

SAC121
SSAC Briefing on Routing Security

Preface

This is a report to the ICANN Board, the ICANN organization staff, the ICANN community, and, more broadly, the Internet community from the ICANN Security and Stability Advisory Committee (SSAC) about routing security.

The SSAC focuses on matters relating to the security and integrity of the Internet's naming and address allocation systems. This includes operational matters (e.g., pertaining to the correct and reliable operation of the root zone publication system), administrative matters (e.g., pertaining to address allocation and Internet number assignment), and registration matters (e.g., pertaining to registry and registrar services). SSAC engages in ongoing threat assessment and risk analysis of the Internet naming and address allocation services to assess where the principal threats to stability and security lie, and advises the ICANN community accordingly. The SSAC has no authority to regulate, enforce, or adjudicate. Those functions belong to other parties, and the advice offered here should be evaluated on its merits.

Table of Contents

Executive Summary	5
1 Introduction	6
1.1 Intended Audience and Use	6
1.2 Background	6
1.3 The Border Gateway Protocol	8
1.4 Hypothetical Route Hijack	9
2 Examples of Routing Incidents	10
2.1 Youtube / Pakistan Telecom	10
2.2 MyEtherWallet / Route53	10
2.3 Safe Host / China Telecom	11
2.4 3ve	11
2.5 Issues with Determining Intent	12
3 The Relevance of Routing Security for the DNS	13
3.1 The DNS is Susceptible to Routing Incidents	13
3.2 Portions of the DNS Infrastructure Susceptible to Routing Hijacks	13
3.2.1 Denial of Service	15
3.3 Alleviating Routing Risks to DNS Resolution	16
4. Routing Security Mechanisms	17
4.1 Routing Registries	18
4.2 Resource Public Key Infrastructure	19
4.2.1 RPKI Distribution	20
4.2.2 Use of the RPKI for BGP	21
4.3 AS Path and BGPsec	22
5 Operating Secured Infrastructure	22
5.1 Monitoring	22
5.1.1 Endogenous Monitoring - The View from Inside	23
5.1.2 Exogenous Monitoring - The View from Outside	24
5.2 Operator Coordination	25
6 Summary	25
7 Acknowledgments, Statements of Interest, and Dissents, Alternative Views and Withdrawals	26
7.1 Acknowledgments	27
7.2 Disclosures of Interest	27

SSAC Briefing on Routing Security

7.3 Dissents and Alternative Views	27
7.4 Withdrawals	27
Appendix A: Additional References	28
Appendix B: RPKI	34
Appendix C: BGPsec	35

Executive Summary

Like all other Internet applications, the Domain Name System (DNS) depends on the Internet's routing system, which controls the data paths across the Internet's more than 70,000 autonomously managed networks. A longstanding problem with the routing system is that its key protocol, the border gateway protocol (BGP), does not protect against incorrect routing information. BGP was designed in the late 1980's and early 1990's when the Internet consisted of only a few hundred networks that all trusted one another. As the Internet grew and the number of networks increased, the number of routing incidents increased and this implicit trustworthiness waned. The routing system today is subject to a continuous stream of routing anomalies that affect its integrity and that sometimes cause large DNS outages. For example, in April of 2018 attackers were able to "hijack" routes to Amazon's Route53 DNS services, which resulted in DNS traffic for domains hosted on this service ending up at a different destination network where it was served by malicious DNS servers.

In this report, the SSAC discusses events like these and what impact similar incidents can have on the DNS, surveys the pros and cons of various solutions, and discusses future security extensions of the routing system (e.g., path validation). The main focus of this report is on the security and stability implications for the DNS, although most of it also applies to other types of Internet applications (e.g., email, web, media streaming).

This report provides a tutorial-style discussion accessible to non-technical members of the ICANN community and elsewhere (e.g., policy makers and legal experts). It does not contain any recommendations to the ICANN Board. Because this report is intended to be understandable to a non-technical audience, it sometimes simplifies technical details that are not relevant to the discussion.

1 Introduction

As listed in the ICANN Bylaws, one of the SSAC's roles is to, "advise the ICANN community and Board on matters relating to the security and integrity of the Internet's naming and address allocation systems."¹ This report examines the security and stability of the Internet's routing system, primarily as it pertains to operators and users of the DNS. It documents this space in an effort to help the larger ICANN and Internet policy communities understand these technologies and the issues surrounding them. The report also provides advice to DNS operators on how they can better secure their routing infrastructure.

Much has been written about Internet routing security in other venues, in academia, and in the Internet routing community. The major contribution of this report is to examine Internet routing security from the perspective of those primarily concerned with operating DNS infrastructure, and with a focus on issues that are relevant to the larger ICANN community.

1.1 Intended Audience and Use

This report will make significant references to the Border Gateway Protocol (BGP). While the document gives some background, for readers unfamiliar with BGP, there are many useful overviews of the protocol online.²

The intended audience of this report are those who operate their own DNS, network or routing infrastructure. Some organizations outsource the operation of their DNS or other networking infrastructure to large providers, thereby outsourcing the knowledge required to operate routing infrastructure. This report may still be useful to those who choose to outsource, because it covers topics that someone outsourcing should ensure their vendor is capable of competently handling.

While this report is aimed at the DNS audience, and so implicitly an audience that is familiar and concerned with DNS operations, the operators of other types of Internet infrastructure may also find useful information in this report.

1.2 Background

The Internet is a network of networks: it consists of more than 70,000 autonomously managed networks that collaboratively transport data from one machine on the Internet to another.³ As a result, a data flow that traverses the Internet usually passes through a sequence of intermediate networks. However, the end points are typically unaware of these intermediate networks. To these end points the Internet is a black box where they simply put in packets that come out at the remote destination. End points do not have any control over which networks their data passes through because each intermediate network has its own view of the Internet, and therefore makes its own decision for the next best network to pass the traffic through.

¹ See ICANN Bylaws, Section 12(b)(i), <https://www.icann.org/resources/pages/governance/bylaws-en/#article12>

² See Cloudflare. "What Is BGP? | BGP Routing Explained." Accessed May 20, 2022. <https://www.cloudflare.com/learning/security/glossary/what-is-bgp/>.

³ See "CIDR Report." Accessed May 20, 2022. https://www.cidr-report.org/as2.0/#General_Status.

Figure 1 illustrates how the Internet transports data for a set of DNS messages. The top of the figure (above the solid line) shows the typical DNS behavior:

1. Client C queries recursive resolver D for the IP address of www.example.net.
2. D iteratively queries the authoritative name server(s) in the DNS.
3. D returns the IP address (192.0.2.1) to C.

For readability, Figure 1 combines the multiple request-response interactions between D and the authoritative name servers into a single bidirectional arrow.

The bottom of the figure (below the solid line) shows that each of these DNS entities (client, resolver, and authoritative server) exist in an autonomous system (AS), or network. For client C to query recursive resolver D the query travels from AS1 → AS2 → AS3, and the response travels AS3 → AS2 → AS1. Likewise, for recursive resolver D to query authoritative server A its traffic travels from AS3 → AS6 and the response traffic travels AS6 → AS3. These DNS entities, like nearly all computers on the Internet, are typically unaware of the one or more networks through which their queries, as well as any responses they receive, may traverse.

The three DNS entities in Figure 1 are part of an AS (AS1, AS3, and A6), but they are drawn outside the ASes to emphasize that the DNS uses the Internet’s routing and forwarding system, just like any other application.

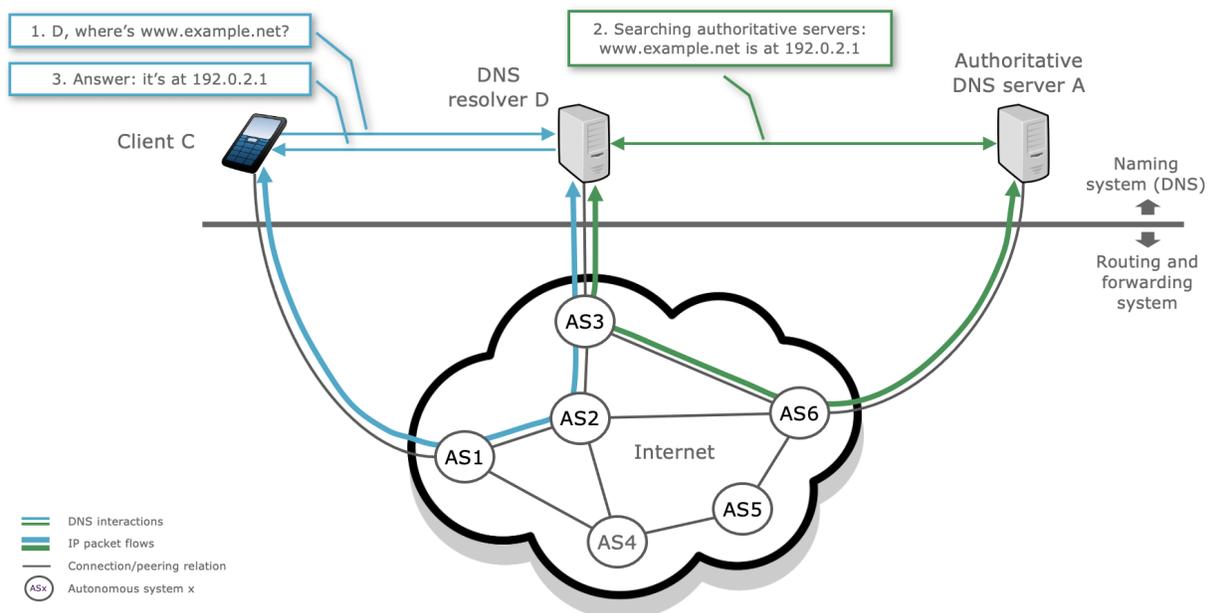


Figure 1: DNS traffic passing through multiple autonomous systems

1.3 The Border Gateway Protocol

BGP is the primary routing protocol used on the Internet. It allows a collection of network devices all running the BGP protocol (BGP speakers) to each learn which connected network provides the best path to a destination (IP address prefix). The basic approach is very simple: each BGP speaker tells all its neighbors about what it has just learned, but only if this new information alters its local view of the network (i.e., it uses the information itself). This is a lot like a social rumor network, where every individual who hears a new rumor that interests them immediately informs all their friends. With BGP, each time a neighbor informs another BGP speaker about a change in reachability to an IP address prefix, that BGP speaker compares this new information against its current knowledge of that prefix. If this new information describes what the local device considers a more preferred path to the prefix, then the BGP speaker tells all its immediate BGP neighbors of this more preferred path, implicitly citing itself as the next hop or viable path to reach the destination. If this new information describes the removal (i.e., withdrawal) of the IP prefix, then the BGP speaker removes the withdrawn path and informs its neighbors.

There is an implicit trust model in the BGP protocol, in that every BGP speaker trusts every other BGP speaker to only announce destinations for which they intend to provide reachability. This is the critical security factor for the Internet's routing system. If any BGP speaker does not adhere to this process and tells its neighbors incorrect information, then receiving BGP routers may use this incorrect information when calculating packet forwarding decisions. This can result in passing packets past observation points or disrupting network connectivity altogether.

How can other BGP speakers tell if a BGP speaker is not behaving as it should? This is the central question for routing security.

1.4 Hypothetical Route Hijack

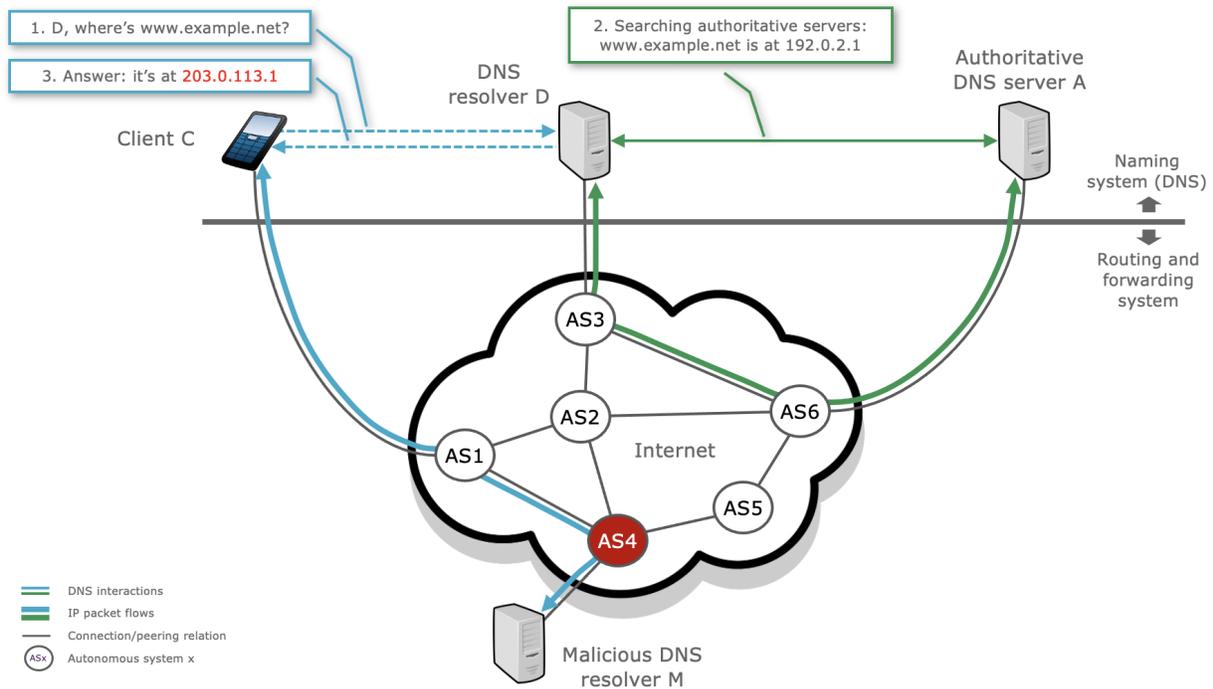


Figure 2: Hypothetical Route Hijack affecting the DNS

Figure 2 shows a routing hijack in progress. The user of the mobile phone C wishes to visit `www.example.net`, and must first perform a DNS lookup to discover its IP address. The mobile device first sends a DNS query to DNS resolver D, its configured recursive resolver. The correct route for this query is $AS1 \rightarrow AS2 \rightarrow AS3$. However, in Figure 2, AS4 has announced the IP address space of resolver D as well, and because of this AS1 believes that their best route to resolver D is AS4, not via AS2.

This means the query from client C travels along the path of $AS1 \rightarrow AS4$, and is then sent to the malicious recursive resolver M in AS4. M responds to the DNS query indicating that `203.0.113.1` is the IP address of `www.example.net` instead of the legitimate address `192.0.2.1` resolver D would have provided. As a result, C contacts the address that M returned, where there is likely a malicious server provisioned, enabling the user to be phished.

A routing hijack such as the one in Figure 2 thus has the effect of compromising the integrity of the DNS. Unless the client performs DNSSEC validation, which is rarely implemented in clients, it will trust the DNS response sent to it by the malicious resolver. Users have little or no visibility into the routing system, or the DNS, and therefore will not notice this kind of attack, which can make this kind of route hijack particularly insidious.

2 Examples of Routing Incidents

This section contains some examples of incidents where propagation of spurious routes have negatively affected users and operators of Internet infrastructure. The small set of examples here should not be misconstrued to indicate that routing incidents are uncommon. A continual stream of observed routing anomalies are noted by BGP observatories,⁴ and the selection criteria used for these examples illustrates the diversity of the incidents, and the diversity of intent behind them. Most of the observed incidents can be attributed to mistakes rather than intentional actions, and the intent behind these actions is often difficult, if not impossible, to ascertain. It is also difficult to assess the damage these incidents cause. The implications of many routing incidents are seldom fully understood.

2.1 Youtube / Pakistan Telecom

On February 24, 2008, the popular online video streaming site YouTube (youtube.com) was globally unreachable for approximately two hours due to a misconfiguration at Pakistan Telecom. Engineers at Pakistan Telecom were directed to block Pakistan users from accessing YouTube. They attempted this by creating a more specific route in their own network for YouTube's IP addresses. Unfortunately, this more specific route leaked out of Pakistan Telecom to their upstream provider, PCCW Global, who then preferred this route and propagated it to the rest of the Internet, thereby directing all traffic intended for YouTube to Pakistan Telecom.^{5,6}

2.2 MyEtherWallet / Route53

On April 25, 2018, the Ethereum trading site MyEtherWallet (myetherwallet.com) was attacked by unidentified criminals using a BGP hijacking attack. Ethereum is a blockchain-based cryptocurrency similar to Bitcoin. The criminals were able to steal about \$150,000 in Ethereum from users in the approximately two hours during which they were able to redirect users to their malicious server.

The attackers injected more specific routes for some of Amazon's Route53 servers, the authoritative DNS provider for MyEtherWallet. Unlike the YouTube / Pakistan Telecom incident described in Section 2.1, this was an attack on the authoritative DNS servers of the actual target. By hijacking the routes to the target's authoritative DNS servers hosted at Route53, the attackers were able to answer DNS queries to direct victims to the attacker's server that impersonated the real myetherwallet.com.

⁴ See Testart, Cecilia, Philipp Richter, Alistair King, Alberto Dainotti, and David Clark. "Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table." In Proceedings of the Internet Measurement Conference, 420–34. Amsterdam Netherlands: ACM, 2019. <https://doi.org/10.1145/3355369.3355581>.

⁵ See Sarrafzadeh, M. "How YouTube Was 'Hijacked.'" ACM SIGDA Newsletter 20, no. 1 (2009): 91. <https://doi.org/10.1145/378886.380416>.

⁶ See RIPE-NCC. "YouTube Hijacking: A RIPE NCC RIS Case Study." RIPE Network Coordination Centre, March 17, 2008. <https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study>.

The attackers announced the routes via eNet Inc, who then announced the routes to their upstream provider, Hurricane Electric, who then announced them to all of their customers. The attackers then hosted their own authoritative DNS servers for MyEtherWallet that directed users to the attacker's server where users were robbed of their Ethereum.

Many domain names hosted at Amazon's Route53 DNS service were affected by this attack, not just myetherwallet.com. The DNS servers that the attackers used only responded to queries for myetherwallet.com, so all other names that were queried returned SERVFAIL. Thus, for approximately two hours some Internet users could not resolve any names hosted on Amazon's Route53 DNS service.^{7,8} The effects of such an attack are proportional not only to the duration of the hijacking event, but also to the time-to-live (TTL) of the DNS responses (i.e., indicator of how long to hold a DNS record in a resolver cache) provided by the attacker's purportedly authoritative DNS name server.

2.3 Safe Host / China Telecom

On June 6, 2019, traffic from multiple locations on the Internet was routed through China Telecom for roughly two hours. The incident began when Swiss-based hosting company Safe Host erroneously advertised more than 70,000 routes to China Telecom. China Telecom then advertised these routes to its neighboring networks instead of filtering them and began receiving traffic for those 70,000 networks.⁹

2.4 3ve

3ve (pronounced “eve”) was a massive multifaceted operation to defraud advertising networks by using hijacked IP address ranges and malware.¹⁰ Uncovered in 2018, the name 3ve was coined by the researchers who discovered it because their analysis determined it was made up of three distinct sub-operations.

All told, 3ve controlled over 1 million IPs from both residential botnet infections and corporate IP spaces [...]. In aggregate, the operation also produced more than 10,000 counterfeit domains, and generated over 3 billion daily bid requests at its peak. We estimate that portions of the bot operation spanned over 1,000 servers in data centers allocated to various functions needed for this type of large-scale operation.¹¹

⁷ See Poinsignon, Louis. “BGP Leaks and Cryptocurrencies.” The Cloudflare Blog, April 24, 2018. <http://blog.cloudflare.com/bgp-leaks-and-crypto-currencies/>.

⁸ See Siddiqui, Aftab. “What Happened? The Amazon Route 53 BGP Hijack to Take Over Ethereum Cryptocurrency Wallets.” Internet Society, April 27, 2018. <https://www.internetsociety.org/blog/2018/04/amazons-route-53-bgp-hijack/>.

⁹ See Goodin, Dan. “BGP Event Sends European Mobile Traffic through China Telecom for 2 Hours.” Ars Technica, June 8, 2019. <https://arstechnica.com/information-technology/2019/06/bgp-mishap-sends-european-mobile-traffic-through-china-t-elecom-for-2-hours/>.

¹⁰ See Google, and White Ops. “The Hunt for 3ve: Taking down a Major Ad Fraud Operation through Industry Collaboration,” November 2018. https://services.google.com/fh/files/blogs/3ve_google_whiteops_whitepaper_final_nov_2018.pdf.

¹¹ *Ibid.*, page 10

3ve first established fake websites and then used a combination of malware installed on unsuspecting users' computers, and BGP route hijacking to impersonate users visiting these fake websites. These "users" would then trigger what the online advertising industry calls an "ad bid request", which initiates a bidding auction between advertisers wishing to show an advertisement to this "user". The advertiser that wins the auction pays the fake website to display their advertisement to the fake user. Most of the payment for the ad imprint is transferred to 3ve.

To facilitate their route hijacking, 3ve set up two real autonomous systems (ASes). One the researchers nicknamed ALPHA was set up in 2013, and another nicknamed BRAVO was set up in 2017. From a routing system perspective, ALPHA operated normally for four years before beginning to hijack networks in March 2017. This initial period of inactivity allowed them to establish a good reputation and history with other ASes. The networks they chose to hijack were mostly defunct ISPs and other networks that no one would notice were being routed by a different AS.

2.5 Issues with Determining Intent

Two terms are often used to describe unwanted routing events on the Internet, route leaks and route hijacks. Other groups and researchers have defined these and similar terms in different ways.^{12,13,14} In this report their primary difference is the real or perceived intent of the entity initially sending the improper information.

The commonly used term "route leak" is often attributed to a BGP speaker that unintentionally propagates a route without authorization. Often route leaks are attributed to user error, software bugs or improperly applied filters on routers. In most cases where a routing event is classified as a leak the impacted prefixes became unreachable by some portion of the Internet, resulting in a denial of services to the impacted address ranges. Section 2.1 discusses a widely known leak, there are others.^{15,16}

The commonly used term "route hijack", on the other hand, is often attributed to malice, perpetrated by an entity able to send BGP announcements intending that they are accepted and propagated to other BGP speakers. In some cases of route hijacks, the services being operated in the redirected ranges will appear to operate as normal when viewed by an end user or system

¹² See RFC 7908: Problem Definition and Classification of BGP Route Leaks

¹³ See Cho, Shinyoung, Romain Fontugne, Kenjiro Cho, Alberto Dainotti, and Phillipa Gill. "BGP Hijacking Classification." In 2019 Network Traffic Measurement and Analysis Conference (TMA), 25–32. Paris, France: IEEE, 2019. <https://doi.org/10.23919/TMA.2019.8784511>.

¹⁴ See Sriram, Kotikalapudi, and Doug Montgomery. "Resilient Interdomain Traffic Exchange:: BGP Security and DDoS Mitigation." Gaithersburg, MD: National Institute of Standards and Technology, December 2019. <https://doi.org/10.6028/NIST.SP.800-189>.

¹⁵ See Kirk, Jeremy. "Indosat Routing Error Impacts Few but Hits Akamai, Chevron." Network World, April 3, 2014. <https://www.networkworld.com/article/2175839/indosat-routing-error-impacts-few-but-hits-akamai--chevron.html>.

¹⁶ See Siddiqui, Aftab. "Major Route Leak by AS28548 – Another BGP Optimizer?," February 13, 2021. <https://www.manrs.org/2021/02/major-route-leak-by-as28548-another-bgp-optimizer/>.

administrator, but the entity hijacking the route(s) may use, alter or drop some traffic. Sections 2.2 and 2.4 provide examples of malicious intent.

This document uses the generic term “routing incident” to refer to an event where the intent is unknown or not relevant.

3 The Relevance of Routing Security for the DNS

3.1 The DNS is Susceptible to Routing Incidents

The DNS protocol and DNS resolution are susceptible to routing incidents. By design, many authoritative DNS servers are promiscuous and will answer any query they receive. Generally, DNS clients do not authenticate the identity of the server that provides the answer, and do not DNSSEC validate signed responses. Additionally, stub resolvers have no visibility into which authoritative servers provided answers to the recursive resolver they originally queried. A routing attack can thus easily enable substitution of one server for another.

By default DNS transactions occur in the clear, without any transport encryption. There are new technologies being deployed to address this (e.g., DNS-over-TLS, DNS-over-HTTPS, DNS-over-QUIC),^{17,18,19} but currently the vast majority of DNS queries are in the clear and use UDP as the transport protocol. Absent the client validating the server’s identity, as commonly found in transport encryption such as TLS, a routing attack can successfully substitute one server for the intended server without the awareness of the client. This allows the substituting server to behave in a potentially malicious manner, such as returning incorrect data.

Routing attacks can alter the network path of a query, thereby allowing third parties to inspect DNS queries or otherwise eavesdrop on transactions. This has obvious implications for the privacy of DNS queries and other plain text protocols. It also opens up an attack vector where the observer can forge a response. Because DNS resolution is commonly implemented as a single transaction, DNS clients will accept the first response they get as the anticipated response. Subsequent responses will then be ignored. Thus, a malicious actor who is able to respond faster (e.g., an on-path attacker closer to the DNS client) than the legitimate responder may be believed and have a high probability to succeed in their attack.

3.2 Portions of the DNS Infrastructure Susceptible to Routing Hijacks

A typical DNS transaction starts with a stub resolver sending a query to a recursive resolver, causing that recursive resolver to query multiple authoritative servers. This is shown in *Figure 1*: Client C sends an initial query to Recursive Resolver D, which then recursively queries multiple authoritative servers, including Authoritative Server A. Any one of these transaction sets can be hijacked and the response passed back to the client. As the popularity of public recursive

¹⁷ See RFC 7858: Specification for DNS over Transport Layer Security (TLS)

¹⁸ See RFC 8484: DNS Queries over HTTPS (DoH)

¹⁹ See RFC 9250: DNS over Dedicated QUIC Connections

resolvers increases, the traffic between the stub and recursive resolver becomes a more interesting target for attackers. When stubs send queries to recursive resolvers not located in the same AS or otherwise topologically local, there is the possibility of a route hijack interfering with their communications.

Recursive resolvers may perform queries to multiple authoritative servers in order to resolve a single name, with each query determining the name servers of the next zone that needs to be queried, and therefore the intended destination of subsequent packets. If an attacker can hijack one of these queries successfully, they can easily determine where every subsequent query will go, and thereby forge the resulting response.

The increasing use of anycast by DNS authoritative server operators adds a new wrinkle to an already complex situation.²⁰ Anycast is a technology that uses multiple servers to answer traffic sent to the same IP address. All DNS servers in an anycast network should behave roughly the same, and respond in the same way to the queries they receive. If an attacker can perform a route hijack against a single anycast instance it can be very difficult to detect by the anycast operator, since all the other anycast instances will continue to behave normally. If the attacker has a specific target in mind, and they know that this target will always use a specific anycast instance in their DNS resolution path, the attacker only has to disrupt the routing between that single instance and the target. If an anycast network operator uses unique autonomous system numbers (ASNs) for each anycast instance it may be easier for the operator to detect when new instances are introduced by attackers.²¹

Transport Layer encryption on the Internet primarily uses Transport Layer Security (TLS) and TLS end-point identity is rooted in domain names. When a client makes a TLS connection to an end-point it includes in its handshake a Server Name Indication (SNI) which is a domain name. When the end-point responds it sends back a certificate that is linked to this domain name. This way the client can ensure they are connecting to a server that is controlled by the same party that controls the domain name listed in the server's certificate. If an attacker is able to convince a public key infrastructure (PKI) certification authority (CA) that they control a domain they can get a certificate issued for that domain, and thereby convince any client connecting to their malicious server that they are a legitimate server for this domain. Certification authorities rely on the Internet's routing system just like any other Internet end-point. It should be noted that protection against these kinds of attacks can be improved by using DNSSEC validation, DNS CAA records,²² and HTTP Strict Transport Security (HSTS),²³ and similar security mechanisms. Such mechanisms make it significantly harder for an attacker to succeed with route hijacks.

This implies that if an attacker can intercept and reply to DNS queries (e.g., TXT record checks) coming from a certification authority, and the domains being queried are not DNSSEC signed,

²⁰ See Fan, Xun, John Heidemann, and Ramesh Govindan. "Evaluating Anycast in the Domain Name System." In 2013 Proceedings IEEE INFOCOM, 1681–89. Turin, Italy: IEEE, 2013. <https://doi.org/10.1109/INFOCOM.2013.6566965>.

²¹ See RFC 6382: Unique Origin Autonomous System Numbers (ASNs) per Node for Globally Anycasted Services

²² See RFC 6844: DNS Certification Authority Authorization (CAA) Resource Record

²³ See RFC 6797: HTTP Strict Transport Security (HSTS)

the attacker may be able to impersonate the targeted domain and acquire a TLS certificate for that domain.^{24,25}

In addition to the risks posed by routing hijacks on DNS resolution, routing security is also important for the registration of domain names. This goes for the connection between the registrant and registrar, as well as the connection between the registrar and registry. Both of these connections are normally secured using Transport Layer Security (TLS). Registrants connecting to registrars will typically use HTTPS, and registrars connecting to registries should be using Extensible Provisioning Protocol (EPP) over TLS with client certificates.²⁶ If the client and server do not implement mutual TLS validation (mTLS) a route hijack attack could successfully impersonate the client.²⁷

3.2.1 Denial of Service

The DNS was designed for an Internet where uninterrupted connectivity could not be assumed, the DNS often includes a component of redundant provisioning (i.e., multiple name servers) and because of this remains relatively resistant to denial of service (DoS) attacks. Caching and built-in retrying in resolvers accounts for much of this resiliency, and the use of anycast for authoritative servers helps here as well.

Stub resolvers usually have more than one recursive resolver configured. If the first one fails they will use the second, and there are usually multiple authoritative servers as well. Once a recursive resolver has resolved a name, that information should stay cached in that resolver for the length of the time-to-live (TTL) of the record. This redundancy and caching baked into the DNS protocol requires that effective DoS attacks against the DNS be prolonged and broad. An attacker will usually have an easier time attacking something else, such as a web server, if the goal is to interrupt service.

There is an important trade off for operators to consider when configuring TTLs for DNS records their infrastructure depends on.^{28,29} Configuring longer TTLs means that in the case of a DoS attack these records will persist for longer in caches and be more difficult to take offline or otherwise update. However, this comes with the cost of not being able to make quick changes to DNS information. Some organizations may wish to make regular and frequent changes to their DNS records and have those changes propagate quickly through DNS recursive server caches. While short TTLs may facilitate rapid changes, if the authoritative servers hosting these records

²⁴ See Birge-Lee, Henry, Yixin Sun, Annie Edmundson, Jennifer Rexford, and Prateek Mittal. “Using BGP to Acquire Bogus TLS Certificates,” 2017, 2. <https://petsymposium.org/2017/papers/hotpets/bgp-bogus-tls.pdf>.

²⁵ See Talon, S2W, and Sojun Ryu. “Post Mortem of KlaySwap Incident through BGP Hijacking.” *S2W BLOG* (blog), February 17, 2022.

<https://medium.com/s2wblog/post-mortem-of-klayswap-incident-through-bgp-hijacking-en-3ed7e33de600>.

²⁶ See RFC 5734: Extensible Provisioning Protocol (EPP) Transport over TCP

²⁷ See RFC 8446: The Transport Layer Security (TLS) Protocol Version 1.3

²⁸ See RFC 9199: Considerations for Large Authoritative DNS Server Operators

²⁹ See Moura, Giovane C. M., John Heidemann, Ricardo de O. Schmidt, and Wes Hardaker. “Cache Me If You Can: Effects of DNS Time-to-Live.” In Proceedings of the Internet Measurement Conference, 101–15. Amsterdam Netherlands: ACM, 2019. <https://doi.org/10.1145/3355369.3355568>.

become unavailable for even a short period of time the records will also become unavailable. One mitigation for this is for resolver operators to continue serving data past when the authoritative servers they originally received that data from become unreachable. Details such as how long to keep this stale data, and when to discard it are described in RFC 8767.³⁰

The route hijack attack MyEtherWallet / Route53 (See Section 2.2) illustrates this tradeoff well. The attackers executed a route hijack against Amazon's Route53 DNS service and returned SERVFAIL for all queries not related to myetherwallet.com, the attacker's intended target. The attack lasted for roughly two hours. Thus, DNS zones hosted at Route53 with TTLs shorter than two hours did not persist in DNS caches for the entire attack, and their corresponding sites suffered outages as a result. Other sites using DNS records with longer TTLs were more likely to stay online during the Route53 DNS outage.

Authoritative DNS operators can increase their resiliency against DoS attacks and routing incidents by hosting authoritative servers in different networks. Using different autonomous systems, with different network prefixes, and with different providers increases operational diversity.³¹

Unlike the DNS resolution processes, the DNS provisioning processes are not able to make use of caching to protect against DoS attacks. Consider the DNS provisioning processes of how a registrant registers, or initiates changes, to a domain via a registrar: a registrant logs into a registrar's website and initiates a change; that registrar then sends EPP commands to a registry's EPP endpoint; finally, that registry records and makes changes to DNS records it hosts. Any of the Internet-accessible endpoints (e.g., registry's EPP endpoint) are susceptible to a DoS attack by malicious actors wishing to interrupt their service. However, the services hosted at these endpoints are rarely time critical, and usually a short postponement of their availability will not result in dire implications for any party. Thus, a DoS attack against these endpoints must be prolonged to be most impactful.

3.3 Alleviating Routing Risks to DNS Resolution

While routing security is important, it is not a substitute for other technologies also key to securing the DNS. It is only one part of a complete approach to securing a network. Other technologies such as DNSSEC, domain name certificates for TLS and channel encryption remain important as well.

DNS-over-TLS (DoT), DNS-over-HTTPS (DoH), and DNS-over-QUIC (DoQ) are relatively new technologies that provide channel encryption between clients and recursive resolvers.^{32,33} They are currently only defined for the transit leg between the stub and recursive resolver, and both provide transport layer encryption. In cases where the DNS data being queried is DNSSEC

³⁰ See RFC 8767: Serving Stale Data to Improve DNS Resiliency

³¹ See SAC005: DNS Infrastructure Recommendation Of the Security and Stability Advisory Committee

³² See SAC109: The Implications of DNS over HTTPS and DNS over TLS

³³ See RFC 9250: DNS over Dedicated QUIC Connections

signed, and the recursive resolver performs DNSSEC validation, these protocols can help secure the final leg between the stub resolver and the DNSSEC validating recursive resolver.

DNSSEC can ensure the authenticity of data received from DNS responses if the response data is signed and the querying party performs validation. Typically DNSSEC validation occurs in the recursive resolver. A stub sends a query to a recursive resolver, which performs the necessary queries to resolve the name, and then also performs DNSSEC validation on all the responses. This means that if all the data is signed, the recursive resolver is able to validate all the data it receives from authoritative servers before sending any data back to the stub.

If the DNS stub resolver implemented DNSSEC validation of signed responses, the client would detect any compromise of the recursive resolver that results in false responses, independent of the protocol used for the DNS queries. However, few DNS stub resolvers perform DNSSEC validation. Therefore if an attacker is able to subvert the connection between the stub and the recursive resolver this validation exercise is rendered irrelevant. DNSSEC itself is not perfect, and should not be relied upon as the *only* mechanism for defense against all potential attacks.

4. Routing Security Mechanisms

In 1994 when RFC1105 introduced BGP, a rogue BGP speaker was not a substantial concern of the protocol design.³⁴ As the Internet grew, and the number of attached networks increased, the implicit trustworthiness of all networks on the Internet waned. BGP was not designed to operate with assured integrity when the accuracy of some routing information is doubtful, leaving the network operator community with the task of developing mitigations to the issues that have emerged. This section covers the primary ongoing efforts introduced to enhance Internet routing security.

While much of this section focuses on BGP, Internet routing security encompasses more than just the BGP protocol. Internet routing security is contingent upon many things: BGP protocol security, accuracy of routing policy, and robust operations. All of these are important. Securing BGP information is important. However, if an operator does not configure their routing policies correctly they may improperly communicate a route. Likewise, operators should monitor their networks, practice good operational security, and follow sound operational procedures.³⁵

It is safe to assume that attacks that leverage weaknesses in the routing system will continue, including those that exploit new defenses we describe, for example by using insider threats or coercion. Operators must remain vigilant and continue to improve their posture as new tools and tactics emerge.

³⁴ See RFC 1105: A Border Gateway Protocol (BGP)

³⁵ See “MANRS – Mutually Agreed Norms for Routing Security,” May 20, 2022. <https://www.manrs.org/>.

4.1 Routing Registries

A network operator can register their autonomous system (AS) and the prefixes they originate in a routing registry. This allows other operators to consult the registry to see what prefixes and routes a given AS should be announcing. The Internet Routing Registry (IRR)³⁶ pre-dates most other efforts in this space, and dates back to the routing work of the early 1990's. Internet routing registries describe, in advance, all of the routes that a network is capable of announcing so that other networks could reject announcements not in this list.³⁷

Each participating operator submits policy data.^{38,39} Clients may use the routing registry to determine the stated policies for a particular AS, including what ASes are *providers* for this AS, and which ASes are *peers* or *customers* of this AS. Additional information provided to a routing registry by an operator could include policy concerning the configuration of BGP communities and the policy responses associated with particular community settings.

However, the utility of a routing registry approach for securing routing has some limitations. Firstly, a routing registry does not only provide information about currently active routes, but normally also includes a larger set of potential routes. Some potential routes may be validly described according to a routing registry, but undesirable from a more global point of view. Second, the quality of the data in the routing registries relies on sufficient authorization of changes or additions to the registry. In addition, some ASes do not want to publicly expose their peering agreements, and routing registries do not normally implement any form of limited disclosure of registry contents.

The contents of different routing registries are not necessarily mutually consistent and there is no clear way to resolve conflicts between them. There is no common authority model ensuring that only authorized parties may publish routing policy data about their own address prefixes and ASes, there is also no common way to describe the intended lifetime of the information these registries contain. Old information that is no longer current or relevant sits alongside current information, and this sits along with contingency information that may never be actually used.⁴⁰

Routing registries are most useful when they are carefully and continuously managed for consistency, coverage and accuracy. The accuracy and usefulness of the information rapidly declines if the information in the registry is neglected. Routing registries appear to work best in defined and carefully scoped contexts with active curation of the registry content for its

³⁶ See "IRR - Internet Routing Registry." Accessed May 20, 2022. <https://www.irr.net/>.

³⁷ See RFC 7682: Considerations for Internet Routing Registries (IRRs) and Routing Policy Configuration

³⁸ See RFC 2622: Routing Policy Specification Language (RPSL)

³⁹ See RFC 4012: Routing Policy Specification Language next generation (RPSLNg)

⁴⁰ *Contingency information* refers to information intended to be used in case of an emergency.

completeness and accuracy. When they take on a broader scope, or are not actively managed, their consistency and utility falls.^{41,42,43}

It is not uncommon for transit providers and IXPs to require peers and customers to register routing information in IRRs, so that the provider can create their routing filters based on this registered information.⁴⁴ However, such IRR-based protection against route hijacks may be circumvented by malicious updates of IRR data. As described earlier, this is an arms race. As operators develop defenses against existing vulnerabilities attackers will seek ways to subvert those defenses.⁴⁵ Some researchers have proposed that an additional layer of independently verifiable accountability will be necessary to identify trustworthy actors in the Internet resource ecosystem.⁴⁶

4.2 Resource Public Key Infrastructure

An approach to provide protection of the integrity of the content of BGP messages is to use digital signatures to provide a set of credentials that allow relying parties to verify the correctness of the information carried in BGP. This does not protect the BGP message directly, but is intended to provide a way for a receiver of a message to validate certain aspects of the information in a BGP update message. This partial approach resolves some, but not all existing issues related to the security of BGP operation. Discussions on the efficacy of the RPKI are ongoing.^{47,48} At the same time RPKI is strongly recommended or may soon be required by some regulators.⁴⁹

⁴¹ See RFC 2725: Routing Policy System Security

⁴² See Du, Ben, Gautam Akiwate, Thomas Krenc, Cecilia Testart, Alexander Marder, Bradley Huffaker, Alex C. Snoeren, and KC Claffy. “IRR Hygiene in the RPKI Era.” In *Passive and Active Measurement: 23rd International Conference, PAM 2022, Virtual Event, March 28–30, 2022, Proceedings*, 321–37. Berlin, Heidelberg: Springer-Verlag, 2022. https://doi.org/10.1007/978-3-030-98785-5_14.

⁴³ See Kuerbis, Brenden, and Milton Mueller. “Internet Routing Registries, Data Governance, and Security.” *Journal of Cyber Policy* 2, no. 1 (January 2, 2017): 64–81. <https://doi.org/10.1080/23738871.2017.1295092>.

⁴⁴ See RIPE NCC. BGP Security: IRR and Filtering Webinar - 25/03/2021, 2021. <https://www.youtube.com/watch?v=GUyDL0zNUUY>.

⁴⁵ See Dai, Tianxiang, Philipp Jeitner, Haya Shulman, and Michael Waidner. “The Hijackers Guide To The Galaxy: {Off-Path} Taking Over Internet Resources.” In *30th USENIX Security Symposium*, 3147–64. USENIX Association, 2021. <https://www.usenix.org/conference/usenixsecurity21/presentation/dai>.

⁴⁶ See Clark, David D., and K. C. Claffy. “Trust Zones: A Path to a More Secure Internet Infrastructure.” *TPRC48: The 48th Research Conference on Communication, Information and Internet Policy*, November 30, 2020, 20. <https://doi.org/10.2139/ssrn.3746071>.

⁴⁷ See Shrishak, Kris, and Haya Shulman. “Limiting the Power of RPKI Authorities.” In *Proceedings of the Applied Networking Research Workshop*, 12–18. Virtual Event Spain: ACM, 2020. <https://doi.org/10.1145/3404868.3406674>.

⁴⁸ See Heilman, Ethan, Danny Cooper, Leonid Reyzin, and Sharon Goldberg. “From the Consent of the Routed: Improving the Transparency of the RPKI.” In *Proceedings of the 2014 ACM Conference on SIGCOMM*, 51–62. Chicago Illinois USA: ACM, 2014. <https://doi.org/10.1145/2619239.2626293>.

⁴⁹ See Dekker, Marnix, Eleni Vytogianni, and Aggelos Koukounas. “7 Steps to Shore up BGP.” Report/Study. European Union Agency for Cybersecurity, 2019. <https://data.europa.eu/doi/10.2824/66344>.

Digital signatures are used because the number and identities of all eventual recipients of the information are not known in advance, and non-repudiation is desirable.⁵⁰ Verification of the contents of an update message is not only a test of whether the BGP information has been altered in any way during its transit between BGP speakers, but also a test of whether the message represents information relating to a valid address prefix, valid AS numbers, correct origination information, and correct processing of the message during propagation.

This requires a means of verification where the issuer of any security credentials relating to origination and propagation is not necessarily known to the party that is validating the information. This typically uses a form of Public Key Infrastructure (PKI), which involves a chain of trust from a trust anchor to the subject of the verification. The PKI is used to associate a public key with an entity that has functional control of an IP address prefix or an AS number. The certificate issuance practices are intended to support transitive trust in this association.

The Resource Public Key Infrastructure (RPKI) has adopted X.509 public key certificates and a certificate extension that uses a list of IP address resources and AS numbers as the foundation for this number RPKI.^{51,52} The holder of the matching private key is the current functional controller of those IP addresses and AS number and can digitally sign authorities and attestations about such number resources within the context of the RPKI.

The RPKI is different from PKI in other contexts as the requirements related to adding digital signatures to the routing domain are different from those of other PKI deployment environments. The common question that a PKI attempts to answer is, "Is this data authentic?" For a typical PKI this is a singular question relating to one subject drawn from the larger set of certificates that make up the PKI (e.g., TLS validating a single certificate). But when we look at routing and the RPKI, the routing data is not singular, but encompasses the entire collection of information, and the test is applied to the entire set of attestations that compose the RPKI. An RPKI user⁵³ cannot choose, or even know, when an attestation to validate a routing message will be needed, because routes can change at any time. What this means is that every RPKI user needs to have access to the entire collection of signed RPKI attestations at all times, which is the major point that makes the RPKI different.

4.2.1 RPKI Distribution

This requirement for all participating entities to have access to all the RPKI data at all times poses a design challenge in how to manage the RPKI and use it in a routing protocol such as BGP.

A basic approach here is for each Internet Registry to publish their RPKI certificate products in their own publication point. This is analogous to the pre-CDN model of web content publication, where each element is independently published. In this case publication is easy, but the onus is

⁵⁰ *Non-repudiation* in this context refers to the specific inability of the digital signer to deny they are the entity that signed the message.

⁵¹ See RFC 3779: X.509 Extensions for IP Addresses and AS Identifiers

⁵² See RFC 6487: A Profile for X.509 PKIX Resource Certificates

⁵³ A user is generally called a relying party in the nomenclature of PKI.

shifted to the relying party client or BGP validator, who has to assemble a local cache of all RPKI signed data. It becomes the task of RPKI clients to maintain a local cache of the entire RPKI by continuously sweeping across these publication points looking for and retrieving changes, and validating all such signed objects as they are received.

The drawback of this distributed approach is that there is no notification model to advise relying parties of any changes to the set of published products. This implies a need for these clients to constantly sweep all the RPKI publication points to ensure that their local cache is up to date. What “up to date” means is relative here, but it is worth remembering that there is a considerable time lag between publication of updates and RPKI-aware BGP speakers acting on this updated information, potentially of the order of minutes, or even hours. That represents a potentially challenging timing objective if the system has to scale up to the size of the Internet's routing environment.⁵⁴

4.2.2 Use of the RPKI for BGP

Route origin validation (ROV) builds upon the earlier work on Routing Registries (See Section 4.1), where a prefix holder is able to publish information as to how an address prefix is to be announced into the routing system by designating the AS number(s) that are permitted to originate a routing announcement for that prefix. In the RPKI framework this information is published as a signed route origin authorization (ROA).^{55,56} An ROA signed by an address prefix holder denotes permission for an AS to originate a route.

There are a number of additional implications associated with publishing an ROA. The first is that no other AS has permission to announce that prefix when there is a cryptographically valid ROA extant in the RPKI system. If the prefix holder wishes to authorize multiple ASes to originate a route for a particular prefix, then the prefix holder must generate multiple ROAs. This means that an address holder can declare that a prefix should not be routed at all by issuing an ROA that provides permission to AS0 (AS zero). Secondly, the ROA denies permission for any AS to originate a prefix that is more specific than the prefix listed in the ROA. There is a MaxLength attribute of an ROA that may be used to define a range of more-specific prefix lengths that are permitted by an ROA. Thirdly, there is no acknowledgment of the ROA on the part of the designated AS. A prefix holder may publish an ROA providing permission to an AS that is unaware of the permission.

⁵⁴ See Kristoff, John, Randy Bush, Chris Kanich, George Michaelson, Amreesh Phokeer, Thomas C. Schmidt, and Matthias Wählisch. “On Measuring RPKI Relying Parties.” In Proceedings of the ACM Internet Measurement Conference, 484–91. Virtual Event USA: ACM, 2020. <https://doi.org/10.1145/3419394.3423622>.

⁵⁵ See RFC 6482: A Profile for Route Origin Authorizations (ROAs)

⁵⁶ See RFC 6483: Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)

Use of the RPKI for BGP is in the process of being deployed by network operators and there are a few websites where this progress can be tracked.^{57,58,59} See Appendix B for more information and references on the RPKI and its use in BGP.

4.3 AS Path and BGPsec

The BGPsec protocol is an extension to BGP that is intended to allow validation of the AS Path attribute.⁶⁰ It represents a relatively high overhead to pay for a limited set of assurances and a limited protective capability. Furthermore, there is a more extreme view that BGPsec cannot achieve any of the security properties due to the fundamental design principles of BGP and BGPsec. See Appendix C for more discussion of these points.

5 Operating Secured Infrastructure

5.1 Monitoring

Network monitoring is the continual observation of a network to detect anomalies and failures. It is how an operator gains intelligence about the actual operation of their network and establishes a baseline that reflects normal operation. The primary purpose of network monitoring is to detect failures and other abnormalities, so that they can be fixed quickly and efficiently. Monitoring can also help in predicting imminent or intermittent failures before they become permanent, and help spot any anomalous behavior that requires further investigation. Good monitoring of a network provides the engineers responsible for its operation pointers to things that they should investigate further, without overwhelming them with false positives that waste their time.

Anomalous behavior is often first identified in application misbehavior. For example, if an authoritative DNS server starts returning many NXDOMAIN responses or experiences a large drop in the number of queries it is receiving, there is likely a DNS or routing issue that warrants further investigation. It is usually easier and cheaper for operators of DNS infrastructure to monitor the behavior of their DNS infrastructure than it is for them to set up monitoring of their routing infrastructure (e.g., monitoring BGP neighbors). If budget is available it is best to monitor both application behavior and underlying routing infrastructure, since monitoring only DNS will not uncover all routing issues. Waiting for routing issues to manifest as application issues can leave routing issues unattended and will impact service in the meantime as applications may continue to function in a degraded fashion during partial routing or DNS failures.

To spot irregular behavior it is first necessary to identify and note what regular behavior is. Network traffic and applications show patterns of behavior, and a good monitoring system will help reveal those patterns to an observer. As the observer becomes aware of these expected patterns, the observer will then trigger on behaviors that are outside of this comfortable norm.

⁵⁷ See Cloudflare. “Is BGP Safe yet?” Accessed May 20, 2022. <https://isbgpsafeyet.com/>.

⁵⁸ See NIST. “NIST RPKI Monitor.” Accessed May 20, 2022. <https://rpki-monitor.antd.nist.gov/>.

⁵⁹ See RIPE-NCC. “RPKI Statistics.” Accessed May 20, 2022. <https://certification-stats.ripe.net/>.

⁶⁰ See RFC 8205: BGPsec Protocol Specification

There is a certain amount of learned experience and intuition in an observer's ability to ignore expected behavior, and that improves over time. A good monitoring system should facilitate this learning process, and be configurable enough to allow for evolving system behavior.

Every organization responsible for keeping a network up and running will need to find their own balance, and make their own decisions about how much to budget for monitoring. One way to think about monitoring is that it mitigates risk to an organization's business by providing information about its operations. Like other decisions around mitigating risk, such as purchasing insurance for important assets, every organization will have its own appetite for risk and will make its own decisions about how best to manage risk. There is typically a diminishing rate of return for investments in monitoring. As more money is spent on monitoring the marginal capability decreases. Thus, small investments in monitoring can have a large effect, while subsequent investments will provide less useful information than earlier investments.

Network monitoring can be done in-house or outsourced depending on the needs and available budget of the organization responsible for the network. The best solution for most network operators will likely be a combination of both. Since every operation is different, and vendors usually want to sell generic services to as many customers as possible, some customization will inevitably be required for any organization that wishes to comprehensively monitor their network and applications.

The rest of this section discusses monitoring from within a network (endogenous monitoring), and monitoring from outside of a network (exogenous monitoring).

5.1.1 Endogenous Monitoring - The View from Inside

Endogenous monitoring is when an organization monitors their ability to reach other networks from their own network. It is typically cheaper and easier to do than exogenous monitoring since it can be self hosted and there exist off the shelf and open source tools for this kind of monitoring.

Every network will have a different 'view' of the Internet depending on what other networks they connect to, and how they connect to them. Most networks on the Internet have well-defined upstream providers who provide them service to the rest of the Internet. The term 'dual-homed' is used to describe networks that connect to the rest of the Internet via two or more providers for redundancy and best path selection. For these kinds of networks their connectivity to the Internet is entirely dependent on the providers they have chosen. What should be monitored is then determined by their connectivity to their providers' networks, and the connectivity between their providers and the rest of the Internet. The most important connectivity to monitor for these kinds of networks are the connections to their upstream providers. It is the upstream provider's job to ensure connectivity to the greater Internet.

Outsourcing connectivity to an upstream provider does not outsource the responsibility of maintaining connectivity. If some part of the Internet is unreachable it is still the reputation of the organization providing the network that will be harmed in the case of downtime. If an organization knows that many of their customers come from a specific region of the world they

may want to specifically monitor the connectivity between their network and that region, and to important network providers in that region.

5.1.2 Exogenous Monitoring - The View from Outside

Exogenous monitoring is when an organization monitors connectivity with their own network from other networks. It allows an organization to monitor how their network is seen from other networks on the Internet. Without this kind of monitoring an organization may not be able to discover how other networks connect to them nor detect anomalies in those connections.

This monitoring is important, but is usually more expensive because it requires the installation of vantage points outside of the network being monitored, and outside vendors will typically need to be engaged to place these vantage points. This is a complex space with diverse offerings, and vendor solutions tend to be wrapped in externally provided services that are not an ideal match to any particular set of monitoring requirements. There is always a trade off between how much an organization wants to monitor, and how much that organization is willing to spend on monitoring.

Anycast adds considerable complexity to exogenous monitoring. In order for an operator to monitor every anycast instance deployed it is probably necessary to have at least one vantage point close to each anycast instance. However, as routing in the Internet is dynamic and subject to change, a vantage point that connects to one anycast instance today may connect to a different one tomorrow.

While specific DNS queries can be used to determine which anycast instance is responding, and thereby the closest instance, it must be noted that an attacker can forge these responses thereby tricking the monitor into thinking it is connecting to a specific anycast instance when it in fact is not.^{61,62,63} One difficult aspect of monitoring an anycast network is determining whether a monitor's closest instance changed because of a malicious or benign routing change, and if because of this change there is now an instance, or multiple instances, that are no longer being monitored. An attacker may also be able to deploy their own anycast instance that is closer to the target of an attack.

Using a vendor's anycast network is often easier than an organization deploying their own, but the individual instances still need to be monitored, and the monitoring solution provided by the vendor may not be a good fit for the application that needs to be monitored. Once again, there is no ideal set of requirements for monitoring an anycast network and tradeoffs must be made that take into account available budget and security requirements. Understanding these trade offs and making informed decisions regarding them requires that an organization have access to specific expertise on anycast networking, and not just general routing expertise.

⁶¹ See RFC 4892: Requirements for a Mechanism Identifying a Name Server Instance

⁶² See RFC 5001: DNS Name Server Identifier (NSID) Option

⁶³ See Fan, Xun, John Heidemann, and Ramesh Govindan. "Evaluating Anycast in the Domain Name System." In 2013 Proceedings IEEE INFOCOM, 1681–89. Turin, Italy: IEEE, 2013.
<https://doi.org/10.1109/INFCOM.2013.6566965>.

5.2 Operator Coordination

When problems occur it is important for operators to know who to contact outside of their organization to help remediate issues. For networks that do not provide transit services for other networks, and that only interact with one or two upstream providers, this can be as simple as knowing someone at these providers to call when things go wrong. For networks that provide transit services for other networks the list of contacts will be more extensive. It is largely the number and kind of connections a network has with other networks that should determine the level of effort invested in developing good contacts at other network operators.⁶⁴ Regardless of the scale and type of collaboration desired, the important thing is to develop relationships before a problem occurs.

Network operator groups (NOGs) exist partly to help facilitate this relationship building. Any network operator, regardless of how complex their network is, can benefit from participating in their regional NOG and getting to know the people who operate the networks in their region. Regional NOGs exist at many levels from metro, to national, to regional. Not every issue is global in nature, and not every operational issue requires the assistance of all NOGs. Thus local NOGs are generally the first point of contact for many issues. For network operators that will invariably experience downtime, or other serious issues, investing time and social capital into building relationships and trust up front is well worth the cost.⁶⁵

NOGs also facilitate information sharing and provide education for staff of network operators who may not otherwise have access to knowledge sharing outlets, especially in regions where knowledge of networking fundamentals may be lacking. Participating in a NOG, even if only online, enhances the expertise of networkers while also facilitating their network and career growth.

6 Summary

Every transaction that takes place on the Internet relies on the routing system to function. The DNS, since it relies on making transactions over the Internet therefore relies on routing as well. The routing system is complex, and packets often traverse many networks before reaching their destination. Each network, or autonomous system (AS), establishes relationships with each of its adjacent networks and learns of distant networks through a process that can be described as *routing by rumor*. Because of this, each network with more than one connection may have its own view of the Internet, and its own best routes to every other network connected to the Internet.

⁶⁴ See Meier-Hahn, Uta. "Creating Connectivity: Trust, Distrust and Social Microstructures at the Core of the Internet." SSRN Electronic Journal, 2015. <https://doi.org/10.2139/ssrn.2587843>.

⁶⁵ See Mathew, Ashwin, and Coye Cheshire. "The New Cartographers: Trust and Social Order within the Internet Infrastructure," August 15, 2010, 21. <https://papers.ssrn.com/abstract=1988216>.

Attackers can take advantage of this by injecting false routes, or other false information into the routing system. Mistakes can also occur through misconfigurations, and it can be difficult, if not impossible to distinguish between whether a routing incident was an unintentional accident or an intentional malicious act to divert traffic. The DNS protocol is particularly susceptible to routing attacks, but it is also more resilient to denial of service attacks than many other Internet applications through the explicit use of multiple service points in DNS configurations. Performing DNSSEC validation on signed domain names can protect against routing hijacks that attempt to subvert DNS answers by responding to queries with false data.

Internet routing security is a combination of BGP protocol security, accuracy of routing policy, and robust operations. Internet routing registries are registries where network operators voluntarily provide information about the address prefixes they originate and propagate. They can be useful, but their utility degrades when they are not actively maintained, or attempt to take on too broad of a scope. The RPKI builds upon the earlier work of routing registries by associating digital signatures with some of the information found in Internet routing messages.

There is no universal solution to routing security that will function equally well for every organization and the network, or networks, for which it is responsible. Each organization must make its own decisions regarding how best to secure their routing infrastructure. Organizations should monitor their routes and be able to distinguish between normal and anomalous routing before an incident occurs. Both monitoring from inside the network looking out, and monitoring from outside looking in are important and provide distinct vantage points from which to gather intelligence. Whether an organization chooses to outsource their monitoring or not they are still responsible for their network's reputation. Every organization responsible for operating a network needs some access to routing expertise, and it is in an organization's best interest to assist their staff in forming friendly relationships with other networking engineers.

7 Acknowledgments, Statements of Interest, and Dissents, Alternative Views and Withdrawals

In the interest of transparency, these sections provide the reader with information about aspects of the SSAC process. The Acknowledgments section lists the SSAC members, outside experts, and ICANN staff who co-authored or contributed directly to this particular document (Contributors) or who provided reviews (Reviewers). The Statements of Interest section points to the biographies of all SSAC members and invited guests, which disclose any interests that might represent a conflict—real, apparent, or potential—with a member's or invited guest's participation in the preparation of this Report. The Dissents and Alternative Views sections provide a place for individuals to describe any disagreement with, or alternative view of, the content of this document or the process for preparing it. The Withdrawals section identifies individuals who have recused themselves from the discussion of the topic with which this Report is concerned. Except for members listed in the Dissents and Alternative Views or Withdrawals sections, this document has the consensus approval of all of the members of SSAC.

7.1 Acknowledgments

The committee wishes to thank the following SSAC members and invited guests for their time, contributions, and review in producing this report.

SSAC Members

Tim April
Patrik Fältström
Ondrej Filip
Cristian Hesselman
Geoff Huston
kc claffy
Warren Kumari
Barry Leiba
Danny McPherson
Russ Mundy
Rod Rasmussen
Mark Seiden

ICANN Staff

Andrew McConachie (editor)
Danielle Rutherford
Kathy Schnitt
Steve Sheng

7.2 Disclosures of Interest

SSAC member biographical information and Disclosures of Interest are available at:
<https://www.icann.org/resources/pages/ssac-biographies-2022-05-02-en>

7.3 Dissents and Alternative Views

There were no dissents or alternative views.

7.4 Withdrawals

There were no withdrawals.

Appendix A: Additional References

This Appendix contains additional references not included in the body of this publication. The SSAC would like to especially thank Ben Du and KC Claffy for providing many of these references. Links below not beginning with <https://doi.org/> have been archived in the Internet Archive's Wayback Machine at <https://archive.org/>.

- Apostolaki, Maria, Aviv Zohar, and Laurent Vanbever. "Hijacking Bitcoin: Routing Attacks on Cryptocurrencies." *ArXiv:1605.07524 [Cs]*, March 24, 2017. <http://arxiv.org/abs/1605.07524>.
- Ballani, Hitesh, Paul Francis, and Xinyang Zhang. "A Study of Prefix Hijacking and Interception in the Internet." *ACM SIGCOMM Computer Communication Review* 37, no. 4 (2007): 265–76. <https://doi.org/10.1145/1282427.1282411>.
- Battista, Giuseppe Di, Tiziana Refice, and Massimo Rimondini. "How to Extract BGP Peering Information from the Internet Routing Registry." In *Proceedings of the 2006 SIGCOMM Workshop on Mining Network Data - MineNet '06*, 317–22. Pisa, Italy: ACM Press, 2006. <https://doi.org/10.1145/1162678.1162685>.
- Borchert, Oliver, Kyehwan Lee, Kotikalapudi Sriram, Doug Montgomery, Patrick Gleichmann, and Mehmet Adalier. "BGP Secure Routing Extension (BGP-SRx): Reference Implementation and Test Tools for Emerging BGP Security Standards." National Institute of Standards and Technology, September 15, 2021. <https://doi.org/10.6028/NIST.TN.2060>.
- Brouwer, Marius, and Erik Dekker. "The Current State of DNS Resolvers and RPKI Protection," 2020, 8. <https://rp.os3.nl/2019-2020/p04/report.pdf>.
- Bush, Randy, James Hiebert, Olaf Maennel, Matthew Roughan, and Steve Uhlig. "Testing the Reachability of (New) Address Space." In *Proceedings of the 2007 SIGCOMM Workshop on Internet Network Management - INM '07*, 236. Kyoto, Japan: ACM Press, 2007. <https://doi.org/10.1145/1321753.1321756>.
- Chung, Taejoong, Emile Aben, Tim Bruijnzeels, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, et al. "RPKI Is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins." In *Proceedings of the Internet Measurement Conference*, 406–19. Amsterdam Netherlands: ACM, 2019. <https://doi.org/10.1145/3355369.3355596>.
- Cohen, Avichai, Yossi Gilad, Amir Herzberg, and Michael Schapira. "One Hop for RPKI, One Giant Leap for BGP Security." In *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, 1–7. Philadelphia PA USA: ACM, 2015. <https://doi.org/10.1145/2834050.2834078>.
- Cooper, Danny, Ethan Heilman, Kyle Brogle, Leonid Reyzin, and Sharon Goldberg. "On the Risk of Misbehaving RPKI Authorities." In *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*, 1–7. College Park Maryland: ACM, 2013. <https://doi.org/10.1145/2535771.2535787>.
- Del Fiore, Julian M., Pascal Merindol, Valerio Persico, Cristel Pelsser, and Antonio Pescape. "Filtering the Noise to Reveal Inter-Domain Lies." In *2019 Network Traffic Measurement and Analysis Conference (TMA)*, 17–24. Paris, France: IEEE, 2019. <https://doi.org/10.23919/TMA.2019.8784618>.

- Dhamdhere, Amogh, Renata Teixeira, Constantine Dovrolis, and Christophe Diot. “NetDiagnoser: Troubleshooting Network Unreachabilities Using End-to-End Probes and Routing Data.” In *Proceedings of the 2007 ACM CoNEXT Conference on - CoNEXT '07*, 1. New York, New York: ACM Press, 2007. <https://doi.org/10.1145/1364654.1364677>.
- Durand, Alain. “OCTO-014: Resource Public Key Infrastructure (RPKI) Technical Analysis.” ICANN Office of the Chief Technology Officer, September 2, 2020. <https://www.icann.org/en/system/files/files/octo-014-02sep20-en.pdf>.
- E-yong Kim, Li Xiao, K. Nahrstedt, and Kunsoo Park. “Secure Interdomain Routing Registry.” *IEEE Transactions on Information Forensics and Security* 3, no. 2 (June 2008): 304–16. <https://doi.org/10.1109/TIFS.2008.922050>.
- Fawcett, Milly. “Network Hijacking - the Low Down.” *The Spamhaus Project* (blog), January 8, 2018. <https://www.spamhaus.org/news/article/778/network-hijacking-the-low-down>.
- Feamster, Nick, Jaeyeon Jung, and Hari Balakrishnan. “An Empirical Study of ‘Bogon’ Route Advertisements.” *ACM SIGCOMM Computer Communication Review* 35, no. 1 (January 2005): 63–70. <https://doi.org/10.1145/1052812.1052826>.
- Feldmann, Anja, Olaf Maennel, Z. Morley Mao, Arthur Berger, and Bruce Maggs. “Locating Internet Routing Instabilities.” *ACM SIGCOMM Computer Communication Review* 34, no. 4 (August 30, 2004): 205–18. <https://doi.org/10.1145/1030194.1015491>.
- Gilad, Yossi, Avichai Cohen, Amir Herzberg, Michael Schapira, and Haya Shulman. “Are We There Yet? On RPKI’s Deployment and Security.” In *Proceedings 2017 Network and Distributed System Security Symposium*. San Diego, CA: Internet Society, 2017. <https://doi.org/10.14722/ndss.2017.23123>.
- Gilad, Yossi, Omar Sagga, and Sharon Goldberg. “MaxLength Considered Harmful to the RPKI.” In *Proceedings of the 13th International Conference on Emerging Networking Experiments and Technologies*, 101–7. Incheon Republic of Korea: ACM, 2017. <https://doi.org/10.1145/3143361.3143363>.
- Gill, Phillipa, Michael Schapira, and Sharon Goldberg. “Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security.” *ACM SIGCOMM Computer Communication Review* 41, no. 4 (August 15, 2011): 14–25. <https://doi.org/10.1145/2043164.2018439>.
- Goldberg, Sharon. “Why Is It Taking so Long to Secure Internet Routing?” *Communications of the ACM* 57, no. 10 (September 23, 2014): 56–63. <https://doi.org/10.1145/2659899>.
- Goodell, Geoffrey, William Aiello, Timothy Griffin, John Ioannidis, Patrick McDaniel, and Aviel Rubin. “Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing.” *NDSS* 23 (2003): 11. <https://www.ndss-symposium.org/wp-content/uploads/2017/09/Working-around-BGP-An-Incremental-Approach-to-Improving-Security-and-Accuracy-in-Interdomain-Routing-Geoffrey-Goodell.pdf>.
- Haag, William, Doug Montgomery, William C Barker, and Allen Tan. “Protecting the Integrity of Internet Routing:: Border Gateway Protocol (BGP) Route Origin Validation.” Gaithersburg, MD: National Institute of Standards and Technology, June 2019. <https://doi.org/10.6028/NIST.SP.1800-14>.
- He, Yihua, Georgos Siganos, Michalis Faloutsos, and Srikanth Krishnamurthy. “A Systematic Framework for Unearthing the Missing Links: Measurements and Impact.” In *Proceedings of the 4th USENIX Conference on Networked Systems Design & Implementation*, 14.

- NSDI'07. USA: USENIX Association, 2007.
https://www.usenix.org/legacy/event/nsdi07/tech/full_papers/he/he.pdf.
- Hu, Xin, and Z. Morley Mao. "Accurate Real-Time Identification of IP Prefix Hijacking." In *2007 IEEE Symposium on Security and Privacy (SP '07)*, 3–17, 2007.
<https://doi.org/10.1109/SP.2007.7>.
- Hu, Yih-Chun, Adrian Perrig, and Marvin Sirbu. "SPV: Secure Path Vector Routing for Securing BGP." In *Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, 179–92. SIGCOMM '04. New York, NY, USA: Association for Computing Machinery, 2004.
<https://doi.org/10.1145/1015467.1015488>.
- Huston, Geoff, Mattia Rossi, and Grenville Armitage. "Securing BGP — A Literature Survey." *IEEE Communications Surveys & Tutorials* 13, no. 2 (2011): 199–222.
<https://doi.org/10.1109/SURV.2011.041010.00041>.
- Jonker, Mattijs, Aiko Pras, Alberto Dainotti, and Anna Sperotto. "A First Joint Look at DoS Attacks and BGP Blackholing in the Wild." In *Proceedings of the Internet Measurement Conference 2018*, 457–63. Boston MA USA: ACM, 2018.
<https://doi.org/10.1145/3278532.3278571>.
- Karlin, Josh, Stephanie Forrest, and Jennifer Rexford. "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes." In *Proceedings of the 2006 IEEE International Conference on Network Protocols*, 290–99. Fess parker's Doubletree, Santa Barbara, Ca, USA: IEEE, 2006. <https://doi.org/10.1109/ICNP.2006.320179>.
- Katz-Bassett, Ethan, Harsha V. Madhyastha, John P. John, Arvind Krishnamurthy, David Wetherall, and Thomas Anderson. "Studying Black Holes in the Internet with Hubble." In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, 247–62. NSDI'08. USA: USENIX Association, 2008.
https://www.usenix.org/legacy/events/nsdi08/tech/full_papers/katz-bassett/katz-bassett.pdf.
- Kent, Stephen, Charles Lynn, Joanne Mikkelsen, and Karen Seo. "Secure Border Gateway Protocol (S-BGP) — Real World Performance and Deployment Issues." *NDSS*, 2000, 14.
<https://www.ndss-symposium.org/wp-content/uploads/2017/09/Secure-Border-Gateway-Protocol-S-BGP-Real-World-Performance-and-Deployment-Issues-paperStephen-Kent.pdf>.
- Khan, Akmal, Hyun-chul Kim, Taekyoung Kwon, and Yanghee Choi. "A Comparative Study on IP Prefixes and Their Origin Ases in BGP and the IRR." *ACM SIGCOMM Computer Communication Review* 43, no. 3 (July 2013): 16–24.
<https://doi.org/10.1145/2500098.2500101>.
- Konte, Maria, Roberto Perdisci, and Nick Feamster. "ASwatch: An AS Reputation System to Expose Bulletproof Hosting ASes." In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, 625–38. London United Kingdom: ACM, 2015.
<https://doi.org/10.1145/2785956.2787494>.
- Kotikalapudi, Sriram, and Azimov Alexander. "Methods for Detection and Mitigation of BGP Route Leaks." Internet Engineering Task Force, October 24, 2021.
<https://datatracker.ietf.org/doc/draft-ietf-grow-route-leak-detection-mitigation/07/>.
- Kuerbis, Brenden, and Milton Mueller. "Negotiating a New Governance Hierarchy: An Analysis of the Conflicting Incentives to Secure Internet Routing." *Communications & Strategies*, no. 81 (March 21, 2011): 125–42. <https://papers.ssrn.com/abstract=2021835>.

- Lad, Mohit, Dan Massey, Dan Pei, Yiguo Wu, Beichuan Zhang, and Lixia Zhang. "PHAS: A Prefix Hijack Alert System." In *Proceedings of the 15th Conference on USENIX Security Symposium - Volume 15*. USENIX-SS'06. USA: USENIX Association, 2006. https://www.usenix.org/legacy/event/sec06/tech/full_papers/lad/lad.pdf.
- Lad, Mohit, Ricardo Oliveira, Beichuan Zhang, and Lixia Zhang. "Understanding Resiliency of Internet Topology against Prefix Hijack Attacks." In *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07)*, 368–77. Edinburgh, UK: IEEE, 2007. <https://doi.org/10.1109/DSN.2007.95>.
- Li, Jun, Toby Ehrenkrantz, and Paul Elliott. "Buddyguard: A Buddy System for Fast and Reliable Detection of IP Prefix Anomalies." In *2012 20th IEEE International Conference on Network Protocols (ICNP)*, 1–10. Austin, TX, USA: IEEE, 2012. <https://doi.org/10.1109/ICNP.2012.6459962>.
- Li, Qi, Xinwen Zhang, Xin Zhang, and Purui Su. "Invalidating Idealized BGP Security Proposals and Countermeasures." *IEEE Transactions on Dependable and Secure Computing* 12, no. 3 (May 1, 2015): 298–311. <https://doi.org/10.1109/TDSC.2014.2345381>.
- Linssen, Raoul. "Vulnerability of DNS Name Servers against BGP Hijacking." *32th Twente Student Conference on IT*, January 19, 2019, 9. <https://www.utwente.nl/en/eemcs/dacs/assignments/completed/bachelor/reports/b-assignment-Raoul-linssen.pdf>.
- Luckie, Matthew. "Spurious Routes in Public BGP Data." *ACM SIGCOMM Computer Communication Review* 44, no. 3 (July 28, 2014): 14–21. <https://doi.org/10.1145/2656877.2656880>.
- Luckie, Matthew, Robert Beverly, Ryan Koga, Ken Keys, Joshua A. Kroll, and k claffy. "Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet." In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 465–80. London United Kingdom: ACM, 2019. <https://doi.org/10.1145/3319535.3354232>.
- Lychev, Robert, Sharon Goldberg, and Michael Schapira. "BGP Security in Partial Deployment: Is the Juice Worth the Squeeze?" In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, 171–82. SIGCOMM '13. New York, NY, USA: Association for Computing Machinery, 2013. <https://doi.org/10.1145/2486001.2486010>.
- Mahajan, Ratul, David Wetherall, and Tom Anderson. "Understanding BGP Misconfiguration." *ACM SIGCOMM Computer Communication Review* 32, no. 4 (August 19, 2002): 3–16. <https://doi.org/10.1145/964725.633027>.
- Meier, Roland, Petar Tsankov, Vincent Lenders, Laurent Vanbever, and Martin Vechev. "NetHide: Secure and Practical Network Topology Obfuscation." In *Proceedings of the 27th USENIX Conference on Security Symposium*, 693–709. SEC'18. USA: USENIX Association, 2018. <https://www.usenix.org/conference/usenixsecurity18/presentation/meier>.
- Mickens, James W., John R. Douceur, William J. Bolosky, and Brian D. Noble. "StrobeLight: Lightweight Availability Mapping and Anomaly Detection." In *Proceedings of the 2009 Conference on USENIX Annual Technical Conference*, 5. USENIX'09. USA: USENIX Association, 2009. https://www.usenix.org/legacy/event/usenix09/tech/full_papers/mickens/mickens.pdf.
- "Mutually Agreed Norms for Routing Security (MANRS)." Accessed May 20, 2022. <https://www.manrs.org/>.

- RPKI Documentation. “NLNet Labs RPKI Documentation.” Accessed May 20, 2022. <https://rpki.readthedocs.io/en/latest/index.html>.
- Oorschot, P.C. van, Tao Wan, and Evangelos Kranakis. “On Interdomain Routing Security and Pretty Secure BGP (PsBGP).” *ACM Transactions on Information and System Security* 10, no. 3 (July 2007): 11. <https://doi.org/10.1145/1266977.1266980>.
- Orsini, Chiara, Alistair King, Danilo Giordano, Vasileios Giotsas, and Alberto Dainotti. “BGPStream: A Software Framework for Live and Historical BGP Data Analysis.” In *Proceedings of the 2016 Internet Measurement Conference*, 429–44. Santa Monica California USA: ACM, 2016. <https://doi.org/10.1145/2987443.2987482>.
- “Qrator.Radar Blog of Route Leaks.” Accessed May 20, 2022. <https://radar.qrator.net/blog>.
- Ramachandran, Anirudh, and Nick Feamster. “Understanding the Network-Level Behavior of Spammers.” In *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, 291–302. SIGCOMM ’06. New York, NY, USA: Association for Computing Machinery, 2006. <https://doi.org/10.1145/1159913.1159947>.
- Robachevsky, Andrei. “14,000 Incidents: A 2017 Routing Security Year in Review.” *Mutually Agreed Norms for Routing Security (MANRS)* (blog), January 9, 2018. <https://www.internetsociety.org/blog/2018/01/14000-incidents-2017-routing-security-year-review>.
- Schlamp, Johann, Georg Carle, and Ernst W. Biersack. “A Forensic Case Study on AS Hijacking: The Attacker’s Perspective.” *ACM SIGCOMM Computer Communication Review* 43, no. 2 (April 29, 2013): 5–12. <https://doi.org/10.1145/2479957.2479959>.
- Schlamp, Johann, Ralph Holz, Quentin Jacquemart, Georg Carle, and Ernst W. Biersack. “HEAP: Reliable Assessment of BGP Hijacking Attacks.” *IEEE Journal on Selected Areas in Communications* 34, no. 6 (June 2016): 1849–61. <https://doi.org/10.1109/JSAC.2016.2558978>.
- Sermpezis, Pavlos, Vasileios Kotronis, Petros Gigis, Xenofontas Dimitropoulos, Danilo Cicalese, Alistair King, and Alberto Dainotti. “ARTEMIS: Neutralizing BGP Hijacking Within a Minute.” *IEEE/ACM Transactions on Networking* 26, no. 6 (December 2018): 2471–86. <https://doi.org/10.1109/TNET.2018.2869798>.
- Shue, Craig A., Andrew J. Kalafut, and Minaxi Gupta. “Abnormally Malicious Autonomous Systems and Their Internet Connectivity.” *IEEE/ACM Transactions on Networking* 20, no. 1 (February 2012): 220–30. <https://doi.org/10.1109/TNET.2011.2157699>.
- Song, Yang, Arun Venkataramani, and Lixin Gao. “Identifying and Addressing Reachability and Policy Attacks in ‘Secure’ BGP.” *IEEE/ACM Transactions on Networking* 24, no. 5 (October 2016): 2969–82. <https://doi.org/10.1109/TNET.2015.2503642>.
- Spamhaus. “Suspicious Network Resurrections.” *The Spamhaus Project* (blog), November 25, 2020. <https://www.spamhaus.org/news/article/802/suspicious-network-resurrections>.
- Steenbergen, Richard. “Examining the Validity of IRR Data.” Presented at the NANOG 44, Orlando, FL, USA, October 14, 2008. https://archive.nanog.org/meetings/nanog44/presentations/Tuesday/RAS_irrdata_N44.pdf.
- Subramanian, Lakshminarayanan, Volker Roth, Ion Stoica, Scott Shenker, and Randy H. Katz. “Listen and Whisper: Security Mechanisms for BGP.” In *Proceedings of the 1st Conference on Symposium on Networked Systems Design and Implementation - Volume 1*, 10. NSDI’04. USA: USENIX Association, 2004.

- <https://www.usenix.org/legacy/publications/library/proceedings/nsdi04/tech/subramanianListen/subramanianListen.pdf>.
- Testart, Cecilia. “Reviewing a Historical Internet Vulnerability: Why Isn’t BGP More Secure and What Can We Do About It?” *SSRN Electronic Journal*, 2018. <https://doi.org/10.2139/ssrn.3141666>.
- Testart, Cecilia, Philipp Richter, Alistair King, Alberto Dainotti, and David Clark. “To Filter or Not to Filter: Measuring the Benefits of Registering in the RPKI Today.” In *Passive and Active Measurement*, edited by Anna Sperotto, Alberto Dainotti, and Burkhard Stiller, 12048:71–87. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2020. https://doi.org/10.1007/978-3-030-44081-7_5.
- Vervier, Pierre-Antoine, Olivier Thonnard, and Marc Dacier. “Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks.” In *Proceedings 2015 Network and Distributed System Security Symposium*. San Diego, CA: Internet Society, 2015. <https://doi.org/10.14722/ndss.2015.23035>.
- Wahlisch, Matthias, Robert Schmidt, Thomas C. Schmidt, Olaf Maennel, Steve Uhlig, and Gareth Tyson. “RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem.” *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, November 16, 2015, 1–7. <https://doi.org/10.1145/2834050.2834102>.
- Wang, Feng, and Lixin Gao. “On Inferring and Characterizing Internet Routing Policies.” In *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement*, 15–26. IMC ’03. New York, NY, USA: Association for Computing Machinery, 2003. <https://doi.org/10.1145/948205.948208>.
- Xiang, Yang, Zhiliang Wang, Xia Yin, and Jianping Wu. “Argus: An Accurate and Agile System to Detecting IP Prefix Hijacking.” In *2011 19th IEEE International Conference on Network Protocols*, 43–48. Vancouver, AB, Canada: IEEE, 2011. <https://doi.org/10.1109/ICNP.2011.6089080>.
- Yoo, Christopher S., and David Wishnick. “Lowering Legal Barriers to RPKI Adoption.” *SSRN Electronic Journal*, 2019. <https://doi.org/10.2139/ssrn.3308619>.
- Zhang, Ming, Chi Zhang, Vivek Pai, Larry Peterson, and Randy Wang. “PlanetSeer: Internet Path Failure Monitoring and Characterization in Wide-Area Services.” *OSDI 4 (2004)*: 16. https://www.usenix.org/legacy/events/osdi04/tech/full_papers/zhang/zhang.pdf.
- Zhang, Zheng, Ying Zhang, Y. Charlie Hu, Z. Morley Mao, and Randy Bush. “Ispy: Detecting Ip Prefix Hijacking on My Own.” In *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication*, 327–38. SIGCOMM ’08. New York, NY, USA: Association for Computing Machinery, 2008. <https://doi.org/10.1145/1402958.1402996>.
- Zheng, Changxi, Lusheng Ji, Dan Pei, Jia Wang, and Paul Francis. “A Light-Weight Distributed Scheme for Detecting Ip Prefix Hijacks in Real-Time.” *ACM SIGCOMM Computer Communication Review* 37, no. 4 (October 2007): 277–88. <https://doi.org/10.1145/1282427.1282412>.

Appendix B: RPKI

When routing security is discussed, the first tool reference is often Resource Public Key Infrastructure or RPKI. RPKI was introduced in 2012 with its architecture being described in RFC6480 and subsequent components being covered in RFC6481, RFC6482, RFC6483, RFC6484, RFC6485, RFC6486, RFC6487, RFC6488, RFC6489, RFC6490, RFC6491, RFC6492, and RFC6493.

In short, RPKI is a way for entities who are assigned IP Address resources from a regional Internet registry (RIR) to assert which autonomous systems (AS) or systems are allowed to originate a prefix. These assertions are then cryptographically signed by the delegated RIR, using a key the RIR publishes in a defined format called a route origination authorization (ROA). These ROAs are then published by the RIR in a database they operate where operators may look up ROAs in order to validate announcements they receive. ROAs also allow resource holders to limit the maximum prefix length for their announcements if they choose to. When not specified, announcements are restricted to the prefix length published in the ROA.

RPKI ROAs only specify the AS allowed to originate the prefix, the origin AS, and have no way of specifying the path by which traffic should reach that origin. As a result of this limitation, an attacker who intends to propagate a malicious announcement can append the authorized origin AS to their AS and announce that route to their neighbors, potentially allowing them to intercept traffic which would traverse their network.

While not preventing a determined adversary, RPKI provides some protections from many common sources of routing incidents, such as an incorrectly entered prefix, or a route server incorrectly advertising a prefix to neighboring ASes.

In addition to the resource holders needing to generate correct ROAs for their resources, other operators of BGP speakers also need to implement RPKI validation within their networks and to drop any routes which are not validated according to the defined process. In recent years, the adoption of RPKI validation has started to increase, partially driven by efforts of some network operators, in addition to initiatives like Mutually Agreed Norms for Routing Security (MANRS).⁶⁶

⁶⁶ See “MANRS – Mutually Agreed Norms for Routing Security,” May 11, 2022. <https://www.manrs.org/>.

Appendix C: BGPsec

The BGPsec protocol is an extension to BGP that is intended to allow validation of the AS Path attribute.⁶⁷

Unlike route origin validation (ROV), BGPsec is not implemented in an off-router mode, but is implemented through the definition of non-transitive BGP AS Path attributes. These attributes carry the digital signatures produced by the AS that propagates a BGP UPDATE message. These signatures, signed by the AS, provide confidence that every AS listed in the AS Path attribute has handled the propagation of this prefix, that the order in the AS Path is the exact order of propagation of the UPDATE message through the inter-domain routing space, and each AS listed has explicitly authorized the propagation of an UPDATE message to its eBGP peer.

BGPsec appears to be solidly based on the concepts first described in earlier sBGP work.⁶⁸ Each eBGP speaker generates a digital signature that covers the information it received (including that digital signature) and the AS number to whom this UPDATE is to be sent. There is a wealth of detail behind this simple explanation, but it can be summarized by the observation that this mechanism ties the network prefix to the AS Path in the UPDATE message. The IETF Secure Inter-Domain Routing (SIDR) working group provided a detailed exposition of BGPsec's design decisions.⁶⁹

Stepwise AS Path validation cannot tolerate AS Sets in this approach, nor AS Confederation Sets, and are in the process of being deprecated in response to this limitation.^{70,71} In a similar vein BGP Route Reflectors require special processing, as do private AS numbers.

There are a number of consequences of this design approach.

The first, and perhaps the most important consequence, is that piecemeal incremental deployment is simply not possible in BGPsec. When an UPDATE is passed from a BGPsec BGP speaker to a non-BGPsec BGP speaker all BGPsec attributes are lost. Thus, if that UPDATE is further propagated to a BGPsec BGP speaker the initial BGPsec information is unavailable. In today's Internet the consequences of this highly constrained deployment scenario are significant impediments to widespread adoption.

This approach also places a high cryptographic processing load on BGPsec-aware BGP speakers. Some skepticism that this is a feasible proposal for the Internet's routing infrastructure guided the design of the RPKI to router protocol.⁷² However, for BGPsec not only are routers expected to

⁶⁷ See RFC 8205: BGPsec Protocol Specification

⁶⁸ See RFC 4301: Security Architecture for the Internet Protocol

⁶⁹ See RFC 8374: BGPsec Design Choices and Summary of Supporting Discussions

⁷⁰ See RFC 6472: Recommendation for Not Using AS_SET and AS_CONFED_SET in BGP

⁷¹ See draft-ietf-idr-deprecate-as-set-confed-set: Deprecation of AS_SET and AS_CONFED_SET in BGP, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-deprecate-as-set-confed-set-07>

⁷² See RFC 6810: The Resource Public Key Infrastructure (RPKI) to Router Protocol

SSAC Briefing on Routing Security

process the BGPsec messages, but also hold secure private keys to perform signing on the fly for outgoing UPDATE messages.

Thirdly, while this approach can provide some assurance regarding the "correct" operation of the BGP protocol and can detect efforts to tamper with update messages, there is no protection against spurious WITHDRAW messages, no ability to ascertain the alignment of the route object with the network's forwarding state and no protection of alignment of the UPDATE with the policy state. In other words, route leaks can still occur in BGPsec.

In summary, BGPsec represents a relatively high overhead to pay for a limited set of assurances and a limited protective capability. Furthermore, there is a more extreme view that BGPsec cannot achieve any of the security properties due to the fundamental design principles of BGP and BGPsec.