in this new source of routing policy information and its fully automated use in production networks.

4.3.3 Tools

Over the years many different tools and projects have been devised which seek to measure or provide visualization and/or analysis related to the properties of the Internet routing system. Some of these have been developed in the academic community in the course of research. Commercial companies offer others as services.

In this section we attempt to list and catalog these efforts according to the following areas:

- Purpose – What does the project seek to measure and expose?
- Coverage – What part(s) of the problem space does it seek to address?
- Enterprise Type - Commercial or Academic
- Data Sources/Components – What underlying data sources and/or collection methods are used?
- Operational Status – Is the tool/project actively maintained? Are there operators committed to addressing problems and adapting/enhancing if necessary?

Purpose. Different tools have different goals or operational paradigms. Some seek to provide real-time alerts to users; others visualize data over time either at aggregate or specific levels of detail.

Coverage. Tools and projects have sprung up to deal with different aspects of the routing system. For many years various measurements of the routing system have emerged. From measuring the size of the global routing table to the length of more-specific BGP advertisements, some data collection projects have been going on for many years. More recently, measuring the deployment and behavior of protocol extensions has shown significant appeal.

Commercial or Academic. The availability and openness of the tool may be an outgrowth of its funding model.

Data Sources and Components. In order to perform measurements, some data must be acquired. While any particular router or ISP may have a view of the global Internet routing table, that information depends on the particulars of their interconnection relationships. The

University of Oregon began the RouteViews project in the mid-1990s to gather BGP feeds from as many ISPs as possible. As the name suggests, the initial motivation was to provide real-time operational visibility to the BGP advertisements of another network as an aid to troubleshooting. Over time, academic researchers started making use of data from the RouteViews project to do aggregate and longitudinal analyses of the routing system. Initially these academic researchers would use scripted logins to the RouteViews monitor. The RouteViews project eventually started making periodic routing table dumps available for analysis. Other projects emerged with similar methodology. RIPE RIS is another research-oriented project. Some commercial companies have undertaken similar methods as well.

Operational Status. Some academic projects emerge from research related to coursework, in pursuit of a thesis or dissertation, and/or related to grant funding. Once the particular goals of the project are achieved, the data collection and processing may or may not be maintained and continued. If the results of the analysis are going to be used as the basis for policy decisions, some level of confidence in the integrity is necessary. Furthermore, if the information from the system is used for operational purposes, the ongoing availability and reliability is important. Some accountable person or entity must stand behind the data.

4.3.3.1 General BGP Monitoring / Alerting Systems Services

A. BGPMon

- Ref: <http://www.bgppmon.net/>
- Purpose: Route monitoring and alerting service. Anomaly detection, policy violation, instability, ROA/RPKI failures.
- Components: Dedicated monitoring infrastructure and alerting service. API allows integration into other provider specific tools.
- Operational Status: Commercial service / product.

B. Renesys

- Ref: <http://www.renesys.com/products/routing-alarms/>
- Purpose: Route monitoring and alerting service. Route anomaly detection, policy violation, and instability.
- Components: Dedicated global monitoring systems,
- Operational Status: Commercial service.

C. BGP Routing Leak Detection System (Puck)

- Ref: <http://puck.nether.net/bgp/leakinfo.cgi>
- Purpose: Detect route leaks
- Components: dedicated global monitoring systems
- Operational Status: Research / Informational tool

D. PHAS

- Ref: <http://phas.netsec.colostate.edu>
- Purpose: Detect misoriginations
- Components: Route Views
- Operational Status: Defunct

4.3.3.2 Resource Public Key Infrastructure (RPKI) Services

A. ARIN RPKI

- <https://www.arin.net/resources/rpki/>

- Operational commercial RPKI services (hosted model)

B. RIPE RPKI

- <http://www.ripe.net/lir-services/resource-management/certification>
- Operational commercial RPKI services.

C. APNIC RPKI

- <https://www.apnic.net/services/services-apnic-provides/resource-certification/RPKI>
- Operational commercial RPKI services.

D. LACNIC RPKI

- <http://lacnic.net/en/rpki/>
- Operational commercial RPKI services.

E. AfrNIC RPKI

- <http://www.afrinic.net/en/initiatives/resource-certification>
- Operational commercial RPKI services.

4.3.3.3 RPKI Monitoring and Measurement Tools:

A. NIST RPKI Monitor

- Ref: <http://www-x.antd.nist.gov/rpki-monitor/>
- Purpose: Measure completeness, correctness and robustness of global RPKI.
- Components: Routeviews Data, global RPKI repository data, BBN RIPSTR RPKI Validator.
- Operational Status: Research tool.

B. RPKI Spider

- Ref: <http://rpki-spider.verisignlabs.com/>
- Purpose: Monitoring and measurement of global RPKI, focusing on robustness and performance of global repositories.
- Components: Global RPKI data, local analysis tools.
- Operational Status: Research tool from Verisign labs.

C. RPKI Dashboard

- Ref: <http://rpki.surfnet.nl/>
- Purpose: Measuring completeness and correctness of Global RPKI deployment. Per AS, Alexa500 views.
- Components: Global RPKI data, local analysis tools.
- Operational Status: Research / information tool.

D. Cobenian RPKI Browser

- Ref: <https://rpki.cobenian.com/>
- Purpose: Allows easy visual inspection of RPKI repositories.
- Components: Global RPKI repositories, web interface.
- Operational Status: Commercial service.

E. Dragon Research Rcynic

- Ref: <http://www.hactrn.net/opaque/rcynic/>
- Purpose: Displays RPKI validation information both summary and detailed
- Components: Global RPKI data, local validation software, visualization, and web interface

- Operational Status: Grant funded, non-commercial.

5 Recommendations

The recommendations from the working group are listed in this section for the DNS and routing sub-teams. CSRIC Working Group 6 recommends that the FCC encourage the voluntary adoption of the following recommendations.

5.1 Working Group 6 DNS & Routing Recommendation(s)

The recommendations that are applicable to both DNS and routing are listed below.

5.1.1 Recommendation: The Internet community should continue measurements of open recursive DNS servers and of the global routing system.

As noted in SAC065, it is recommended that members of the Internet community, such as those highlighted in section 3.3.4, should undertake a formal Internet-wide measurement tools and globally coordinated compliance program.⁴⁷ Standardized measurements of the global routing system could also be considered by the Internet community but should follow the multi-stakeholder model for development and assessment.

5.2 DNS Subgroup recommendations

5.2.1 Recommendation: Network Operators and Service Providers should follow the recommendations published in SAC065⁴⁸ as applicable

Network Operators and Service Providers should review the advisory from ICANN's Security and Stability Advisory Committee (SSAC) SAC065 – *SSAC Advisory on DDoS Attacks Leveraging DNS Infrastructure* with respect to five (5) recommendations in SAC065 that are applicable to operators of Internet infrastructure (DNS operators and network operators) and manufacturers.

As described in SAC065, the SSAC strongly recommends⁴⁹:

1. All network operators should take immediate steps to prevent network address spoofing.
2. Recursive DNS server operators should take immediate steps to secure open recursive DNS servers.
3. Authoritative DNS server operators should support efforts to investigate authoritative response rate limiting.
4. DNS server operators should put in place operational processes to ensure that their DNS software is regularly updated and communicate with their software vendors to keep abreast of the latest developments.
5. Manufacturers and/or configurators of customer premise networking equipment, including home networking equipment, should take immediate steps to secure these devices and ensure that they are field upgradable when new software is available to fix security vulnerabilities, and aggressively replace the installed base of non-upgradeable devices with upgradeable devices.

⁴⁷ See SAC065: *SSAC Advisory on DDoS Attacks Leveraging DNS Infrastructure*, 18 February 2014 at: <http://www.icann.org/en/groups/ssac/documents/sac-065-en.pdf>

⁴⁸ See SAC065: *SSAC Advisory on DDoS Attacks Leveraging DNS Infrastructure*, 18 February 2014 at: <http://www.icann.org/en/groups/ssac/documents/sac-065-en.pdf>

⁴⁹ Ibid., page 4.

5.2.2 Recommendation: Network Operators and Service Providers should segregate authoritative and recursive DNS servers.

Recursive resolvers and authoritative name servers should be run on separate servers.⁵⁰ Alternatively, they should be run on separate virtual machines (VMs) when deployed in virtualized environments, such as public clouds.

5.2.3 Recommendation: Network Operators and Service Providers should configure the DNS servers to limit the amount of information returned in queries.

Network operators and Service Providers should configure the DNS servers to return the minimum information when possible.

5.2.4 Recommendation: Network Operators and Service Providers should manage detected end-users of open recursive DNS servers.

Network Operators and Service Providers should manage, as deemed necessary; end-users detected operating un-managed open recursive DNS servers to ensure the security of the underlying network. Multiple methods exist for managing (i.e. notify, walled garden, rate limit, etc.) end-users-operated un-managed open recursive DNS servers with the goal being to reducing the attack surface of open recursive DNS servers presented by the network operator.

5.2.5 Recommendation: Network Operators and Service Providers should disable recursion on authoritative servers.

As described in RFC 5358, Network Operators and Service Providers should disable recursion on DNS servers that are meant to be authoritative only.⁵¹

5.2.6 Recommendation: Network Operators and Service Providers should investigate encouraging the use of the TCP protocol for DNS on Internet facing interfaces.

RFC 5966⁵², updates the requirements for the support of TCP as a transport protocol for DNS implementations. Network Operators and Service Providers servers should work to ensure that their DNS servers support the use of TCP as a transport protocol and encourage the use of TCP.

5.2.7 Recommendation: Equipment suppliers that include DNS resolvers in their equipment should make sure they not default to being an open recursive DNS server and make it easy to verify its proper configuration

Equipment suppliers that include DNS resolvers in their equipment including those that are embedded into customer premise equipment (CPE) devices such as home routers should have their out of the box configuration default to not operating as an open recursive DNS server. Equipment supplier that include an embedded DNS resolver should provide instructions for how to properly configure the DNS resolver so as NOT to operate as an open recursive DNS server

⁵⁰ Domain Name System (DNS) Security Reference Architecture,
http://www.dhs.gov/sites/default/files/publications/dns_reference_architecture_0.pdf

⁵¹ BCP 140 – RFC 5358: Preventing Use of Recursive Nameservers in Reflector Attacks, page 3, <https://www.rfc-editor.org/rfc/rfc5358.txt>

⁵² RFC 5966: DNS Transport over TCP – Implementation Requirements

and/or include a mechanism to get auto updates to the firmware and configurations.

5.3 Routing Subgroup recommendations

5.3.1 Recommendation: Service Providers and Networks Operators should review and use where applicable the best practices from CSRIC III Working Group 4

The Working Group recommends that Service Providers and Network Operators should review and use where applicable the Best Practices identified by CSRIC III Working Group 4's final report on BGP Security Best Practices.⁵³ Particularly critical is that ISPs maintain per-prefix route filters on customer BGP sessions and additionally consider adding some kind of AS-path filter and max-prefix limits to these and other sessions as an extra line of defense against unintentional leaks or re-originations. Other best practices recommended by that working group include:

- Network operators continue to make their own determinations in using BGP session-level counter-measures
- Network operators continue to make their own determinations in using denial-of-service against router infrastructure (interface, resource, crafted packets) counter-measures.
- Stub networks should filter traffic at the borders to ensure IP ranges assigned to them do not appear in the source-address field of incoming packets and only those ranges appear in the source-address field of outgoing packets
- Transit networks should use features such as unicast Reverse Path Filtering (uRPF) where applicable

It should also be noted that source-address validation serves to help decrease the possibility of control-plane attacks against BGP sessions, route processors, and other infrastructure components.

5.3.2 Recommendation: Service Providers and Networks Operators should monitor the development of protocols by the IETFs SIDR Working Group

Working Group 6 recommends Network Operators and Service Providers should keep monitoring the status of the protocols under development by the IETF's SIDR (Secure Inter-Domain Routing) Working Group⁵⁴. Additionally, follow GROW (Global Routing Operations Working Group)⁵⁵ which discusses operational issues that fall outside of particular protocol development. The Operations Security (OPSEC) Working Group⁵⁶ is developing a document on BGP Operations Security.⁵⁷

⁵³ "CSRIC III Working Group 4 - Network Security Best Practices: Final Report - BGP Security Best Practices." March 1, 2013. Accessed August 11, 2014.

⁵⁴ "Secure Inter-Domain Routing (sidr) - Charter." Accessed August 11, 2014.
<http://datatracker.ietf.org/wg/sidr/charter/>.

⁵⁵ "Global Routing Operations (grow) - Charter." Accessed August 11, 2014.
<http://datatracker.ietf.org/wg/grow/charter/>.

⁵⁶ "Operational Security Capabilities for IP Network Infrastructure (opsec)." - Charter. Accessed August 11, 2014.
<http://datatracker.ietf.org/wg/opsec/charter/>.

⁵⁷ Durand, J., and G. Doering. "BGP Operations and Security." Draft-ietf-opsec-bgp-security-04.txt. July 24, 2014. Accessed August 11, 2014. <http://tools.ietf.org/html/draft-ietf-opsec-bgp-security-04>.

5.3.3 Recommendation: Service Providers and Networks Operators should continue to take part in the development of the global RPKI

Working Group 6 recommends that Service Providers and Network Operators should continue to take part in the development of the global RPKI including software implementations and support by Regional Internet Registries (RIRs) of which they are members. As a follow-on to the recommendations from CSRIC III's Working Group 6⁵⁸, we recommend continued participation in the development and population of the RPKI (Resource Public Key Infrastructure) in order to gain more familiarity with its uses and properties. The decision of how to make use of the information should still remain at the discretion of Network Operators.

A guide to getting started with RPKI is located in Appendix A where a few main takeaways are:

- Make initial decision on whether you will run your own RPKI repositories or whether you will be using a hosted RPKI service. This will depend on existing resources. Information on specific hosted RPKI services are listed in Appendix A.
- Decide on how you will use RPKI information, as a relying party. This refers to using information from the RPKI to make BGP routing decisions, for example to take action such as adding a route filter or marking a route as no longer valid in some circumstances. It is recommended that no immediate action be initially taken with routes that are marked as invalid but rather, to observe the initial behavior and learn how RPKI can be actionable in your environment.

6 Conclusions

CSRIC Working Group 6 spent more than six months researching, analyzing, and evaluating industry recommendations with regards to open recursive DNS servers and the vulnerabilities associated with BGP. During this time members of this working group participated in dozens of conference calls, identified industry recommendations, and researched new recommendations, plus dedicated countless hours editing and revising the final report.

In conclusion, members of this working group feel this Final Report is a fair and accurate representation of their collective viewpoints and perspectives and hopes this report will help the Internet community.

⁵⁸ "CSRIC III Working 4 Network Security Best Practices: Final Report - BGP Security Best Practices." March 1, 2013. Accessed August 11, 2014.
http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG4_Report_March_2013.pdf.

Appendix A: Getting Started with RPKI

The RPKI (Resource Public Key Infrastructure) [RFC 6840] is a publicly available repository that provides information that can be used to help improve the security of BGP. The RPKI has been standardized by the IETF and adopted by the Regional Internet Registries (RIRs) as of 2013. The purpose of the RPKI is to provide a trusted database mapping from an IP prefix to a set of autonomous systems (ASes) that are authorized to *originate* (i.e. claim to be the destination for) this prefix in the interdomain routing system. This trusted mapping can then be used to protect against some of the most damaging attacks on interdomain routing with BGP; namely, *prefix and subprefix hijacks*, where an AS originates (“hijacks”) routes for IP prefixes that it is not authorized to originate, causing the traffic intended for those prefixes to be intercepted by the hijacker's network.

A network operator can interact with the RPKI in two ways.

- Registering BGP routes originated by the operator’s AS, and
- Using the information in the RPKI to detect problems with and/or filter routes originated by other ASes.

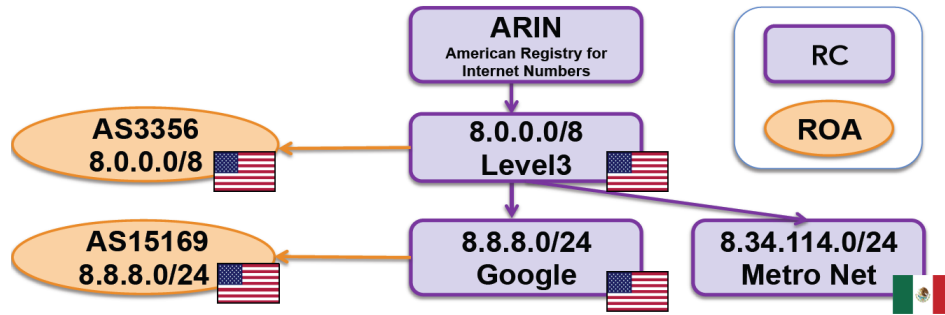
We start with some background and then discuss the options for network operators in each of the two phases.

Overview. The RPKI itself does not require changes to BGP; it operates entirely out-of-band with respect to the BGP protocol. RPKI objects are stored in public repositories. Each RPKI authority has its own *publication point* (i.e. directory in a file system) where it publishes every object it issued. *Relying parties* (“RP”s) download RPKI objects from publication points to their local caches, validate the objects, push information to their routers, and use it to inform routing decisions in BGP. RPKI repositories contain various cryptographic objects, the most important of which are *route origin authorizations*, or *ROAs*, that are used to authorize routes in BGP; specifically, each ROA authorizes a specified AS to originate a set of prefixes, and its subprefixes up to a specified length called *maxlength*, in BGP. ROAs protect BGP from prefix and subprefix hijacks.⁵⁹

Structure of the RPKI. RPKI repositories are operated by a hierarchy of *RPKI authorities*; the RPKI hierarchy mirrors the IP address allocation hierarchy. The roots of the RPKI are the five *Regional Internet Registries (RIRs)* – ARIN, RIPE, APNIC, LACNIC and AfriNIC – each holding subsets of the IP address space. The RIRs use the RPKI to sub allocate IP addresses to national/local Internet registries (NIRs or LIRs) or ISPs, who further allocate subsets to other ISPs or customers. Each RPKI authority has a *resource certificate (RC)* that contains its cryptographic public key and its set of allocated IP addresses. An RPKI authority uses its cryptographic key in its RC to issue signed cryptographic objects for IP addresses covered by its allocation, specifically: (1) an RC that suballocates a subset of its addresses to another RPKI authority, or (2) a *ROA* to authorize any AS to originate any IP prefix (where the prefix’s length may go up to a specified length called “maxlength”) that is *covered* by the IP prefixes in its own certificate. All the objects issued by an RC are stored in the RC’s *publication point* (i.e. folder or directory) in a publicly-available repository. The RC also signs a *manifest* that logs the hash

⁵⁹ An IP prefix P *covers* prefix p if p is a subset of the address space in P or if P=p. For example, 63.160.0.0/12 covers 63.160.1.0/24. Also, prefix 63.160.0.0/12 has *length* 12.

of every object present in the publication point; manifests provide allow *relying parties* to check that they correctly received all the objects in the publication point.



The figure above presents a hypothetical RPKI hierarchy, showing RC and ROAs. The ARIN RIR issues an RC that delegates 8.0.0.0/8 to Level3. Level3's RC issues an RC subdelegating 8.8.8.0/24 to Google and an RC subdelegating 8.34.114.0/24 to Metro Net, and a ROA authorizing AS 3356 to originate 8.0.0.0/8 in BGP; these objects would be stored in Level3's publication point. Google uses its RC issues a ROA authorizing AS 15169 to originate 8.8.8.0/24 in BGP; this ROA would be stored in Google's publication point. The "maxlength" parameter can be used to "collapse" multiple ROAs together; for example, Level3 could issue a ROA for prefix 8.0.0.0/8 with maxlength 9 to AS 3356; this means that AS 3356 is authorized to announce the prefixes 8.0.0.0/8 and 8.0.0.0/9 and 8.128.0.0/9 in BGP.

Hosted vs. Delegated model. Today each of the RIRs operates its hierarchy of RPKI delegations and authorizations in two different models: hosted and delegated.

In the delegated model the holder of an RC (e.g. Level3 or Google or Metro Net in the Figure) holds the private cryptographic keys for corresponding to the public key contained in its RC. The holder of the RC would also operate its own RPKI repository server, which hosts its publication point (where it publishes its RPKI objects). The advantages of the delegated model are that it allows resource holders to maintain control of their RPKI objects and cryptographic keys. On the other hand, running the delegated model is more complex, and requires the deployment and configuration of RPKI software.

The hosted model is designed to make using the RPKI easier for network operators. Instead of requiring the holder of an RC to keep track of its own cryptographic keys and to host its own RPKI repository server, all this effort is "hosted" at the RIRs. Specifically, the RIR holds all the private cryptographic keys for all "hosted" objects it issues, and makes these objects available from its own repository server that it runs. Importantly, because resource holders do not hold their own cryptographic keys in the hosted model, they are ceding all control to issue and revoke RPKI object to the RIRs. On the other hand, the hosted model requires little RPKI expertise from operators, and does not require to them to implement and maintain their own RPKI repositories; instead, they can interact with the RPKI through a simple web interface.

For some RIRs, (e.g. ARIN) only direct resource holders can participate in the hosted RPKI, and organizations that are sub allocated resources must have their parent submit ROAs on their behalf. As an example of the hosted model, suppose Level3 in the figure was registering ROAs using ARIN's hosted RPKI. ARIN would issue an RC for Level3's prefix 8.0.0.0/8; ARIN, rather than Level3 would hold the private cryptographic key whose corresponding public cryptographic key is contained in the RC for 8.0.0.0/8. Level3 would then use a web interface to request that ARIN issue a ROA authorizing AS 3356 to originate prefix 8.0.0.0/8 in BGP; ARIN would then use the private cryptographic key in the RC for 8.0.0.0/8 to issue this ROA. Because

Google is not a direct resource holder, Google have to make sure that Level3 also asked ARIN to issue a ROA authorizing AS 15169 to originate prefix 8.8.8.0/24; in contrast to what is shown in the figure, this ROA would be issued directly from the RC for 8.0.0.0/8, using the private key held by ARIN.

Operators using the RPKI can choose between the hosted and delegated model; however, for some RIRs, the delegated model is still in its experimental stages.

Relying parties. A *relying party* is any party (e.g. an AS, a router) that downloads and uses information in the RPKI. A relying party downloads all the information available in the RPKI repositories to its local cache, and then cryptographically validates them, and then uses this data to inform routing decision made in BGP. Specifically, a relying party uses its "local cache of the complete set of valid ROAs" [Sec. 2, RFC6483] to classify each route learned in BGP into one of three *route validation states*:

- **Valid:** The route has a valid *matching ROA*. A matching ROA has (1) the same origin AS as the BGP route, and (2) a prefix P that covers prefix p in the BGP route, and (3) the specified *maxlength* is no shorter than the length of the prefix p in the BGP route.
- **Unknown:** The route has no valid *covering ROA*. A covering ROA is any ROA for a prefix that covers the prefix p in the BGP route.
- **Invalid:** The route is neither unknown nor valid. That is, there is a valid covering ROA, but no matching ROA.

For example, the RPKI in the Figure makes a BGP route for 8.0.0.0/8 originated by AS 3356 *valid* (because of the matching ROA), and a route for 8.0.0.0/8 originated by AS 666 *invalid* (because of the covering but not matching ROA for 8.0.0.0/8). Any route for 12.0.0.0/8 is *unknown* (because there is no covering ROA). Meanwhile, any route for 8.34.114.0/24 is *invalid* (because of the covering ROA for 8.0.0.0/8).

These route validity states can then be used to inform routing decisions made in BGP. (We discuss different options for how network operators can use this information as relying parties in Step 2.)

Status of the RPKI. As of January 13, 2014, the production RPKI contains ROAs for about 20K prefix-to-origin-AS pairs. (About 488K prefixes were announced in BGP that day.) At depth 0 there are trust anchor certificates for each RIR (i.e. ARIN, RIPE, APNIC, LACNIC, AfriNIC). The trust anchors are long-lived certificates, which issue a handful of shorter-lived intermediate RCs that are also held by the RIRs. Intermediate RCs issue leaf RCs to organizations (Sprint, Swisscomm, etc.) that then issue ROAs. ARIN has an extra layer of intermediate RCs. ROAs usually contain one AS and many IP prefixes. Operators can use following RPKI looking glasses to view the objects present in the RPKI, and see how these objects tag the validity state of production BGP routes:

<http://rpki-monitor.antd.nist.gov/>

<http://localcert.ripe.net:8088/trust-anchors>

http://www.labs.lacnic.net/rpkitools/looking_glass/

Step 1: Register your own routes

When registering routes in the RPKI, resource holders need to choose between using a hosted

RPKI service and operating their own RPKI repositories. The hosted model is easier to use, but requires operators to cede all control to issue and revoke RPKI objects to the RIRs.

Deciding what routes to register. Before starting to register ROAs with the RPKI, the operator should make sure he/she understand how all prefixes and subprefixes of their IP address allocations are being used in BGP.

To see why this is important, we can refer back to the Figure, and the route validity rules discussed in the previous section. The valid ROA for 8.0.0.0/8 that covers all sub prefixes of 8.0.0.0/8, means that all BGP routes for sub prefixes of 8.0.0.0/8 are automatically *invalid* unless they have their own matching ROAs. For example, if MetroNet in the Figure forgets to issue valid ROAs for its own routes for prefix 8.34.114.0/24, all these routes will be *invalid* (rather than unknown) because of the covering ROA issued by Level3. Thus, every ROA that is issued for a prefix immediately makes all routes for its sub prefixes invalid; thus, if an operator wishes to issue a ROA for a prefix P , it must make sure it also issue ROAs for *each and every route* for a sub prefix of P that can legitimately be originated in BGP. This is true for both the hosted and delegated model of the RPKI.

Thus, an operator that adds a ROA for its routes to the RPKI should do so carefully, otherwise it could mistakenly cause other route (for all sub prefixes of its prefix) to become invalid. More discussion on this important issue that can commonly lead to misconfigurations is available in this blog post:

<https://labs.ripe.net/Members/waehlich/one-day-in-the-life-of-rpki>

Procedure for registering, using the hosted model. Here is a suggested procedure for adding ROAs to the RPKI using the **hosted model** for the RPKI.

1. Determine which prefixes will be registered with the RPKI.
2. For each prefix P , figure out all the routes originated in BGP for **all prefixes π covered** by prefix P . This can be done using this tool:

<http://localcert.ripe.net:8088/bgp-preview>

For each prefix π , determine the set of ASes that are allowed to originate that prefix in BGP. You should now have a set of pairs (π, a) where prefix π is covered by prefix P and AS a is an AS that is authorized to originate IP prefix π in BGP. Make sure you have every prefix and AS pair covered by prefix P !

3. Go to the RPKI website of the RIR who allocated prefix P to your organization and request a ROAs for each set of pairs (π, a) , following the instructions provided by each RIR. Here are the websites for each RIR:

<https://www.arin.net/resources/rpki/>

<http://www.ripe.net/lir-services/resource-management/certification>

<https://www.apnic.net/services/services-apnic-provides/resource-certification>

<http://www.lacnic.net/en/web/lacnic/certificacion-de-recursos-rpki>

<http://www.afrinic.net/en/initiatives/resource-certification>

4. Visit an RPKI looking glass, and make sure that you have not forgotten to register a route. You can do this by visiting the RIPE validator and typing in prefix **P** into the search box. You should make sure that any route covered by **P** has not become “invalid” as a result of registering the ROA for **P**.

<http://localcert.ripe.net:8088/bgp-preview>

Procedure for registering, using the delegated model. This model is still in an experimental stage at many RIRs. Advice on how to do this is available, for example, at:

<http://www.ripe.net/lir-services/resource-management/certification/managing-certificates>
(see “running your own certificate authority)

<https://www.arin.net/resources/rpki/delegatedrpki.html>

Once you have registered ROAs. To ensure that the RPKI does not become a repository for stale and useless information, operators have a responsibility to make sure that they ROAs they added to the RPKI are up to date; if a particular AS is no longer meant to be originating a sub prefix of **P** in BGP, then the ROA for that AS should be deleted from the RPKI. Also, if a new AS starts originating a sub prefix of **P**, then a new ROA for that AS should be registered immediately in the RPKI; otherwise, the new route will be classified as “invalid” by the RPKI. Also, the RPKI has a concept of expiry; each item in the system has a validity period expressed as a data range. If the item is not updated with a new date range before it expires, the result will be that the object and others that have a hierarchical reliance on it will become invalid.

Step 2: Decide how you will use RPKI information, as a relying party.

Network operators can also use the RPKI as “relying parties”, using information from the RPKI to inform routing decisions made in BGP; importantly, when using the RPKI as a relying party, an operator can use information registered by other operators, for other prefixes

When a BGP security mechanism such as RPKI indicates that a route is invalid, Service Providers have a variety of actions they can take based on that information. The action a Service Provider takes can vary depending on some details of the route. For example, if a route is associated with a customer the Service Provider may take different actions than if the route is associated with a peer. The country associated with the prefix, the prefix’s length, or other characteristics may also be used to determine a course of action.

One option is for Service Providers to log violations (e.g. routes learned in BGP that the RPKI tags as “invalid”) and take no action. This allows a service provider to gather information about how frequent alerts are, or it may be an appropriate action for routes that do not affect the Service Providers customers.

Alternatively, Service Providers can use violations to trigger an alert to the Network Operations Center or Security Operations Center for either immediate human analysis/ response or for a post incident analysis. The human can take a number of actions, such as adding a route filter, calling peers, announcing more specific routes, or contacting the organization associated with the route.

Violations can also trigger automated technical responses. These could include modifications of preference on the violating route (modify local-pref to make the route less preferred for example), announcement of defensive routes, such as more specifics, or automatic

implementation of filters on peers sending bad routes. The automated technical responses can be either embedded in routers, or implemented from a separate system that interfaces with the router control plane or management plane.

One set of automated technical responses recommended by the RPKI specifications involves filtering or depreferencing “invalid” routes. Filtering invalid routes in an automated way is one possible policy; which this policy prevents prefix and sub prefix hijacks, it also comes with the risk that a misconfiguration or error in the RPKI can legitimate BGP routes to become unreachable. An alternative, more lenient policy [see RFC6483] is to *depref* “invalid” routes for a given prefix; that is, a router should prefer “valid” routes over “invalid” routes (this policy implies that a router still selects an “invalid” route when there is no “valid” route for the exact same IP prefix!). Thus, the router may still be able to reach routes that are wrongly classified as “invalid” as a result of problems with the RPKI. On the other hand, this policy does **not** automatically prevent subprefix hijacks; see Section 5 of RFC 7115.

When deploying a system each Service Provider will need to determine what actions meet its individual network requirements. Those actions will likely change over time. For example, a Service Provider can start by logging all violations and reviewing that list monthly, investigating a subset of violations in detail. After a period of time, that Service Provider could start sending some violation alerts to the NOC for immediate response. Eventually, the Service Provider could automate the most common violation types.

Today, for example, RPKI looking glasses observe a large number of invalid routes; many of these are likely due to misconfigurations, when resource holders register a ROA for a prefix, but forget to register ROAs for its subprefixes. Misconfigurations of this type suggest that more operational experience is required, and that Service Providers should start by being cautious about using RPKI information as relying parties.