How Effective is Multiple-Vantage-Point Domain Control Validation?

Grace H. Cimaszewski[†] * Henry Birge-Lee[†] * Liang Wang[†] Jennifer Rexford[†]
Prateek Mittal[†]
[†] *Princeton University*

Abstract

Multiple-vantage-point domain control validation (multiVA) is an emerging defense for mitigating BGP hijacks against Web PKI certificate authorities. While adoption of multiVA is growing, little work has quantified its effectiveness against BGP hijacks in the wild. We bridge the gap by presenting the first analysis framework that measures the security of multiVA deployment under a confluence of real-world routing and networking practices (namely, DNS and RPKI). Our framework accurately models the attack surface of multiVA by 1) considering attacks on DNS nameservers involved in domain validation, 2) incorporating deployed practical security techniques such as RPKI, 3) performing fine-grained Internet-scale analysis to compute resilience (i.e., how difficult it is to launch a BGP hijack against a domain and get a bogus certificate under multiVA). We apply our framework to perform a rigorous security analysis of the multiVA deployment of Let's Encrypt, compiling a dataset of 31 billion DNS queries for about 1 million domains over the course of four months. Our analysis shows that while DNS does enlarge the attack surface of multiVA, Let's Encrypt's multiVA deployment still offers an 88% median resilience against BGP hijacks, a notable improvement over the 76% resilience offered by single-vantage-point validation. RPKI, even in its current state of partial deployment, effectively mitigates BGP attacks and improves security of the deployment by 15%. Exploring over 11,000 different multiVA configurations, we find that Let's Encrypt's deployment can be further expanded to achieve a resilience of over 97% with only two additional vantage points in different public cloud providers. In addition to adding these vantage points, moving to a full quorum policy can achieve a maximal resilience of over 99%, motivating a rethinking of multiVA design parameters.

1 Introduction

Certificate Authorities (also known as Certification Authorities or CAs) serve as root of trust for secure TLS communications by signing digital certificates that tie a notion of a server's identity (like a domain name) to its public key. However, the process that CAs use to verify domain ownership, known as domain control validation, is vulnerable to BGP attacks [8, 12]. By hijacking traffic to a victim's domain during domain control validation, an adversary can fool a CA into signing a certificate on its behalf for a domain it does not control [8]. The first real-world instances of these attacks were observed in 2022 when attackers used them to steal millions' of dollars worth of cryptocurrency [7, 29]; millions of other websites (including those observing best security practices) may also be vulnerable [7]. This presents a potentially devastating attack in the hands of repressive regimes that have been accused of launching strategic BGP attacks in the past [20].

To mitigate the risk of BGP attacks on domain control validation, several CAs (including the world's largest Web PKI CA Let's Encrypt and Google Trust Services) have deployed a countermeasure known as multiVA [9, 31].[1] In multiVA, domain control validation for each certificate request is performed multiple times from distinct vantage points spread throughout the Internet. This technique detects BGP attacks by exploiting the fact that many BGP announcements are localized and only affect a portion of Internet traffic [9]. If a CA performs domain control validation from a vantage point whose routing is unaffected by the adversary's attack, this vantage point's traffic will be routed to the victim's server and connect to the "real" domain, enabling detection of the attack.

While this countermeasure is sound in principle (i.e., more remote vantage points with greater geographic spread have more potential to catch localized BGP attacks [8, 9]), a key question is: *how effective is multiVA in practice?* For mul-

*Both authors contributed equally to this work.

[1] Consistent with Birge-Lee *et al.* [9], we use the term multiVA to refer to this technology, which has also been called Multiple-Vantage-Point Domain Validation [8], Multi-Path Domain Validation [33], and Multi-Perspective Domain Validation [1] in some prior works.

tiVA to be effective, enough vantage points must route to the victim during a BGP attack so that the CA can detect the attack. There is currently an active debate over multiVA's effectiveness to strengthen domain control validation against BGP attacks at the CA/Browser Forum (the governing body for CAs) [55]. An inherent challenge lies in the fact that strategic BGP attacks can target the weakest link in multiVA deployments, including any vulnerable IP prefixes involved in the domain validation process. Thus, it is not obvious that a multiVA deployment constrained by financial and operational costs can significantly reduce the BGP attack surface [13]. MultiVA's effectiveness in light of real-world domain hosting practices and routing dynamics is a critical open research question for the security community [13, 17] and our work aims to inform this debate via a rigorous quantitative study.

There are several components of the Internet and routing ecosystem that impact the effectiveness of multiVA, including the Domain Name System (DNS) [13, 17] and the ongoing deployment of the Resource Public Key Infrastructure (RPKI). While previous work [13] has only begun to address some of these factors (as we discuss in § 9), we present the first work to rigorously compute the resilience of multiVA against BGP attacks under real-world network and DNS dynamics.

Contributions. In this paper, we contribute a novel analysis framework that simulates multiVA under real-world conditions to rigorously understand its effectiveness in the wild. Notably, our work models the DNS infrastructure of over one million of Let's Encrypt's customer domains and incorporates existing RPKI deployment via Internet-scale topology simulations. We also present the first analysis of vantage points deployed across different cloud providers.

1. Using log data from the world's largest Web PKI CA Let's Encrypt, we perform an extensive study of the DNS configurations of over a million customer domains using **31 billion DNS queries** from geographically-distributed vantage points. We develop a custom full-graph DNS resolver to capture factors that impact the resilience of domains to BGP attacks, including IP prefixes of associated DNS infrastructure and the deployment of DNSSEC and RPKI.
2. We perform fine-grained Internet topology simulations under both current and hypothetical RPKI conditions using CAIDA Internet topology data [14], BGP data collected by public route monitors, and traceroute data from cloud datacenters to understand the routing of a diverse set of vantage points during attacks from 1 K sampled adversaries against all 800 K IP prefixes on the Internet (which required four hundred million full topology simulations containing 24 trillion simulated routes).
3. We complete a quantitative multi-faceted security evaluation to measure the impact of factors in the routing and DNS ecosystems on over 11 K different hypothetical multiVA deployments, including deployments using multiple cloud providers and Let's Encrypt's existing deployment. By computing effective resilience[2] for Let's Encrypt customer domains under different conditions, we understand how to optimize the security of a multiVA deployment.

The results of this analysis (also summarized in Table 1) reveal several insights to inform multiVA deployments:

1. Considering DNS as an attack surface multiplies the number of IP prefixes an adversary can target with a BGP attack by a factor of five. As only 5% of domains are DNSSEC-signed, BGP attacks on DNS resolution are highly viable.
2. 60% of target IPs are covered by RPKI Route Origin Authorizations (ROAs) which helps to neutralize the attack surface introduced by DNS.
3. Considering both the DNS and the RPKI infrastructures, the resilience of the median customer domain under Let's Encrypt's current multiVA deployment is still 88.6% (only 6% less than when both of these factors are omitted).
4. Let's Encrypt's deployment can be further strengthened by adding additional cross-cloud provider vantage points or by switching to a full quorum policy which requires validation consensus across all vantage points. These changes can boost the resilience of the median domain to above 98.6% or 96.5% respectively. Both of these changes together can yield a 99.3% median domain resilience.

Our findings present critical empirical evidence for promoting the adoption of multiVA. Furthermore, our work has already enhanced the security of the Web PKI ecosystem: per this work's recommendations, Let's Encrypt plans to deploy an additional remote vantage point in northern Europe (Stockholm) to augment resilience [2]. The experimental results presented in this work have also been cited in discussions at the CA/Browser Forum as evidence to support the push to require multiVA in the issuance of domain validated-certificates at all CAs [22]. We have open sourced our framework to facilitate CAs' multiVA deployments and future research.

Ethical considerations. Our analysis uses certificate issuance logs provided by Let's Encrypt that have been sanitized of private client information. All the information used from these logs is publicly available in certificate transparency logs. In our measurement, we rate limit our queries to avoid overwhelming public DNS servers.

2 Background

In this section we start with a brief background on interdomain routing insecurity and then provide an overview of the DNS and how it is susceptible to routing attacks.

[2]The fraction of ASes on the Internet that cannot use an equally-specific BGP attack to obtain a certificate for a given domain [8, 9].

Significant Findings	Section
Importance of DNS and RPKI	
- Considering DNS causes a five-fold increase in the number of prefixes an adversary can target	§ 5.4
- Only 5.6% of domains are fully protected by DNSSEC	§ 5.4
- 60% of target IPs associated with domains are covered by ROAs and benefit from RPKI	§ 5.4
Effect on Domain Resilience	
- Considering the DNS attack surface drops the resilience of domains by 20% (from 95% to 75%) under multiple vantage point validation but 40%+ (from 83% to 42%) under single vantage point validation	§ 7.1
- 76.3% of domains have some RPKI coverage of their DNS resolution graphs; current RPKI coverage improves the resilience of domains by 15% to 90% and full coverage could improve the resilience by 20% to 95%	§ 7.1
- The resilience of Let's Encrypt's current deployment is still 88.6% considering these factors	§ 7.1
Ways to Improve the Deployment	
- Adding a single vantage point in Let's Encrypt's current cloud provider improves resilience to 93.2%	§ 7.2
- If two vantage points are added, a cross-cloud strategy is optimal and can boost resilience to 97.5%	§ 7.2
- Implementing a full quorum policy even with existing vantage points can boost resilience to 96.5%	§ 7.2

Table 1: Significant results from our analysis framework.

2.1 Routing System: BGP Attacks

Interdomain routes between Autonomous Systems (ASes) are negotiated via the Border Gateway Protocol (BGP). When an AS announces its IP prefix via BGP and lists its Autonomous System Number (ASN) as the origin of that BGP announcement, neighboring ASes propagate that announcement and append their own ASNs to it, providing a list of ASes on the path to reach the origin AS in each announcement. However, BGP route messages are unsigned, unauthenticated, and thus vulnerable to hijacks. The stealthiness and financial damages of BGP hijacks are on the rise, as seen in the repeated strategic use of BGP attacks to steal cryptocurrency [7, 29].

Equally-specific prefix hijacks. One of the simplest types of BGP attacks is an equally-specific prefix hijack where an adversary makes a BGP announcement containing the same prefix belonging to a victim AS, in effect claiming ownership of the victim's IP prefix. In this attack, affected ASes believe this malicious announcement and route their traffic destined for the victim to the adversary instead. When undetected, this is largely effective at attracting traffic.

Sub-prefix hijacks. In sub-prefix or more-specific prefix hijacks, the adversary announces a sub-prefix (i.e., a longer and more preferred prefix) of the victim's IP prefix. Because of longest-prefix-match forwarding, the adversary's sub-prefix route is preferred over the victim's route, often enabling these attacks to affect the entire Internet. While sub-prefix attacks are highly effective, they are not always viable as ASes typically filter BGP announcements for prefixes longer than 24 bits (or 48 bits in IPv6). Thus, 24-bit IPv4 network prefixes usually cannot be attacked with sub-prefix hijacks.

RPKI makes BGP hijacks more difficult. Resource Public Key Infrastructure (RPKI) [38] cryptographically attests to which ASes own which IP prefixes. Thus, when adversaries illegitimately try to originate an IP prefix which was not allocated to them, RPKI-based filtering (known as Route Origin Validation or ROV) can block these announcements and prevent them from propagating. Additionally, RPKI also specifies the prefix length allowed in BGP announcements for that prefix. This enables ASes performing ROV to filter sub-prefix attacks regardless of the ASN listed as the origin.

While RPKI is a major improvement to routing security, it is not a panacea because it only validates the origin AS of a BGP announcement, not the other ASes that claim to forward traffic to the origin. This allows an adversary to claim a potentially non-existent route to the true origin AS in a BGP announcement. We refer to such malicious announcements as equally-specific-prefix prepend attacks. These attacks can evade ROV but tend to affect a smaller portion of the Internet, as this strategy makes the adversary's announcement less preferable in BGP route selection.

2.2 DNS Name Resolution

When a web client connects to a domain, the domain name must first be resolved to an IP address. This is done using the Domain Name System (DNS). The DNS is hierarchically composed of delegations of zones between nameservers. To resolve a domain, the client first queries the root nameserver, then follows the graph of nameserver delegations to reach the authoritative nameserver(s) for that domain, which provide the actual record containing the domain name's IP address.

Most DNS queries are sent via unencrypted, unauthenticated UDP packets, making them a target for network attacks. In addition to off-path vulnerabilities like packet fragmentation-based attacks [52], DNS is vulnerable to BGP attacks which allow adversaries to answer DNS queries with malicious records that can point users to adversary-controlled servers instead of a victim domain [11].

DNSSEC adds cryptographic protections to DNS. DNSSEC is a DNS extension that requires all records from authoritative nameservers to be cryptographically signed. This

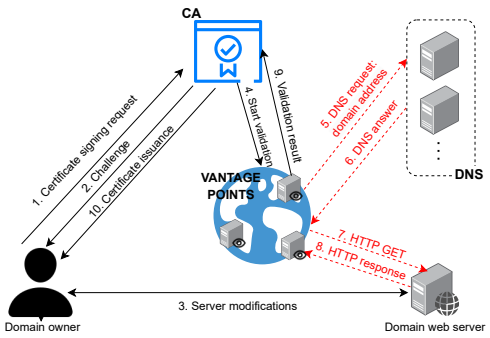

Figure 1: Certificate issuance process under multiVA. (Red text and lines indicate connections vulnerable to BGP hijack.)

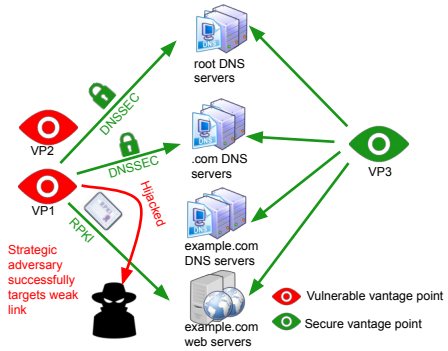

Figure 2: Example of vantage points and a domain vulnerable to an adversary's BGP hijack. The adversary strategically hijacks the most vulnerable link in light of countermeasures like RPKI and DNSSEC. To obtain a certificate for the target domain, the adversary must hijack enough vantage points to satisfy the CA's quorum policy (in our example, 2 of 3).

prevents attacks on DNS because even if an adversary compromises a victim's DNS query, it cannot generate a valid signed response without the victim's private zone-signing key. While many CAs validate DNSSEC when resolving domain names, not all domains register DNSSEC records.

3 Adversary Model

Our threat model extends the original multiVA threat model [9] to include potential BGP hijacks on DNS resolution. The adversary aims to obtain a certificate for a domain it does not control, which can enable more sophisticated attacks such as domain impersonation or man-in-the-middle attacks on encrypted TLS traffic between the victim domain and domain visitors.

Trusted system components. We assume that the CA and all of its remote vantage points are trusted. The CA is not acting maliciously and has complete control over its remote vantage points. We do not consider non-BGP attacks in our threat model, such as attacks that exploit off-path vulnerabilities in the DNS protocol [12, 30]. We also do not consider vulnerabilities introduced by bugs/misconfigurations in the software run by the CA or its vantage points. Lastly, we assume that DNSSEC and RPKI provide their stated security properties.

Adversary attack strategy. We consider an adversary that tries to use a BGP attack to fool the domain control validation process. Because domain control validation is used to bootstrap trust for the first time, domain control validation must be performed over unauthenticated channels. By launching a BGP attack against a victim domain that affects domain control validation traffic from a CA and sufficiently many remote vantage points, an adversary can pose as the victim domain and obtain a certificate (see Figure 2).

BGP capabilities. We consider an adversary with control of a single AS that can make malicious BGP announcements for any prefix(es) it chooses to target. We consider that the adversary may launch the following two types of BGP attacks:

1. **Equally-specific-prefix attack:** an adversary announces an equal-length prefix as the victim domain's prefix.
2. **Equally-specific-prefix prepend attack:** an adversary claims reachability to the victim's prefix via a non-existent connection by inserting itself on a valid path. This attack attempts to circumvent RPKI-ROV detection by keeping the valid originating AS as the route announcement origin.

We focus our analysis on equally-specific prefix attacks and largely omit sub-prefix hijacks from our analysis to remain faithful to the primary threat model motivating the design of multiVA [8, 9]. Sub-prefix hijacks are usually global in scope, violating the underlying principle of multiVA which involves detecting attacks by finding unaffected portions of the Internet. However, the wide spread of sub-prefix attacks and the clear signal that occurs when a new prefix enters the global routing system render these attacks highly detectable. They can be effectively detected through BGP monitoring, which can be implemented by either CAs [8] or domain owners/hosting providers [6, 35, 50]. Furthermore, sub-prefix hijacks can be mitigated by ISPs through the implementation of RPKI ROAs. Prominent ISPs such as Amazon Web Services (AWS) [32] and other top ISPs [15] filter RPKI-invalid BGP announcements. Another strategy to mitigate sub-prefix attacks is running security-critical services on 24-bit prefixes [8, 42, 53] as many networks (e.g., AWS) filter BGP prefixes longer than 24 bits [4, 16, 45], making prefixes longer than 24 bits invisible on a global scale [5]. While we do not position multiVA as a defense against sub-prefix hijacks, we do perform an analysis of the viability of sub-prefix attacks against multiVA in § 8.

Incorporating the DNS attack surface. We extend the original multiVA threat model to consider that an adversary can attack not only the domain server but also the DNS infrastructure associated with a victim's domain, using BGP attacks (see

Figure 2). When a CA performs domain control validation, the CA must resolve the domain to an IP address via DNS before contacting the domain to complete validation. If an adversary reroutes a DNS query with a BGP hijack, the adversary can generate a spurious DNS response that directs the CA to an adversary-controlled server (instead of the victim domain) and fraudulently complete validation. This significantly expands the multiVA attack surface as the many DNS nameservers are potential targets of BGP attacks.

We consider all authoritative nameservers involved in the DNS lookup of a victim's domain potential targets with the exception of those protected by DNSSEC, because the adversary cannot forge DNSSEC-signed DNS responses with BGP hijacks. We conservatively assume that an adversary:

- can retry validation several times to allow it to fool validation even if it can only compromise one of several A record or nameserver IP addresses (e.g., Let's Encrypt's validation rate limits are per-account and can easily be bypassed by creating multiple accounts [21]).
- can launch BGP attacks for a duration longer than the time period for DNS caching (e.g., longer than the 60 second DNS caching used by Let's Encrypt [17] as is the case with many real-world BGP attacks [7, 27, 29]).

Following from these assumptions, we consider that the attacker can successfully hijack traffic between the CA and domain if any IP address contacted during resolution (starting from DNS root) of a domain is vulnerable.

Let's Encrypt uses locally-run DNS resolvers at each vantage point [17], which removes the risk of BGP attacks on DNS queries between Let's Encrypt and its local recursive resolver. Our threat model also excludes off-path attacks on DNS like transaction id guessing [30] and packet-fragmentation attacks [52] as they can be prevented with best DNS operational practices [30, 52].

4 Analysis Framework

Prior work on the security of multiVA does not, or only to a limited extent, consider DNS as a potential attack vector [8, 9, 13]. Our study aims to understand the security of multiVA under a more realistic setting, which not only considers additional DNS-based attack vectors, but also the relevant deployed security measures (e.g., RPKI and DNSSEC) and operational practices. The former may allow more attack strategies and degrade the security of multiVA, while the latter could improve it. Toward this goal, we design a new analysis framework to fill the gaps in prior work. We develop our framework as a general-purpose, automated tool capable of modeling arbitrary deployment configurations, so that a CA can use it to evaluate the security of its multiVA deployment and explore potential enhancements to it.

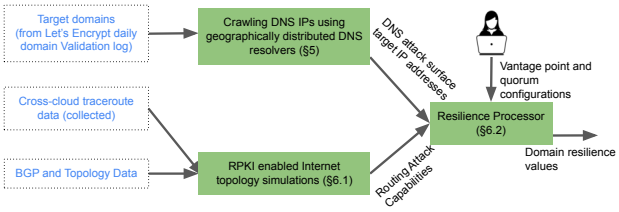

Figure 3: An overview of the analysis framework.

The framework takes as input a set of domain names and configuration of a multiVA deployment, and outputs the domain's resilience. A domain's resilience measures the fraction of the Internet from which that domain is *immune* to equally-specific BGP attacks. ASes in this fraction of the Internet cannot fool enough CA vantage points with an equally-specific BGP attack for the CA to potentially sign a malicious certificate for that domain. We measure the impact of various multiVA configurations by seeing how the resilience of customer domains signed by the CA changes. We defer the full definition of resilience to § 6.2.

As shown in Figure 3, our framework consists of three major components: (1) the geographically-distributed DNS resolvers, (2) the Internet topology simulator, and (3) the resilience processor.

Geographically-distributed DNS resolvers. To get a more complete view of the potential DNS attack surface, the *geographically-distributed DNS resolvers* resolve the input domains (extracted from Let's Encrypt's certificate issuance logs) in near-real-time, and trace the dependencies of the resolved domains (i.e., full-graph DNS resolution). Over the course of our experiment (which studied approximately 1.4 million domains) we sent 30.7 billion DNS requests and recorded 520 billion nameserver IP addresses and 30.6 billion A record IP addresses. We pair this data with BGP hijack simulations to understand which IP prefixes - and by extension, domains - are vulnerable to BGP attacks. The domain resolution routine is executed each time a daily issuance log is received from Let's Encrypt to allow for domains to be resolved within 24 hours, reducing the impact of domains being taken down or transient DNS configurations.

Fine-grained topology and accurate BGP attacks simulation incorporating RPKI. To produce accurate BGP attack simulations, the *Internet topology simulator* performs fine-grained, prefix-level (instead of AS-level) routing simulation.

To better model the routing decisions of datacenters being considered for vantage points, we needed to capture their BGP connectivity. We used the bdrmap tool [40] to infer peer lists for 19 unique datacenters spread across three cloud providers: AWS, Microsoft Azure, and Google Cloud Platform (GCP) (see Appendix Table 6 for full list). This enabled us to simulate the datacenters' routing in the face of BGP hijacks.

Leveraging this traceroute data, public topology informa-

tion [14], and BGP RIB dumps [43,48], the simulation engine simulates the interdomain routing of roughly two hundred thousand groups of IP prefixes, and further simulates equally-specific prefix BGP attacks against these prefixes launched from 1,000 random ASes. Different from prior BGP attack simulations [9], our simulation engine takes deployed BGP security practices, i.e., RPKI, into account, and can simulate attacks against RPKI-protected prefixes. These simulations are repeated under both RPKI and non-RPKI conditions resulting in roughly four hundred million simulations and over twenty-four trillion AS-level paths calculated.

Multifaceted quantitative security estimation. Finally, using the DNS data collected by the measurement engine and the results produced by the simulation engine, the *resilience processor* estimates the security of a multiVA deployment against routing attacks, which could target any vulnerable DNS servers found by the geographically-distributed DNS resolvers, under various scenarios.

Using the security estimation result as an indicator, our research addresses the following real-world deployment questions for CAs to maximize their resilience to BGP hijacks:

- **Selection of vantage point location.** What are the optimal vantage point locations that make the CA most secure against BGP attacks?
- **Measuring impact of quorum policy.** Birge-Lee *et al.* [9] observed that a stricter quorum policy may have negative operational implications for a CA's issuance (e.g., benign failures). By what extent do stricter quorum policy configurations improve resilience against BGP hijacks?
- **Comparison of cloud provider routing resiliency.** Does validation using vantage points hosted in multiple cloud providers provide additional security benefits beyond concentrating VPs in one cloud provider?

5 Measuring the DNS Attack Surface

In order to calculate domain resilience, we develop an experimental system to DNS resolve domains in certificates signed by Let's Encrypt. The design of our resolver tool was driven by our primary research motivation: **to quantify the extent of the DNS attack surface that may be targeted by interdomain routing attacks**. Our tool collects comprehensive details of a domain's DNS lookup graph, *to record every IP address that may be contacted in the resolution of that name.*

5.1 Defining the DNS Attack Surface

We denote the *full-graph DNS attack surface* of a domain d as $Q(d)$, which is the set of all IP addresses that may be the targets of BGP hijacks for d. Intuitively, $Q(d)$ includes all IP addresses of the webservers that host d, as well as all IP addresses that may be contacted to resolve the DNS zone that

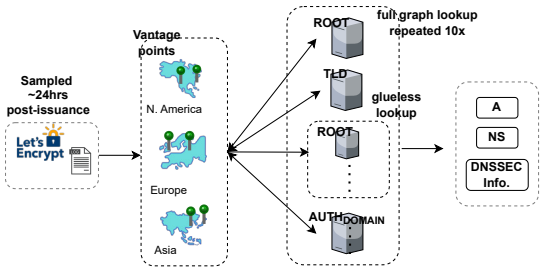

Figure 4: Design of daily DNS resolution framework.

hosts d's DNS records. To illustrate our mathematical formulation, we will refer to `example.com` as a running example.

Define the DNS zone delegation chain of d as $Z_0, Z_1, \ldots, Z_{k_d}$, where Z_0 represents the root zone and Z_{k_d} the authoritative zone for d's DNS records. Each zone Z_i is hosted by a set of nameservers N_{Z_i} ($n \in N_{Z_i}$ indicates a nameserver name). Lastly, denote the set of IP addresses contained in the A record of a name d as $A(d)$. To resolve domain d's IP address $A(d)$, a DNS resolver iteratively queries nameservers in each successive zone, receiving either a delegation to the next zone or an answer (or error) from the authoritative nameserver(s) $N_{Z_{k_d}}$. We use $I(Z)$ to denote the IP addresses of the nameservers in zone Z, as well as any IPs that may be queried to resolve the namerserver names themselves.

For the case of d = `example.com`, `example.com` is delegated across DNS zones Z_0 = ., Z_1 = .com, and Z_2 = example.com. N_{Z_0} consists of the 13 IANA-defined root nameservers; N_{Z_1} consists of the .com TLD authoritative nameservers; N_{Z_2} = {`ns1.provider.net`, `ns2.provider.net`} consists of the nameservers authoritative for example.com.

We formalize the attack surface of d, $Q(d)$, as follows:

$$Q(d) = I(Z_{k_d}) \cup A(d)$$

The DNS resolver can learn the IP address of d's nameservers by two methods: either by sending a separate DNS request for the nameserver's A records with its own associated $Q(n)$, or using the "glue" records provided in the additional section of the parent zone's response (if included). We refine the expression of Q for this case and call this Q':

$$Q'(n) = \begin{cases} A(n) & \text{if glued record exists} \\ Q(n) & \text{otherwise ("glueless")} \end{cases}$$

The DNS attack surface of d's authoritative zone Z_{k_d} can be defined recursively:

$$I(Z_{k_d}) = I(Z_{k_d-1}) \cup \bigcup_{n \in N_{Z_{k_d}}} Q'(n)$$

Applying this algorithm to compute Q(`example.com`), we must resolve the A records of `example.com` as well as

enumerate all IP addresses needed to resolve its authoritative nameservers {ns1.provider.net, ns2.provider.net}. The resolver first sends a DNS query to a root nameserver, which responds with a delegation to the .com nameservers in the authority section and a list of these nameserver IPs in the additional section. Next, the query is sent to the .com nameservers, which respond with delegation to the {ns1.provider.net, ns2.provider.net} authoritative nameservers. Separate DNS queries for the A records of both ns1.provider.net and ns2.provider.net are sent to resolve the authoritative nameservers' IP addresses. Using the answer from this "glueless" lookup, the resolver queries the authoritative nameservers for example.com's A records.

Next, we consider DNSSEC records. We refine this expression for the case where the preceding zone's response contains a DS signing key, if DNSSEC records are registered:

$$I^{+DNSSEC}(Z_i) = \begin{cases} \emptyset & \text{if DS key in } Z_{i-1} \text{ response} \\ I(Z_i) & \text{otherwise} \end{cases}$$

For the base case of the root zone Z_0, the full-graph DNS attack surface is the empty set: the IP addresses of the root nameservers are hardcoded in operating system root hints files and the root nameserver responses are signed by the global DNS root key, so the root zone is safe from BGP hijacks:

$$I^{+DNSSEC}(Z_0) = \emptyset$$

Combining these equations, the quantity $Q^+(d)$ represents the set of IP addresses that may be BGP hijacked when DNS resolving and connecting to the domain d. $Q^+(d)$ can be derived by performing an iterative DNS query starting at the root nameservers and resolving/recording the glued IP addresses of *every nameserver* in each successive zone.

$$Q^+(d) = I^{+DNSSEC}(Z_{k_d}) \cup A(d)$$

5.2 Designing a Full-Graph DNS Resolver

While prior work had developed the notion of full-graph DNS and its implications for domain name hijacking, this work focused on DNS-level exploits and did not consider the effect of insecure Internet routing [47]. There has also been little exploration of the usage of DNSSEC and its implication within the context of the PKI and routing security. (A more thorough comparison of prior work can be found in § 9.)

Our DNS resolver tool differs from classic open-source resolver implementations and prior work in these aspects:

Full graph DNS lookups considering DNSSEC. Referring to the definition in § 5.1, the resolver tool computes the full-graph DNS attack surface $Q^+(d)$ of the domain name d, which is a set of IP addresses including d's A records and all its nameserver dependencies. The tool also records the presence of DNSSEC DS records in each successive zone response, as the DS key provides a cryptographic authenticity check on the DNS response that can be validated by the CA. DNS responses from the root zone and many top-level domain zones (as well as the authoritative resolver on about 5.6% of domains) have DS key records securing their responses. Popular open-source tools like KnotDNS or Unbound lack full-graph DNS resolution capability.

DNS lookups near realtime to cert issuance. Because TLS certificates do not certify fields related to routing or DNS, a domain's DNS configuration and records can change any time after certificate issuance. Previous work found that DNS queries for a high proportion of Let's Encrypt certified-domains returned NXDOMAIN errors when queried weeks after certificate issuance time [9]. To mitigate these effects, we perform DNS lookups shortly after issuance to ensure DNS conditions closely match those at the time Let's Encrypt performed domain validation.

Distribution DNS lookups geographically. A domain name may resolve to different IP addresses when queried from different geographic regions. This behavior is indicative of Content Distribution Network (CDN) use, and we found that paths to local content replicas are often much shorter than the path to a single origin server potentially in another region, which increases resilience to hijacking.

Repeated DNS lookups. We observed that around 3% of domains resolved to different IP addresses when a DNS lookup was performed multiple times. The two major factors influencing this were 1) DNS load-balancing systems where only a random subset of IP addresses (drawn from a large pool of replicas) were returned on each query and 2) path-dependent domains where queries to different authoritative nameservers returned answers with different A records. Since an adversary can realistically re-attempt obtaining a certificate several times (and even impact which nameserver is used by the CA), we consider any IP address that could be returned in a lookup part of the BGP attack surface and performed each lookup 10 times to discover as many IP addresses as possible.[3]

5.3 Global Near-Real-Time Domain Resolution Infrastructure

In support of our work, Let's Encrypt provided us access to daily certificate issuance and domain validation logs, which contained information for all certs issued in the preceding 24 hours (an average of 1.08 million certs per day). Over the period of April 20 - August 13 2022, we randomly sampled 10K certificates daily from these logs, yielding a total dataset of 1.354 million domains from 810,000 certificates.

Daily DNS resolution jobs using the tool described in § 5.2

[3] After 10 lookups, there was a significant reduction in the number of new A record IP addresses returned, making 10 a logical cap for repeated queries.

Feature	Figure
Total number of certificates	810,000
Number of certs. successfully resolved	755,942
Total number of domains	1,354,318
% successful A record resolution	97.3
% successful AAAA record resolution	12.3
IP prefixes of domains	
Median number of prefixes in A records	1
Median number of prefixes in NS records	3
Little use of DNSSEC	
% domains full DNSSEC-signed	5.6
Registration of RPKI-ROA records	
% domains with at least 1 ROA-covered prefix	76.3
% domains with all ROA-covered prefixes	26.2
% target IPs with ROA records	60.0

Table 2: An overview of DNS dataset, summarized.

were run in three continents at six AWS server locations [4]: (1) **North America:** Ohio, Oregon; (2) **Europe:** Paris, Frankfurt; (3) **Asia:** Singapore, Tokyo, as close to certificate issuance time as possible (5 hours after log upload). Chiefly, we are interested in performing iterative queries for the A (IPv4 address) and AAAA (IPv6 address) records of domains in the subject names of certificates issued by Let's Encrypt.

Implementation details. The primary implementation challenge lay in operation at scale: our tool needed to efficiently conduct millions of DNS requests per day, tolerating various malformed records and edge cases, and output a concise log that captured the full history of the lookup and that could be plugged into our analysis engine (see Figure 3). We release our collected DNS dataset and DNS resolver tool as open-source software on our Github repository[5].

5.4 Profile of LE-certified Domains

In this section, we present key statistics on the routing implications of domains' DNS and webserver configurations, summarized in Table 2. This represents the first study to date of features of domains included in Let's Encrypt certificates.

Multiple prefixes for hijack targets. Of the 1.248 million domain names sampled, 97.8% and 12.3% could be resolved to valid A and AAAA records successfully. Note that the low rate of AAAA record retrieval is not a consequence of the tool itself, but because of low IPv6 usage by the domain names surveyed. On average, the domains surveyed had webservers hosted across 1.158 distinct IP prefixes (median 1.0) and DNS nameservers hosted across 4.72 prefixes (median 3.0). 86.2% (11.0%), 11.7% (29.1%), 0.73% (9.4%), 0.56% (11.1%), 0.18% (3.5%) of webservers (DNS nameservers in parentheses) are associated with 1 to 5 IP prefixes, respectively. A detailed distribution is in Appendix Figure 7. Overall, DNS presents nearly 5x more potential routing hijack targets to a hijacker than the website servers alone. 41.9% of nameservers used are hosted on /24 prefixes, meaning they are resistant to sub-prefix hijacks. Webservers tended to be hosted on shorter prefixes: only 35.0% of webserver IPs are hosted on /24 prefixes. We analyze the implications of these statistics for sub-prefix hijacks in § 8 and Appendix § D.

[4]Let's Encrypt's current multiVA infrastructure hosts vantage points in Ohio, Oregon, and Frankfurt AWS regions (region codes listed in Table 6).

[5]https://github.com/inspire-group/routing-aware-dns

Provider	Prop. of NS (%)	ROA coverage (%)
CLOUDFLARENET	28.3	98.2
AMAZON-02	14.4	98.9
AKAMAI-ASN2	7.1	100
NSONE	3.1	50.0
GODADDY-DNS	2.8	100
UltraDNS	2.8	11.1
Google, US	2.7	100
Total	61.2	-

Table 3: Top nameserver hosting providers and the proportion of their network prefixes with valid ROA records.

Sparse adoption of DNSSEC. We find that DNSSEC registration is sparse: only 5.6% of A records returned were fully signed by DNSSEC at every link in the nameserver delegation chain, AAAA records were somewhat more fully signed at 16.0%. Overall, DNSSEC ensures integrity of DNS records for only a very small minority of domains.

Growth of RPKI. Our study gives encouraging evidence for increasing ROA registration, especially at the nameserver level. 63.6% of all nameservers surveyed are hosted on IP prefixes with valid ROAs. ROA registration is slightly lower for webserver IPs, with 51.8% of webserver IPs having ROAs.

DNS nameserver centralization. Our measurement shows the global DNS infrastructure is highly centralized in terms of prefix space and operational ownership. Only 100 prefixes, owned by 20 distinct organizations, represented 56.1% of non-DNSSEC nameservers used across the domain dataset: Cloudflare alone operates over 24% of these nameservers. Five Cloudflare-hosted prefixes alone accounted for 11.9% of all nameserver prefixes surveyed. The top nameserver hosting providers and their ROA usage statistics are summarized in Table 3. One advantage of this centralization is that it facilitates ROA adoption as a sizeable portion of the DNS ecosystem can be protected with ROA deployment at only a handful of top providers. The proportion of ROA coverage at these providers tends to be higher than the overall routing table average of 41% [44], which is in part why ROA helps to offset the attack surface introduced by DNS.

Implications for the attack surface. These findings outline several new opportunities for BGP hijacking using DNS. DNS nameservers present a higher number of potential hijack targets with greater geographical diversity, affording attackers more chances to conduct a hijack localized to a CA's vantage point(s) and nameserver that evades global BGP detection. Overall, the additional attack targets may allow some adversaries to succeed in hijacking a domain when they would not have been able to target the domain's A records directly. Next,

we will use Internet topology simulation and a quantitative metric to evaluate the security of multiVA deployments.

6 Internet Topology Simulations: A Multi-Faceted Approach

Given the potential target IP addresses for each domain from our DNS measurements, we estimate each domain's vulnerability to BGP hijacks via simulation. We focus on simulation due to ethical and modeling concerns. Launching BGP attacks against the real production infrastructure hosting TLS domains is clearly unethical. Previous work [9] has conducted ethical BGP attacks, but these were carried out on target domains specifically created by researchers for BGP experiments and using nodes that allow researchers to make BGP announcements (e.g., PEERING testbed [49]). The setup does not accurately replicate the real-world network topology and configurations of web and DNS infrastructure. Given that our study aimed to evaluate the effectiveness of multiVA against hijacks on genuine TLS domains, we have instead opted to simulate attacks on real hosting infrastructure.

We divide this task into two parts: 1) Internet-wide topology simulations that model the routing of every IP prefix on the Internet under various RPKI conditions and simulated attacks from 1 K randomly sampled adversaries and 2) domain resilience computation using data of target IPs for each domain and routing information from the simulator. This approach elucidates how vulnerable a domain is to BGP attacks.

6.1 Simulation Methodology

To measure the impact of BGP attacks on the PKI, we run global equally-specific BGP simulations. We model Gao-Rexford route preferences [23] over the CAIDA AS topology model [14], and perform IP prefix-level (instead of AS-level) simulations using RIB data from RouteViews [48] and RIPE RIS [43]. Given a specific adversary AS and a specific victim AS, we simulate whether a potential set of vantage points would route data to the victim or the adversary.

Novel collection of cross-cloud datacenter peering connections. In addition to using prefix-level simulations, we needed to accurately model the BGP connections of different potential vantage points used in our simulation. Because public BGP data sources only have a partial view of global routing, many (particularly peering) links that are heavily used by cloud data centers are missed in both public topologies and BGP data. We augment these data sources with neighbors found by running traceroute and the bdrmap [40] tool at all cloud vantage points considered. This work is the first to consider multiVA deployments spread across multiple cloud providers (as previous work considered vantage points hosted solely in AWS datacenters [9]). We collected bdrmap data from 19 distinct datacenters spread across three cloud providers (AWS, Azure, and GCP) and four continents (Europe, Asia, North America, and South America).[6] By collecting bdrmap data across different cloud datacenters, we can determine the optimal providers for hosting vantage points and consider the effectiveness of cross-cloud multiVA deployments.

Simulation of RPKI-enabled ROV. In addition to non-RPKI simulations, we also ran RPKI-enabled simulations for every prefix in the route table. In the RPKI simulations we modeled the adversary's announcement as having the true victim's origin ASN prepended to evade RPKI-based Route Origin Validation (or ROV). All of Let's Encrypt's current AWS vantage points perform ROV [32]. Any announcements for an equally-specific IP prefix covered by RPKI that does not contain the true origin ASN will be filtered and not used by Let's Encrypt's vantage points. Even in a worst-case scenario where no other ASes other than AWS perform ROV, AWS would have no RPKI-valid routes to the victim and end up with no route to the victim's prefix in its routing table, resulting in traffic being null routed. In this event (where no AWS vantage points use the adversary's route) the adversary cannot succeed in obtaining a certificate. Thus, **the only way for an adversary to succeed in an equally-specific prefix attack against an RPKI-protected prefix is to evade ROV and prepend the true origin ASN to its announcement**. While this is a viable strategy (which has been observed in the wild to evade ROV [29]), it does cause the adversary's malicious BGP announcement to be longer and less-preferred.

Our resilience processor is also capable of incorporating the concrete usage of ROAs and ROVs in the Internet today. It does this by loading both the RPKI and non-RPKI simulations and then determining on a per-IP basis which set of simulation results to use based on one of three different modes:

1. No RPKI adoption: This mode assumes that no IP addresses are covered by RPKI and the non-RPKI simulation results are used. We include it as a baseline for comparison with previous work (which did not consider RPKI).

2. Current RPKI adoption: This mode uses data from Routinator [34] to determine which IP addresses are currently covered by RPKI, enabling the resilience processor to use the appropriate (i.e., RPKI or non-RPKI) simulation results for each IP address. Routinator is open-source RPKI Relying Party software that downloads RPKI data from the five RIR trust anchors and produces a dataset of IP prefixes covered by RPKI, identical to the procedure used by real routers when producing filtering rules for ROV. An IP address is counted as RPKI-protected if its prefix and origin AS match an ROA record in the downloaded dataset.

3. Full RPKI adoption: This mode represents the "best case" scenario assuming RPKI adoption has expanded to the entire routing table. In this mode only the RPKI simulation results are used. While current RPKI adoption models the cur-

[6] Let's Encrypt's two primary data centers in Denver and Salt Lake City operate out of a provider that does not lease cloud services preventing us from running bdrmap. We use public BGP data for these providers.

rent routing table, we expect that RPKI adoption will increase over time, which will improve the resilience of domains and bridge the gap to resilience of the full RPKI adoption results.

Finally, we simulated BGP attacks on every prefix seen in the public BGP data from a set of one thousand adversaries randomly sampled from all ASes in the topology.

6.2 Defining a Domain's Resilience

Considering the output of the Internet topology simulations and DNS lookups, we compute an effective resilience for each domain contained in the DNS data. The effective resilience measures the fraction of ASes on the Internet that are topologically *incapable* of acquiring a certificate for a given domain by launching an equally-specific BGP attack [9,13,36,56]. Intuitively, domains with weak BGP connectivity have a lower effective resilience as it is more likely an adversary's attack will succeed. Domains with richer BGP connectivity have a high effective resilience. In addition to domain connectivity, the number of CA vantage points and quorum policy impact the effective resilience of domains as stricter quorum policies and more vantage points allow a CA to detect more BGP attacks and reduce the likelihood that an attack will be viable.

We compute the effective resilience of each domain by considering each vantage point and seeing which adversaries are able to successfully hijack communication from that vantage point to *any* target IP address in $Q^+(d)$ (defined in § 5.1) for the domain in question (see Appendix Algorithm 1). If, based on our simulations, an adversary succeeds in hijacking traffic from a vantage point to either an A record IP address or a vulnerable DNS nameserver IP address, we consider the adversary to be capable of hijacking the validation request from that vantage point. We then consider the quorum policy and compute, for a given domain, which adversaries are capable of hijacking validation from enough vantage points to satisfy the quorum policy. If an adversary cannot hijack enough validation requests to satisfy the quorum policy, the domain is considered to be resilient to attacks from that adversary. The fraction of randomly sampled adversary ASes that a domain is resilient to is the domain's resilience value. We present the mathematical definition of resilience in Appendix A.

7 Quantifying the Impact of Real-World Dynamics on multiVA

We run resilience simulations on the DNS dataset in § 5.4 to quantify the change in resilience resulting from considering DNS nameservers, RPKI-ROA, and DNSSEC in the BGP hijack attack surface. Our simulations aimed to model a wide array of possible vantage point configurations as well as Let's Encrypt's current multiVA deployment. Our results show that DNS significantly broadens the attack surface available to a BGP hijacking-adversary: for a single vantage point, DNS translated to an average resilience drop of over 26%. At the same time, RPKI-ROA records provided a resilience gain of 9.9% on average across all configurations tested.

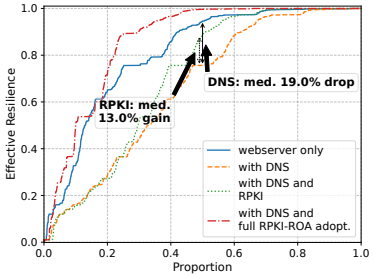

(a) multiVA

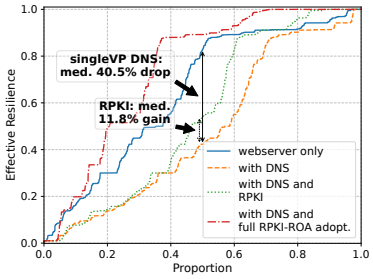

(b) singleVA

Figure 5: Comparing the resilience of **(a)** Let's Encrypt's multiVA and **(b)** singleVA deployment under different conditions.

7.1 Re-evaluating Let's Encrypt's multiVA Deployment

Resilience of LE deployment drops when considering DNS. Figure 5a compares the resilience of Let's Encrypt current deployment of multiVA domain validation (with three vantage points) when considering case (1) BGP hijacks on only the domain's webserver; case (2) BGP hijacks on the domain's webserver and non-DNSSEC signed-nameservers; case (3) BGP hijacks on all targets in (2) while accounting for the usage of RPKI-ROV by the ASes hosting Let's Encrypt's vantage points. Including DNS as an hijack target reduces median domain resilience by over 20%, from 94.6% to 75.6%. Approximately one-third of domains had a resilience of less than 0.5, meaning that they can be hijacked by the adversary in the average case. The result suggests DNS indeed enlarges the attack surface of multiVA and affects its security.

We also plot the effect of cases (1) and (2) for Let's Encrypt's original single vantage point-deployment prior to 2020 in Figure 5b. The drop in resilience due to NS records is dramatically higher for the single vantage point case: here resilience drops by over 33%. Thus, although multiVA is an incomplete remedy for DNS-based BGP hijacks, it minimizes the resilience drop compared to the singleVA case.

RPKI improves resilience even more so for multiVA. The

# Additional VPs	Quorum	Resilience (median)	
		Single-cloud	Multi-cloud
0	**n – 1**	**88.6**	—
1	n-1	93.2	93.2
2	n-1	94.6	97.5
3	n-1	95.8	97.7
4	n-1	96.9	97.8
5	n-1	97.0	98.6
0	n	96.5	—
1	n	97.5	98.4
2	n	98.2	99.3

Table 4: Effect on resilience with varying VP counts/locations and quorum policies under current RPKI-ROV deployment. (Boldface indicates LE's present deployment.) Extra VPs are *in addition* to LE's present deployment of 3 remote VPs.

resilience degradation presented in § 7.1 is pessimistic and does not consider the enhancement in routing security attributable to RPKI-ROA deployment. When considering deployed RPKI-ROA on top of the DNS model, resilience increases significantly, from 75.6% to 88.6%. When considering the scenario of full RPKI-ROA adoption (i.e., where all prefixes in the routing table have a valid ROA registered), the median domain resilience shoots up to 99.6%. Of note is that the resilience benefit from RPKI-ROA is significantly higher for multiVA than the single vantage point case: the median domain resilience under the full RPKI-ROA adoption regime for singleVA is only 89.2%. The variance in resilience measurements highlights the importance of considering DNS and RPKI data in the domain resilience estimation. Still, the registration of RPKI records for domain-hosting prefixes is largely beyond the CA's control; in light of this, we explore more multiVA configurations in § 7.2 to augment CA resilience.

Takeaways. Our results suggest the necessity of considering deployed security measures in evaluation for future multiVA security research. This produces a more nuanced and accurate estimate of a multiVA deployment's security. To protect their domains, domain owners should choose hosting providers that enable RPKI [15] and turn on DNSSEC if possible.

7.2 Strategies to Enhance multiVA Security

In this section, we explore techniques to improve the security of multiVA deployments. We consider four factors that a multiVA deployment can control: the number of vantage points, vantage point location, vantage point cloud provider, and quorum policy to issue certificates. We compute the resilience of Let's Encrypt's multiVA under 11,110 combinations of these parameters, and summarize our observations from the optimal configurations. We use the resilience of Let's Encrypt current deployment while modeling DNS and RPKI as the baseline.

Adding vantage points using AWS. Let's Encrypt's current deployment of multiVA uses remote vantage points hosted exclusively in AWS datacenters in addition to its primary datacenter hosted by the datacenter provider Flexential. Let's Encrypt chose this approach to avoid the complexity and engineering overhead incurred with deploying vantage points in multiple cloud providers [9]. Thus, a logical and low-cost extension to LE's deployment would be to add additional AWS vantage points. We find that a single cloud provider deployment requires a significant number of vantage points for increased effectiveness (see Figure 6a). For example, the AWS deployment flatlines at a maximal 97% resilience after adding 5 more vantage points to Let's Encrypt's deployment.

Considering all AWS vantage points surveyed, we found the optimal locations (i.e., those with the largest resilience increase) to be concentrated in Asia (Tokyo, Mumbai, and Singapore), along with another perspective in Northern Europe (Stockholm) and South America (Sao Paulo).

Adding one extra AWS vantage point under the same quorum policy in Europe (Stockholm) shows promise, improving resilience from 88.6% to 93.2%. We posit this result is an artifact of Let's Encrypt's current deployment and quorum policy, which allows issuance to proceed if one vantage point fails validation. Because two of three of Let's Encrypt's remote vantage points are located in the US, domain validation can still pass if only the Frankfurt vantage point (which provides greater routing diversity than the US vantage points [9]) detects an attack. Adding Stockholm as a vantage point supports Frankfurt in representing the European vote and protecting against attacks where the European perspective is crucial.

However, adding another AWS vantage point only improves resilience to 95.1%, and a third brings resilience to 96.6%. A fourth vantage point yields an incremental resilience gain of 0.3% more to 96.9%, and a fifth maximizes at 97.0%. Although the resilience benefit of adding more AWS nodes saturates, this approach still does increase resilience by nearly 10% with minimal engineering setup and management costs.

Adding vantage points using additional cloud providers. A more substantial resilience gain can be achieved by adding additional vantage points in alternate cloud providers. If one more vantage point is added, Microsoft Azure offers the optimal next vantage point location in West Europe (Amsterdam), which improves resilience by a slight 0.1% over the first optimal EC2 location. If two additional vantage points are deployed, the optimal configuration becomes GCP, with nodes in Asia Northeast (Tokyo) and Asia Southeast (Singapore); this boosts median resilience to 97.5% representing a 8.9% improvement over Let's Encrypt's current deployment and a 2.9% improvement over an all-AWS based deployment with the same number of vantage points. These gains can be further maximized with three additional vantage points spread across AWS and GCP, which scores a 97.7% resilience figure. Finally, a deployment of 5 nodes across all three cloud providers studied reaches a resilience of 98.6% (see Figure 6b).

The advantage of a cross-cloud deployment can be found in the peers for various datacenters found with bdrmap. If two

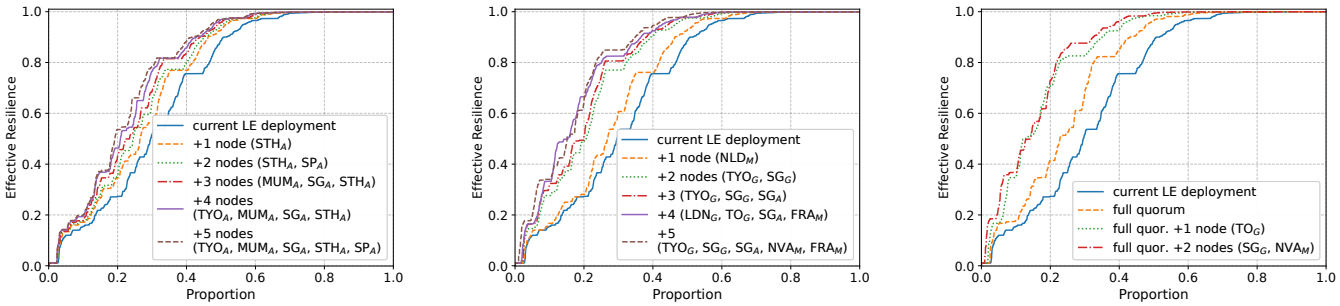

(a) (b) (c)

Figure 6: Comparing resilience of Let's Encrypt's deployment when **(a)** adding AWS vantage points; **(b)** adding cross-cloud vantage points; **(c)** changing to a full quorum policy.

Provider	AWS	GCP	Azure
AWS	0.34	0.10	0.15
GCP	0.10	0.54	0.25
Azure	0.15	0.25	0.78

Table 5: Average fraction of overlapping peers between datacenter pairs of different providers. Because of homogeneity within providers, pairs of datacenters from different providers have a lower average overlap than pairs from a single provider.

datacenters have a higher percentage of overlapping peers, they are likely to have a higher degree of similarity in BGP routing tables as routing by cloud providers relies heavily on peer-learned routes [3]. We calculated the fraction of overlapping peers between two datacenters (with sets of peers P_1 and P_2) as $\frac{|P_1 \cap P_2|}{max(|P_1|,|P_2|)}$. We computed this overlap fraction between all pairs of datacenters we studied and grouped the results by the cloud providers. We found that on average datacenters of the same provider had a higher overlapping peer fraction than datacenters of different providers (see Table 5). This explains why a cross-cloud deployment can achieve more route diversity even with fewer vantage points.

Choosing a more stringent quorum policy. Another approach to improve resilience is to change the quorum policy from allowing one vantage point to fail to requiring complete consensus among vantage points. This makes the adversary's task significantly harder because it must successfully hijack traffic from all of Let's Encrypt's vantage points, as a single vantage point routing to the victim will block issuance. While this does significantly improve resilience, it also incurs additional false positives (benign failures in the absence of an attack) as discussed in prior work [9]. The CA needs to strike a balance between security and accuracy. One way to achieve this while still strengthening the quorum policy is to encourage users and hosting providers to wait until DNS propagation is complete before requesting a certificate as partial DNS propagation is a leading cause of false positives [9]. In addition, distinguishing benign failure modes (like NXDOMAIN errors at remote vantage points not indicative of malicious activity) from the types of responses that can arise during an attack may reduce the false positive rate and make changing the quorum policy viable [9]. While additional research of these engineering options may be needed to deploy full quorum policy at scale, we believe that understanding their security implications is a critical next step for multiVA.

Changing the quorum policy with the current vantage point set improves median domain resilience to 96.5%. This near-8% boost is likely due to the dynamics of Let's Encrypt's current deployment, where Frankfurt (the only non-US vantage point) can be ignored in the current quorum policy, allowing quorum to be met with only US-based vantage points. If we further consider adding a vantage point under the full quorum policy, **a 98.4% resilience can be achieved, counterbalancing the extra attack surface introduced by DNS**. Adding two vantage points raises resilience to 99.3% (see Figure 6c).

8 Discussion

Viability of sub-prefix hijacks. Although we do not position multiVA as a defense against sub-prefix attacks, our dataset sheds light on their viability. In the absence of BGP monitoring or other countermeasures to mitigate sub-prefix attacks, an IP prefix may have two properties that make it impervious to this class of attacks: a prefix length of 24 bits, or coverage by properly configured RPKI ROAs and routed on a path with full ROV enforcement. Under this definition, 29% of domains are immune to sub-prefix attacks as *all of their associated target IPs are immune to sub-prefix attacks*. In Appendix D we discuss the details of our sub-prefix attack analysis and ecosystem improvements (such as additional ROV enforcement and proper maxLength configuration) that can increase the proportion of domains immune to sub-prefix attacks to 62%.

Effectiveness of multiVA against real-world BGP hijack-

ers. We evaluate the performance of multiVA against a set of 13 ASes identified as "serial hijackers" (ASes with a history of performing BGP hijacks) [54]. with the goal of capturing multiVA's effectiveness against more powerful known attackers. Our analysis illustrates multiVA's strength: multiVA outperforms singleVA by nearly a factor of two, showing a resilience improvement from 39% to 69%. Furthermore, the proposed deployment extensions in § 7.2 yield sizable resilience gains against these serial hijackers: adding just one additional vantage point raises median resilience to 84.6%. We observe that these ASes are more effective attackers compared to our random sample of 1K ASes, a suspected consequence of sampling bias in the serial hijackers dataset which contains more richly connected and prominent ASes. Full results from the serial hijacker simulations are presented in Appendix C.

9 Related Work

Network attacks on domain control validation. Domain control validation is a critical service and is vulnerable to several different types of network attacks. Gavrichenkov *et al.* discussed the vulnerability of domain control validation to BGP attacks [24] and Birge-Lee *et al.* first ethically demonstrated such attacks in the wild [8], which inspired the design of multiVA and its subsequent deployment at Let's Encrypt [9].

Several works (like those by Birge-Lee *et al.* [8, 9]) have examined the domains' resilience to BGP attacks, but with a significantly different (and limited) attack surface that did not consider DNS or RPKI. A more recent poster by by Brandt *et al.* on DNS attacks on domain validation did a preliminary investigation of the impact of DNS [13] but did not fully capture the DNS attack surface by not considering full DNS lookup graphs, DNSSEC, or geographical variance of DNS. Furthermore, it does not consider RPKI or vantage points beyond Let's Encrypt's current deployment, and it does not offer recommendations for multiVA enhancement through expansion to different cloud providers or quorum policy changes [13].

Beyond BGP attacks, domain control validation is vulnerable to other classes of network attacks, including non-BGP-based DNS attacks [17]. Brandt *et al.* recommend a multiVA-based approach as a mitigation for DNS attacks on domain control validation [12]. Dai *et al.* discuss DNS attacks against Let's Encrypt's multiVA and demonstrate a novel technique to influence a vantage point to query a chosen authoritative nameserver [17]. We encompass this work in our analysis by assuming that an adversary capable of compromising **any** non-DNSSEC authoritative nameserver associated with a victim's domain can obtain a fraudulent certificate.

Recent work shows that ROV can be undermined by attacks on DNS (including BGP attacks) as ROV may depend on insecure DNS nameservers [28]. While these attacks and their interactions with other BGP attacks deserve further research, we omit them here as they can be mitigated by changes in ROV deployments at ROV-enforcing networks [28].

BGP attacks and defenses. Our work builds upon foundational prior work in BGP security. The notion of domain resilience (similar to the constructions in [8–10, 53]) was originally adapted from AS-resilience as defined by Lad *et al.* [36]. Additionally, the algorithm behind our simulation engine is based on the algorithm presented by Gill *et al.* [26].

Prior work in BGP security has also shaped our analysis. Crucially we incorporate RPKI [38] at its currently deployment in the routing table. We also note Gilad *et al.*'s discussion of how the maxLength attribute in RPKI [25] can undermine its security if misset; however, this vulnerability is fundamentally a sub-prefix attack, which is outside our threat model. We further acknowledge the impact of BGP security practices like prefix filtering (as recommended by the MANRS project [41]) and peer locking [51]. While these practices have a positive impact on network security, we do not consider these in our work. The MANRS prefix filtering techniques are not foolproof and have been evaded by real-world attacks [29], and peer locking is only viable for ASNs directly connected to a handful of participating providers [51].

DNS Mapping. Prior work [47] introduced the notion of full-graph DNS and the implications of transitive trust between nameserver delegations for hijacking DNS names. Other work [18, 19, 46] has leveraged full-graph DNS resolution to understand DNS resolution dependency to study the security, availability, and robustness of DNS. The full-graph resolvers in prior work do not record DNSSEC, consider the names in glued records as dependencies, and are not open-sourced, thus necessitating us to develop our own.

10 Conclusion

We develop a novel analysis framework which captures the most complete attack surface of BGP hijacks on CAs to date. Our work provides rigorous analysis and exhaustive quantitative data about the effectiveness of both current and future multiVA under real-world conditions. The CA/Browser Forum has begun discussion of mandating multiVA industry-wide based on results from this work [22]. Based on this analysis, Let's Encrypt plans to expand their deployment by adding an additional AWS vantage point in Stockholm (the optimal location identified in our vantage point search) [2].

Our analysis framework that encompasses both the DNS attack surface and protections offered by RPKI has broad applicability to network security and privacy domains beyond the PKI and constitutes a primary contribution of this paper. Collecting enough data to model the intricacies of routing in the context of certificate issuance entailed significant financial resources and engineering efforts which can benefit future research in the security community. For example, our methodology can aid developers of anonymity systems and cryptocurrencies to assess security against BGP attacks [42, 53].

11 Acknowledgments

We would like to thank our shepherd for their guidance, Joon Kim, Mona Wang, and Sophia Yoo for their detailed feedback, and the anonymous USENIX reviewers for their suggestions and comments. We extend special thanks to Let's Encrypt for their partnership on this work, in particular Jillian Karner for facilitating data access, the engineering team (including Aaron Gable, James Renken, and others) for their work on Let's Encrypt's multiVA deployment, and Josh Aas for his continued support. This work was supported in part by the National Science Foundation under grants CNS-1553437, CNS-1704105, CNS-2131938, the Open Technology Fund, and by the United States Air Force and DARPA under Contract No. FA8750-19-C-0079. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force, DARPA, or any other sponsoring agency.

References

[1] Josh Aas, Daniel McCarney, and Roland Shoemaker. Multi-Perspective Validation Improves Domain Validation Security. https://letsencrypt.org/2020/02/19/multi-perspective-validation.html, 2020.

[2] Anonymous. Private conversations with Josh Aas, Executive Director of Let's Encrypt, March 2023.

[3] Todd Arnold, Jia He, Weifan Jiang, Matt Calder, Italo Cunha, Vasileios Giotsas, and Ethan Katz-Bassett. Cloud Provider Connectivity in the Flat Internet. In *Proceedings of the ACM Internet Measurement Conference*, 2020.

[4] AWS. Settlement Free Peering Policy Interconnection with AS16509. https://aws.amazon.com/peering/policy/.

[5] Devin Bayer. Visibility of Prefix Lengths in IPv4 and IPv6. https://labs.ripe.net/author/dbayer/visibility-of-prefix-lengths-in-ipv4-and-ipv6/, 2010.

[6] BGPmon. https://bgpmon.net/.

[7] Henry Birge-Lee. Attackers exploit fundamental flaw in the web's security to steal $2 million in cryptocurrency. https://freedom-to-tinker.com/2022/03/09/attackers-exploit-fundamental-flaw-in-the-webs-security-to-steal-2-million-in-cryptocurrency/, 2022.

[8] Henry Birge-Lee, Yixin Sun, Anne Edmundson, Jennifer Rexford, and Prateek Mittal. Bamboozling Certificate Authorities with BGP. In *USENIX Security Symposium*, 2018.

[9] Henry Birge-Lee, Liang Wang, Daniel McCarney, Roland Shoemaker, Jennifer Rexford, and Prateek Mittal. Experiences Deploying Multi-Vantage-Point Domain Validation at Let's Encrypt. In *USENIX Security Symposium*, 2021.

[10] Henry Birge-Lee, Joel Wanner, Grace H. Cimaszewski, Jonghoon Kwon, Liang Wang, François Wirz, Prateek Mittal, Adrian Perrig, and Yixin Sun. Creating a secure underlay for the internet. In *USENIX Security Symposium*, 2022.

[11] Russell Brandom. Hackers Emptied Ethereum Wallets by Breaking the Basic Infrastructure of the Internet. The Verge, 2018.

[12] Markus Brandt, Tianxiang Dai, Amit Klein, Haya Shulman, and Michael Waidner. Domain Validation++ For MitM-Resilient PKI. In *ACM SIGSAC Conference on Computer and Communications Security*, 2018.

[13] Markus Brandt, Haya Shulman, and Michael Waidner. Evaluating Resilience of Domains in PKI. In *ACM SIGSAC Conference on Computer and Communications Security*, 2021.

[14] CAIDA. The CAIDA AS relationships dataset. https://www.caida.org/catalog/datasets/as-relationships/, 2022.

[15] Cloudflare. Is BGP safe yet? https://isbgpsafeyet.com/, 2023.

[16] Cogent. Cogent Customer User Guide. https://www.cogentco.com/files/docs/customer_service/guide/global_cogent_customer_user_guide.pdf.

[17] Tianxiang Dai, Haya Shulman, and Michael Waidner. Let's Downgrade Let's Encrypt. In *ACM SIGSAC Conference on Computer and Communications Security*, 2021.

[18] Casey Deccio, Chao-Chih Chen, Prasant Mohapatra, Jeff Sedayao, and Krishna Kant. Quality of name resolution in the domain name system. In *2009 17th IEEE International Conference on Network Protocols*, pages 113–122. IEEE, 2009.

[19] Casey Deccio, Jeff Sedayao, Krishna Kant, and Prasant Mohapatra. Measuring availability in the domain name system. In *2010 Proceedings IEEE INFOCOM*, pages 1–5. IEEE, 2010.

[20] Chris C Demchak and Yuval Shavitt. China's maxim–leave no access point unexploited: The hidden story of

china telecom's bgp hijacking. *Military Cyber Affairs*, 3(1):7, 2018.

[21] Let's Encrypt. Rate Limits. https://letsencrypt.org/docs/rate-limits/, 2021.

[22] CA/Browser Forum. Minutes of the F2F 58 Meeting in Ottawa, Canada, 28 Feb-2 March 2023 – Validation SC (2 March). https://cabforum.org/2023/03/02/minutes-of-the-f2f-58-meeting-in-ottawa-canada-28-feb-2-march-2023-validation-sc-2-march/, 2023.

[23] L. Gao and J. Rexford. Stable Internet Routing without Global Coordination. *IEEE/ACM Transactions on Networking*, 9(6):681–692, 2001.

[24] Artyom Gavrichenkov. Breaking HTTPS with BGP Hijacking. *Black Hat*, 2015.

[25] Yossi Gilad, Sharon Goldberg, Kotikalapudi Sriram, Job Snijders, and Ben Maddison. The use of maxLength in the RPKI. Internet-Draft draft-ietf-sidrops-rpkimaxlen-15, IETF Secretariat, August 2022.

[26] Phillipa Gill, Michael Schapira, and Sharon Goldberg. Modeling on quicksand: Dealing with the scarcity of ground truth in interdomain routing data. *ACM SIGCOMM Computer Communication Review*, 42(1):40–46, jan 2012.

[27] Dan Goodin. How 3ve's BGP Hijackers Eluded the Internet – and Made $29M. https://arstechnica.com/information-technology/2018/12/how-3ves-bgp-hijackers-eluded-the-internet-and-made-29m/, 2018.

[28] Tomas Hlavacek, Philipp Jeitner, Donika Mirdita, Haya Shulman, and Michael Waidner. Behind the Scenes of RPKI. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, CCS '22, 2022.

[29] Peter Kacherginsky. Celer Bridge incident analysis. https://www.coinbase.com/blog/celer-bridge-incident-analysis, 2022.

[30] Dan Kaminsky. The Great DNS Vulnerability of 2008 by Dan Kaminsky. DUO Security, 2016.

[31] David Kluge and David Warner. Google trust services acme api available to all users at no cost. https://security.googleblog.com/2023/05/google-trust-services-acme-api_0503894189.html, 2023.

[32] Fredrik Korsbäck. How AWS is helping to secure internet routing. https://aws.amazon.com/blogs/networking-and-content-delivery/how-aws-is-helping-to-secure-internet-routing/, 2021.

[33] Dina Kozlov and Gabbi Fisher. Securing Certificate Issuance using Multipath Domain Control Validation. https://blog.cloudflare.com/secure-certificate-issuance/, 2019.

[34] NLnet Labs. Routinator 3000 Software. https://www.nlnetlabs.nl/projects/rpki/routinator/.

[35] Mohit Lad, Daniel Massey, Dan Pei, Yiguo Wu, Beichuan Zhang, and Lixia Zhang. PHAS: A Prefix Hijack Alert System. In *USENIX Security Symposium*, 2006.

[36] Mohit Lad, Ricardo Oliveira, Beichuan Zhang, and Lixia Zhang. Understanding resiliency of internet topology against prefix hijack attacks. In *Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 368–377, 2007.

[37] H. Landau. Certification authority authorization (caa) record extensions for account uri and automatic certificate management environment (acme) method binding. RFC 8657, RFC Editor, November 2019.

[38] Matt Lepinski and Stephen Kent. An Infrastructure to Support Secure Internet Routing. RFC 6480, 2012.

[39] Weitong Li and Tijay Chung. RoVista. https://rovista.netsecurelab.org/.

[40] Matthew Luckie, Amogh Dhamdhere, Bradley Huffaker, David Clark, and kc claffy. bdrmap: Inference of Borders Between IP Networks. In *ACM Internet Measurement Conference (IMC)*, 2016.

[41] MANRS. Mutually Agreed Norms for Routing Security. https://www.manrs.org/.

[42] Apostolaki Maria, Zohar Aviv, and Vanbever Laurent. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *IEEE Symposium on Security and Privacy (S&P)*, 2017.

[43] RIPE NCC. RIS Raw Data – RIPE Network Coordination Centre. https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data, 2021.

[44] NIST RPKI Monitor. https://rpki-monitor.antd.nist.gov/.

[45] NTT. Routing Policies NTT-GIN. https://www.gin.ntt.net/support-center/policies-procedures/routing/.

[46] Vasileios Pappas, Zhiguo Xu, Songwu Lu, Daniel Massey, Andreas Terzis, and Lixia Zhang. Impact of configuration errors on DNS Robustness. In *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 319–330, 2004.

[47] Venugopalan Ramasubramanian and Emin Gün Sirer. Perils of transitive trust in the domain name system. In *ACM SIGCOMM Conference on Internet Measurement*, 2005.

[48] University of Oregon Route Views Project. http://www.routeviews.org/routeviews/.

[49] Brandon Schlinker, Kyriakos Zarifis, Italo Cunha, Nick Feamster, and Ethan Katz-Bassett. PEERING: An AS for Us. In *ACM Workshop on Hot Topics in Networks (HotNets)*, 2014.

[50] Pavlos Sermpezis, Vasileios Kotronis, Petros Gigis, Xenofontas Dimitropoulos, Danilo Cicalese, Alistair King, and Alberto Dainotti. Artemis: Neutralizing bgp hijacking within a minute. *IEEE/ACM Trans. Netw.*, 26(6):2471–2486, Dec 2018.

[51] Job Snijders. Practical everyday BGP filtering with AS_PATH filters: Peer locking. *NANOG-67*, 2016.

[52] Carsten Strotmann. IP fragmentation and the DNS — vulnerable DNS servers. https://blog.apnic.net/2022/09/29/ip-fragmentation-and-the-dns-vulnerable-dns-servers/, 2022.

[53] Yixin Sun, Anne Edmundson, Laurent Vanbever, Oscar Li, Jennifer Rexford, Mung Chiang, and Prateek Mittal. RAPTOR: Routing Attacks on Privacy in Tor. In *USENIX Security Symposium*, 2015.

[54] Cecilia Testart, Philipp Richter, Alistair King, Alberto Dainotti, and David Clark. Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. In *ACM Internet Measurement Conference (IMC)*, 2019.

[55] Ben Wilson. Protection against BGP hijacking. https://groups.google.com/a/mozilla.org/g/dev-security-policy/c/lxiA7zcKLws/m/-1imDKu4AQAJ, 2020.

[56] Matthias Wübbeling and Michael Meier. Improved Calculation of AS Resilience against IP Prefix Hijacking. In *IEEE Conference on Local Computer Networks Workshops (LCN Workshops)*, 2016.

A Domain Resilience Including Vulnerable DNS IP addresses and A Records

We begin by considering our topology simulation results to produce an attack bit (that we will notate as α) for each adversary, destination, and vantage point tuple. For a specific adversary (a), destination IP prefix under attack (p), and vantage point (v), α indicates whether the adversary's attack is successful in attracting traffic from the vantage point to the destination (which is indicated by $\alpha(a,p,v) = 1$).

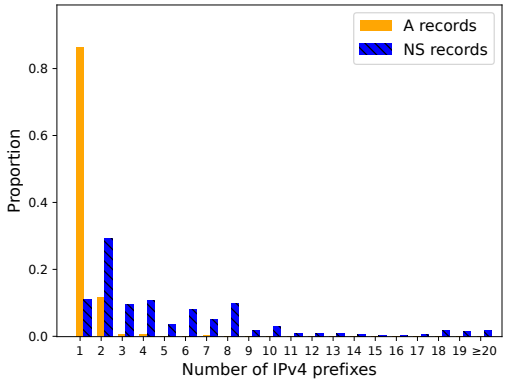

Figure 7: Distribution of number of prefixes for A and NS records, with respective median of 1.0 and 3.0 prefixes.

Algorithm 1: Resilience computation for a domain d. $Q(d)$ are the target IP addresses of d as defined in Section 5.1; V is the set of vantage points the adversary can attack.

```
function getResilience (d);
  Input  : Domain name d
  Output : res(d), where 0 ≤ res(d) ≤ 1.0
successCount ← 0;
Adversaries ← {Sampled ASes};
foreach a ∈ Adversaries do
    V ← {};
    foreach vp ∈ VantagePoints do
        if a is capable of hijacking traffic from vp
          to ANY IP addresses in Q(d) then
          V ← V ∪ vp ;
    if quorum(V) == true then successCount++ ;
return 1 − successCount / |Adversaries| ;
```

Including DNS lookup data. As prior work has established [17], DNS server selection techniques in popular DNS software (like that used by Let's Encrypt) can be exploited to allow an adversary to *select* which DNS server (or servers in a multi-level DNS lookup) a CA uses. We model this by assuming that an adversary's ability to compromise *any* DNS server in a domain's DNS lookup graph can be used to compromise the lookup of that domain. With this in mind, we analyze our DNS data to extract a target IP list for each domain, vantage point pair that contains 1) the IPs of all non-DNSSEC protected DNS servers in the DNS lookup for the domain from the region of the vantage point, appended with 2) any A records found for the domain in that region.

Given this list of target IPs for each domain, we define α^* as a function of the adversary a, the domain name d, and the

Region	Provider		
	Amazon Web Services	Microsoft Azure	Google Cloud Platform
Asia	ap-southeast-1 ap-northeast-1 ap-south-1	japan-east-tokyo	asia-southeast1 asia-northeast1
Europe	eu-west-3 eu-central-1 eu-north-1	west-europe germany-west-central	europe-west2
Americas	us-east-2 us-west-2 sa-east-1	us-east-2	us-east4 us-west1 northamerica-northeast2

Table 6: Full list of cloud datacenter locations surveyed.

vantage point v as

$$\alpha^*(a,d,v) = \bigvee_{p \in IP(d)} \alpha(a,p,v)$$

where $IP(d)$ is the list of target IP addresses for a name d. Under this definition, a domain name is vulnerable to attack at a vantage point (i.e., $\alpha^*(a,d,v) = 1$) if *any* of the target IPs for that name at that vantage point are susceptible to hijack by the adversary.

Given this definition, we incorporate the CAs quorum policy in a similar manner to Birge-Lee *et al.* [9]. By defining a quorum policy $q(\mathcal{W})$ that evaluates to 1 when $\mathcal{W}$ (the set of vantage points controlled by the adversary) is sufficient to issue a certificate, we can define $\alpha^+(a,d,q)$ for an adversary a, name d, and quorum policy q as:

$$\alpha^+(a,d,q) = q(\{v \in \mathcal{V} \,|\, \alpha^*(a,d,v) = 1\})$$

where $\mathcal{V}$ are the set of all vantage points operated by a CA.

Finally, we define the effective resilience for a domain name d, a quorum policy q, a set of vantage points $\mathcal{V}$ and a set of adversaries $\mathcal{A}$ as:

$$\gamma(d,q,\mathcal{V},\mathcal{A}) = 1 - \frac{\sum_{a \in \mathcal{A}} \alpha^+(a,d,q)}{|\mathcal{A}|}$$

Importantly, because we liberally define $\alpha^+(a,d,q)$ (where $\alpha^*(a,d,v) = 1$ if the adversary capable of attacking communication from vantage v to *any* target IP address of d), this formula underestimates resilience and incorporates all nameserver selection techniques discussed in previous work [17].

B Cloud Providers and Locations Used

Table 6 contains the list of cloud providers and locations used in the DNS data collection in § 5 and simulations in § 7.

C Effectiveness Against Serial Hijackers

The aim of these simulations was to evaluate multiVA resilience in the case of a powerful, well-known (not random) attacker AS. Simulation results for Let's Encrypt's present multiVA deployment and a singleVA deployment are presented in Figure 8. MultiVA substantially outperformed singleVA achieving a median resilience of 69% compared to 39% (for singleVA). As is the case for random adversaries, resilience can be substantially improved by adding vantage points with adding even a single vantage point increasing the median resilience to 85%. These results underscore multiVA's ability to mitigate attacks even from "strong" BGP hijackers.

D Protections Against Sub-Prefix Attacks

Sub-prefix attacks are excluded from the threat model for multiVA because they are often global in scope and multiVA is not designed to detect global attacks. Sub-prefix hijacks can be detected by alternative techniques like BGP monitoring [9, 35]. Furthermore, the rise of RPKI deployment, along with the amount of infrastructure run on /24-length prefixes substantially reduces the viability of sub-prefix attacks; we expect the viability of these attacks to continue to drop as additional security best-practices are deployed.

Sub-prefix attacks on IP addresses routed via /24 prefixes are not viable as many ASes filter prefixes longer than 24 bits [5]. Additionally, RPKI can mitigate sub-prefix attacks: ROAs specify the allowed length for announced prefixes and ROV filtering drops adversarial sub-prefix announcements whose prefix length does not match that specified in the ROA.

There are some important considerations regarding RPKI's ability to prevent sub-prefix hijacks. First, ROA records can be specified with an optional "maxLength" attribute that allows sub-prefix announcements up to the specified maxLength. If the maxLength of an ROA is longer than the prefix length used to announce that address block in BGP, sub-prefix attacks are still viable as an adversary can announce a longer prefix than the victim while still having the announcement compliant with the ROA [25].

Second, although AWS (the provider of Let's Encrypt's vantage points) performs ROV filtering [32], a "hidden hijack" can occur if AWS forwards validation traffic on a benign, RPKI-valid route to the victim, but an AS on this route does not perform ROV filtering and has the adversary's malicious sub-prefix route installed. Due to longest-prefix-match routing, this AS will forward packets via the adversary's route even though other ASes perform ROV filtering. We model the possibility of these hidden hijacks by computing the AS paths from all vantage points to all target IP prefixes and checking if each hop performs ROV according to a public dataset [39].

Formally, we consider an IP address i (with prefix p) immune to sub-prefix hijacks if it satisfies either of the following properties:

1. i's prefix length is 24
2. i's prefix is protected by RPKI, more precisely:

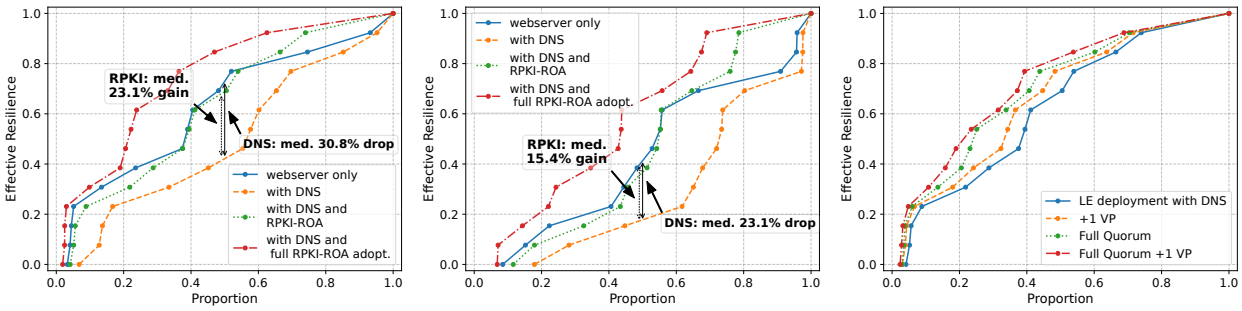

(a) (b) (c)

Figure 8: Resilience of vantage point deployments for serial hijacker experiments for **(a)** LE deployment and **(b)** singleVP deployment under various attack surface assumptions; **(c)** under current RPKI usage for potential expanded multiVA deployments.

Property	**Prop. of domains (%)**		
	A	NS	A and NS
1. Domains where all target IPs on /24 prefix	25.5	46.6	15.8
2a. Domains where all target IPs are covered by a valid ROA	61.8	63.7	44.7
2a & 2b. with properly-set maxLength	39.9	51.7	25.8
2a & 2b & 2c. with properly-set maxLength and full path ROV-filtering	28.9	42.4	18.2
Sub-prefix hijack-immune: 3. All target IPs on either /24 prefix OR ROA with correct maxLength and full path ROV-filtering	46.7	55.7	29.2

Table 7: Statistics on exposure to sub-prefix hijacks at the domain level with current ROV and ROA deployment.

Domains immune to sub-prefix attacks under	**Prop. of domains (%)**		
	A	NS	A and NS
universal ROV deployment	49.7 (+3.0)	63.1 (+7.4)	38.1 (+8.9)
correct maxLength	49.7 (+3.0)	51.3 (+4.4)	33.4 (+4.2)
correct maxLength and universal ROV	71.3 (+24.6)	78.0 (+22.3)	62.0 (+32.8)
correct maxLength, universal ROV, and universal ROA	100	100	100

Table 8: Resilience of domains, both absolute and amount of increase, under various routing security improvements.

(a) *i*'s prefix is covered by a matching ROA record AND
(b) if present, the maxLength attribute of the ROA is equal to the prefix length of *i* (i.e., the ROA is "minimal") AND
(c) for Let's Encrypt's vantage points, every AS on the path from the vantage point to *i*'s prefix either performs ROV, or all its neighbors perform ROV[7]

[7]Even if an AS itself does not perform ROV, it can be protected from "hearing" RPKI-invalid routes if all of its neighboring ASes perform ROV.

We apply these checks to the set of target A IPs, the set of target NS IPs, and the set of all target IPs (i.e., the union of A and NS) for each domain; the statistics are presented in Table 7. Currently, **29% of domains in our dataset are immune to sub-prefix attacks and benefit from the added security of multiVA, even in the absence of BGP monitoring or any other countermeasures for sub-prefix attacks**.

The fact that an adversary can attack either a webserver or a DNS server IP address to compromise domain validation substantially increases domains' vulnerability to sub-prefix hijacks (as it does with equally-specific prefix hijacks). When nameserver or webserver target IP addresses are considered independently, the fraction of domains immune to sub-prefix hijacks increases to 56% and 47%, respectively. This serves as potential justification for the use of CAA record extensions [37] that specify validation method (e.g., DNS validation) and eliminate the dependence of domain validation on communication with a domain's webserver.

Immunity to sub-prefix attacks can increase substantially as more of the Internet implements ROV filtering and ROA configuration practices improve. The proportion of domains immune to sub-prefix attacks increases to 38% if ROV is universally deployed. Correcting instances of ROA records with misconfigured maxLength with the existing ROV deployment makes 33% of domains immune. If both practices are implemented together (universal ROV and correct maxLength), that proportion more than doubles to 62% (see Table 8).

We note that security-conscious domain owners can immediately benefit from multiVA's full protections by ensuring that their infrastructure is hosted on either /24 prefixes or protected by properly configured ROA records. In the future, ubiquitous deployment of RPKI (including ROA records with minimal maxPrefix length for all IP prefixes along with full ROV enforcement) will negate sub-prefix attacks for all domains. In this scenario, only equally-specific BGP attacks will remain viable, vastly expanding the benefits of multiVA.