

Experiences Deploying Multi-Vantage-Point Domain Validation at Let’s Encrypt

Henry Birge-Lee, Liang Wang, Daniel McCarney, Roland Shoemaker, Jennifer Rexford, Prateek Mittal

Abstract

An attacker can obtain a valid TLS certificate for a domain by hijacking communication between a certificate authority (CA) and a victim domain. Performing domain validation from multiple vantage points can defend against these attacks. We explore the design space of multi-vantage-point domain validation to achieve 1) *security* via sufficiently diverse vantage points, (2) *performance* by ensuring low latency and overhead in certificate issuance, (3) *manageability* by complying with CA/Browser forum requirements, and requiring minimal changes to CA operations, and (4) a low *benign failure* rate for legitimate requests. Our open-source implementation was deployed by the Let’s Encrypt CA in February 2020, and has since secured the issuance of more than 300 million certificates. We show our approach has negligible latency and communication overhead, and a benign failure rate comparable to conventional designs with one vantage point. Finally, we evaluate the security improvements using a combination of ethically conducted real-world BGP hijacks, Internet-scale traceroute experiments, and a novel BGP simulation framework. We show that multi-vantage-point domain validation can thwart the vast majority of BGP attacks. Our work motivates the deployment of multi-vantage-point domain validation across the CA ecosystem to strengthen TLS certificate issuance and user privacy.

1 Introduction

Certificate Authorities (CAs) establish trust on the Internet by vouching for the identity of remote servers through digitally signed certificates. A certificate contains domain name(s) and the associated public key. The CAs must ensure that a certificate is only given to the legitimate owner of a domain. Otherwise, an adversary with a bogus certificate could circumvent the assurances of confidentiality and integrity offered by encryption, and then decrypt or modify sensitive user data (e.g., logins, financial information, cryptocurrency credentials [18]). To prevent these attacks, CAs perform *domain control validation* that requires the domain owner to demonstrate control of a core resource associated with the domain (e.g., a web server, email address, or DNS record).

1.1 Domain Validation Attacks and Defenses

Domain control validation is vulnerable to localized and targeted Border Gateway Protocol (BGP) attacks that allow adversaries to obtain bogus certificates [17, 27]. These




Figure 1: Domain control validation by Certificate Authority.

attacks are possible because validation is often performed over insecure HTTP connections. The BGP attack puts the adversary on the path between the CA and the victim domain or the victim domain and the CA. Then, the adversary asks the CA for a certificate for the victim domain. The validation request (e.g., an HTTP GET to the victim domain) is directed to the adversary where it is maliciously answered (Figure 1). With its validation request successful, the adversary can obtain a certificate for the domain. These attacks are particularly effective because the BGP attack is *localized* (affecting the target CA) and *short-lived* (during domain validation). These attacks are well within the capability of repressive regimes which have been accused of launching BGP attacks [22] and have motive to intercept encrypted communications [1].

To prevent these attacks, CAs need to defend themselves from routing attacks on domain control validation. A promising approach is to perform validation from *multiple* diverse vantage points, to make it hard for the adversary’s attack to “fool” all (or many) of the vantage points [17, 19]. Vantage points unaffected by the BGP attack reach the legitimate victim domain and observe that domain control validation has *not* been completed (see figure 2). This would stop the CA from issuing a certificate to the adversary. With effective multi-vantage-point validation in place, successful attacks would require announcing BGP routes with broad scope (readily visible in public BGP monitoring platforms [10, 11]), such as advertising smaller sub-prefixes (which is infeasible for /24 IP prefixes).

1.2 MultiVA Design, Deployment, & Analysis

This paper presents the design and evaluation of multiVA, the first real-world deployment of the multi-vantage-point countermeasure to secure domain control validation.




Figure 2: A localized BGP attack affects a portion of the Internet. If the CA has only one vantage point (A), the adversary successfully gets the certificate. With multiple vantage points (B), the CA detects the attack as two vantage points reach the legitimate server and fail the validations.

MultIVA design. Validating from multiple vantage points seems like a simple idea. Yet, creating a production-grade system is challenging, due to competing trade-offs:

- *Security.* The multiple vantage points must offer sufficiently diverse perspectives on routing to ensure that some vantage points can reach the legitimate domain. Also, the *quorum policy* (i.e., the “vote” among the vantage points before signing a certificate) must be strong enough to thwart attacks, without sacrificing performance and robustness.
- *Manageability.* Validating from multiple vantage points requires more server and network resources, spread across more locations. This may require billing arrangements with multiple cloud providers. In addition, the CA/Browser forum, which governs the operations of publicly trusted CAs, places security and auditing requirements on the data centers CAs use [21]. Thus, a system with multiple vantage points may require the maintenance and auditing of multiple data centers.
- *Performance.* The latency introduced by additional vantage points should not significantly slow down the overall domain control validation process. Similarly, the approach should have low communication overhead.

The implementation of multi-vantage-point domain validation should also be incrementally deployable. Performance constraints are particularly sensitive when deploying on a live production system, requiring careful system monitoring and a phased deployment.

- *Benign failure.* A benign (validation) failure is a non-malicious validation request that should have been successful but was blocked because of validation failures caused by external factors in some vantage points. The failures are mostly caused by DNS propagation delay and configuration errors; see §4.2. A multi-vantage-point validation system should not throw a significant number of benign failures.

We explore the complex design space of multi-vantage-point domain validation to balance the trade-offs among these challenges. We propose to use a deployment of multiple vantage points within a single cloud provider to achieve good performance and manageability, as the site reliability engineering (SRE) and billing departments only need to interface with a single cloud provider. We satisfy the compliance requirements imposed by the CA/Browser forum by carefully tracking validation results from the original CA and the cloud vantage points, respectively. We carefully select cloud vantage points across diverse geographic locations to ensure sufficient diversity and system security, and connect them to existing CA components using mutually-authenticated TLS. Our design balances the number and location of vantage points to control the trade-off between security, manageability, performance, and benign failures. More vantage points would improve security, but may increase validation overhead and manageability difficulty. Finally, we incorporate a configurable domain validation quorum policy to strike a balance between security and benign failures.

MultIVA deployment. We develop a fully open-source implementation of our multiVA design. Notably, our implementation does not require any changes to the Automated Certificate Management Environment (ACME) [15]. We build upon the Boulder ACME implementation [6] and only modified software components relevant to domain validation. Our open-source implementation was deployed by the Let’s Encrypt CA in its live production environment in February 2020. Since then, our multiVA deployment has secured the issuance of over 300 million TLS certificates, and validates domain control for approximately 1.5 million certificates a day [24]. Our work thus demonstrates the feasibility of the multiVA approach at Internet scale and represents a major step in strengthening the Internet PKI against BGP attacks.

Evaluating system performance and benign failures. We obtained operational data for the multiVA deployment from Let’s Encrypt and use it to analyze system performance and benign failures. We find that the system incurs negligible latency overhead since (1) validations from multiple vantage points occur in parallel and (2) validation time from well

connected cloud-based vantage points is much faster than the validation time from the existing CA vantage point. We measured the communication overhead of the deployment for typical certificate issuance rates as 0.5 Mbps per remote vantage point (far below the saturation point of 100 Mbps upstream links). Finally, we show that the rate of benign failures incurred by the multiVA deployment is around 1%. These benign failures are typically caused due to DNS propagation delays and configuration issues, and typically can be overcome by retrying the certificate issuance request.

Evaluating system security. Unlike performance metrics like latency and benign failure rate which can *only* be measured with real data from an active deployment, the security offered by multi-vantage-point domain validation cannot be understood from deployment data alone. Many important security questions, like the fraction of attacks that can be mitigated, quorum policies to use, and placement of additional vantage points, cannot be answered solely by relying on deployment data, and instead require combining multiple analytic perspectives.

Our approach is to analyze multi-vantage-point validation *holistically* across many different analytical frameworks including ethically conducted real-world attacks, deployment metrics, Internet-scale traceroute measurements, and novel BGP attack simulations. By comparing the data produced in these different settings, we can reason holistically about how to optimally evaluate and deploy multi-vantage-point validation. We show that our multiVA deployment that uses a single cloud provider can successfully mitigate a vast majority of BGP attacks. Finally, we also make recommendations for further enhancing the multiVA deployment via additional diverse vantage points.

Our work motivates the industry-wide deployment of the multiVA approach in the CA ecosystem to strengthen the PKI and protect user privacy.

2 Let’s Encrypt Certificate Management

Let’s Encrypt is an automated CA that requires no human interaction to request or renew a certificate. Thus, Let’s Encrypt simply exposes a standardized API—the Automated Certificate Management Environment (ACME) [15]—for third-party clients (e.g., EFF’s Certbot [7]) to access.

The Boulder ACME implementation. Let’s Encrypt is powered by an open-source ACME implementation called Boulder [6], designed for security, reliability, and performance. To date Boulder has issued over a billion trusted certificates.

Boulder is subdivided into components based on their role in the certificate issuance process and to minimize the attack surface between components and the global Internet. Each component is designed to be deployed in isolated network segments with strict firewall rules limiting ingress/egress traffic based on the API exposed by the components.

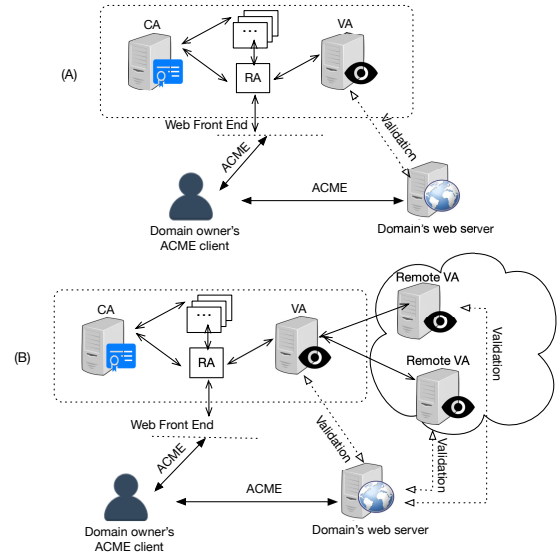


Figure 3: Boulder with a single VA (A) and multiple VAs (B).

As in Figure 3 (A), the primary system components of Boulder include Web Front End (WFE), Registration Authority (RA), Validation Authority (VA), and Certificate Authority (CA). We omit additional components not specific to validation/issuance (e.g., storage). Of all the components in Boulder, only the web front end requires inbound traffic from the wider Internet allowed for the ACME API interface. Similarly only the VA requires outbound traffic to arbitrary hosts on the wider Internet to perform domain validation.

During certificate issuance ACME clients interact with the web front end to submit a to-be-signed certificate, following the process described in the ACME RFC [15]. The web front end interacts with the resource authority to associate accounts with authorizations and other resources. The resource authority in turn interacts with the VA to request domain validation. The VA performs the validation and sends the validation result to the resource authority. The resource authority then asks the CA to sign the certificate and returns the signed certificate (or error messages if the validation fails) to the ACME client through the web front end.

Single-VA domain control validation. During domain validation, Let’s Encrypt challenges the client to demonstrate its ownership of the domain(s) requested in the certificate. When the client is ready for the challenge and asks Let’s Encrypt to initiate the domain validation challenge, the request is directed to one of the available data centers of Let’s Encrypt by a CDN layer in front of Boulder. A randomly selected VA instance within the data center performs the domain validation. Note that though there could be many VAs in different data centers, only one VA is selected to perform domain validation for a given domain.

The VA resolves the requested domain to IP addresses using a recursive DNS resolver colocated with the VA. The VA performs HTTP, DNS, or TLS based domain control validation as specified by RFC 8737 [39]. In HTTP and TLS based validations, the VA uses DNS to look up an A or AAAA record for the requested domain, and initiates an HTTP or TLS connection to the domain’s web server. In DNS based validations, the VA checks for the validation response in a DNS TXT record for the requested domain.

Overall, in a CA service, certificate issuance involves complex interactions between different system components. A minor modification to the ACME protocol standard or the existing Boulder implementation may affect the reliability and security of the CA. Therefore, when designing the multiple vantage point validation system, we want to modify neither the ACME protocol nor the non-VA components in Boulder, and minimize the changes to the existing VA component.

3 Multi-Vantage-Point Validation Design

Below we outline our design of multiVA, an incrementally deployable domain validation system that leverages multiple vantage points to mitigate BGP attacks against CAs. We first discuss how our design complies with CA/Browser Forum policies without requiring auditing on remote data centers. Then, we discuss how we simplify management by using a single cloud provider to host vantage points. Next, we present our method for scalable and secure communication with vantage points. We also introduce our configurable quorum policy that can balance the trade-off between security and benign failures. Last, we present our phased deployment strategy that Let’s Encrypt used to deploy our design in their live production environment.

Satisfying CA/Browser Forum compliance without sacrificing manageability. The CA/Browser Forum governs the operations of publicly trusted CAs, and imposes requirements (e.g., physical security and access requirements) on CAs that may constrain the deployment of multi-vantage-point validation (see CA/Browser Forum Baseline Requirements [21] Section 5). Under a naive system, all the validation authorities (or VAs) would require independent security audits, increasing the cost of adding VAs. The cloud provides a logical location for additional VAs that can significantly improve manageability while offering many beneficial features such as geographic diversity. However, independent security audits and physical access to the cloud infrastructure may not be easily obtained.

To resolve this conflict, multiVA utilizes the concept of primary and remote validation authorities, as shown in Figure 3(B). A primary VA is located in a data center that is fully compliant with the CA/Browser Forum Baseline Requirements and, based on the current requirements, is authorized to validate a certificate independently without any

information from other vantage points. For a CA who plans to upgrade to multiVA, the primary VA would be the existing VA being used by the CA. A remote VA can be deployed in a network or machine managed by a third party (e.g., the cloud). As a key deployment insight, multiVA requires that:

A primary VA’s validation must succeed for a customer’s validation request to be successful, regardless of the responses of the remote VAs.

Thus, all certificates signed under multiVA are a *subset* of the certificates that would be signed without multiVA.

Using a single cloud provider for manageability. Even without security audits or physical access to remote data centers, using multiple data centers with different cloud providers would be difficult for a CA to manage. Each cloud provider has different billing policies and requires a separate cost analysis, along with requiring different tooling for the Site Reliability Engineering (SRE) team. We resolve this by hosting all multiVA VAs in a single cloud provider: Amazon Web Services (AWS). To justify this approach, we performed an in-depth security analysis for the AWS-based multiVA (see §6) and found that, despite using only one cloud provider, multiVA *does* achieve a significant level of route diversity. This result significantly improves the manageability of multiVA and reduces the implementation burden of multi-vantage-point domain validation.

Communicating securely with remote validation authorities. Each component in multiVA is associated with a certificate bound to a specific role (e.g., primary or remote VAs), and communicates with other components using *gRPC* [8] over mutually-authenticated TLS for confidentiality, authorization, and authentication. By examining certificates and the associated roles, a component can confidently determine the legitimacy of components it communicates with. For instance, the primary VA will not accept fraudulent communications from nodes pretending to be remote VAs.

Besides, using *gRPC* saves round trips between the primary and the remote VAs. In the *gRPC*-based setup, the primary VA sends one RPC message to the remote VA which in return sends back a validation result. One validation only involves one round trip of communication between the primary and remote VAs. An alternative choice is VPN, but it may incur additional round trips and introduce more latency than *gRPC*. We demonstrate that multiVA introduces negligible latency compared to single-VA domain validation in §4.

Balancing security and benign failure rate with a configurable domain validation quorum policy. A quorum policy specifies the number of remote VAs required to agree on the validation result before signing a certificate. While requiring domain validation to succeed at a large number of remote VAs would enhance security, it would also increase the risk that certificate requests get rejected incorrectly, since in practice a non-trivial number of challenge requests may fail




Figure 4: Cumulative number of issued certificates on each day since multiVA was deployed with domain exception list in February 2020; the blue line shows when full enforcement was enabled.

due to uncontrollable factors (e.g., DNS propagation delay and configuration errors).

Our multiVA system adopts a configurable quorum policy that enables the CA to strike a balance between different trade-offs. Given n remote VAs, the k - n quorum policy allows $n - k$ remote VA challenges to fail to return an answer or return an inconsistent answer during a validation, while still allowing certificate issuance (assuming the primary VA has successfully validated the domain). The initial deployment uses $n = 3$ and $k = 2$, i.e., allowing at most one of three remote VAs to fail. Our analysis shows that this setup can effectively improve the security against domain validation attacks (§5 and §6) with an acceptable benign failure rate (§4.2).

Open-source implementation and phased deployment of multiVA at Let’s Encrypt. We developed and released an open-source implementation of multiVA that any CA can use. Our implementation does not require any changes to the ACME protocol [15] and only modified software components relevant to domain validation in the Boulder ACME implementation [6]. Our implementation of the multiVA design was deployed by Let’s Encrypt, a non-profit which is the world’s largest CA. Let’s Encrypt set up three remote vantage points in three AWS datacenters: Oregon (us-west-2), Ohio (us-east-2), and Frankfurt (eu-central-1). Let’s Encrypt’s existing primary VAs are located in two data centers in Denver and Salt Lake City, managed by Flexential [25]. Since Let’s Encrypt issues millions of certificates a day, changing their certificate issuance path required significant care. We collaborated with Let’s Encrypt to develop and apply a multi-stage deployment plan:

- *Staging deployment:* Let’s Encrypt deployed multiVA in a staging environment, which is a functional duplicate of the production environment and is used for internal testing of new features as well as external testing by ACME client developers.

- *Testing in production environment:* Then, Let’s Encrypt introduced multiVA to the production environment in a data-collection mode: the remote VAs performed domain validations, but did not affect the primary VA’s validation decisions. Let’s Encrypt collected detailed information on each validation from all VAs to understand the potential causes of validation failures and performance bottlenecks, which helped us refine our implementation to handle the load (e.g., during traffic spikes) generated by the full volume of production ACME requests.
- *Production deployment with domain exception list:* Next, Let’s Encrypt applied multiVA for most of the domain validation requests, with a domain exception list to temporarily exclude certain domains from multiple-vantage-point validation. Let’s Encrypt populated the exception list with domains that may, based on the logs, have trouble with multi-vantage-point validation to prepare for future certificate renewal. Let’s Encrypt only did this for the domains associated with ACME accounts that have contact information, so it can communicate with the domain owners to inform them that they are on the domain exception list, and troubleshoot the issues.
- *Full production deployment:* And finally, Let’s Encrypt has a completed deployment of multiVA and all certificate requests are now validated by multiple VAs.

Phased deployment minimizes the impact on the ongoing operations of the CA, while allowing us to gradually improve the performance and reliability of the system.

As shown in Figure 4, the multiVA deployment has issued over 300 million certificates since production deployment in February 2020. Our work thus demonstrates the feasibility of the multiVA approach at Internet scale.

4 Real-world Deployment Evaluation

We obtained and analyzed log data for performance and certificate validation from Let’s Encrypt’s production deployment of multiVA. Our analysis demonstrates the viability of multiVA at scale with respect to both performance and benign failures.

4.1 Domain Validation Performance

Latency. In the k - n quorum validation scheme that Let’s Encrypt adopted, we do not expect performance to degrade significantly because remote validations are performed in parallel using the same timeout as the local validation by the primary VAs. The only additional latency come from the RPC round trip between the primary and remote VAs.

Total validation latency is bounded by the $(n - k)^{th}$ slowest remote VAs in the quorum majority. For instance, when only




Figure 5: Let’s Encrypt certificates issued per day since multiVA has been deployed with domain exception list. The blue line indicates when full enforcement was enabled.

two of three remote VAs are required to succeed (i.e., with a 2-3 quorum policy), one slow remote VA would not increase the overall latency. *When choosing a quorum threshold for remote VAs, some consideration must be given to how many slow, or entirely unresponsive, remote VAs the system can tolerate.* If the number of slow VAs makes up a quorum majority then the system performance would degrade.

When the CA chooses remote VAs that have similar performance characteristics as their primary VAs, there should be little change in validation performance. This is seen in the performance data from the Let’s Encrypt deployment between June 2nd and Aug 31st (see Figure 6). In this deployment we see that the remote VAs provide somewhat *better* performance than the primary VA. This is likely due to the fact that primary VAs are located in the data centers of a colocated hosting provider, while the remote VAs are in AWS that has richer peering relationships that provide more performant routing paths. We also see that the RPC overhead contributes minimally to the total latency. So, for the Let’s Encrypt deployment, latency is typically determined by the primary VA (as it would be in the absence of multiVA).

Bandwidth overhead. The multiVA RPC scheme does introduce a small amount of traffic for each validation request. For an issuance rate of about 20 certificates per second, the RPC traffic amounts to around 0.5 Mbps per remote VA. This allows for using a rather large number of remote VAs before the traffic overhead would come anywhere near the saturation point for most network uplinks, which typically have a minimum bandwidth of 100 Mbps.

Certificate issuance rate. Figure 5 shows the daily certificate issuance rate for the Let’s Encrypt multiVA deployment. The vertical blue line indicates when full enforcement of multiVA without the domain exception list was enabled. We can see that the multiVA deployment is able to handle the load of the world’s largest CA and scales to millions of daily certificates. The issuance rate also shows a stable trend over several months, including after the removal of the domain exception

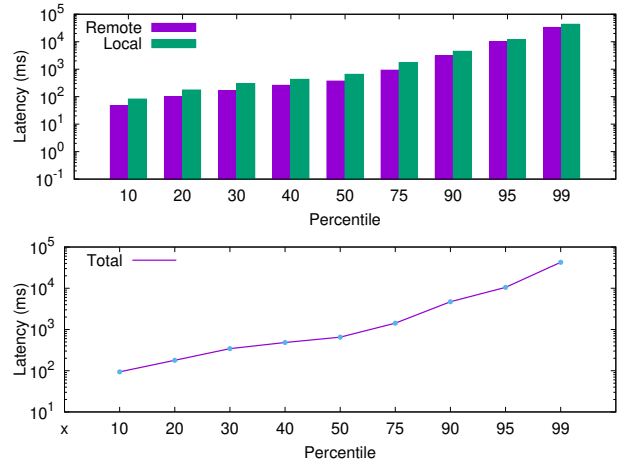


Figure 6: Let’s Encrypt validation latency in milliseconds, with y-axis on log scale. Remote latency is the time it takes to complete a validation request at the remote VA (once it is sent with gRPC); local latency is the time taken to complete the request at the primary VA; and total latency is the time taken to complete the remote validation requests with the gRPC overhead. Since total latency is comparable to the local latency, the overhead for multiVA is minimal.

list and the transition to the full enforcement mode (June 2020).

4.2 Benign Failure Rate

In this section, we study the impact of multiVA on benign failures, i.e., the domain validation failures that are caused by uncontrollable external factors,

Certificate validation dataset. We used the certificate validation dataset from Let’s Encrypt to analyze benign failures. This dataset includes the information of 451 M certificate validations collected from the primary VA (which records the remote validation responses) of Let’s Encrypt over a 20-day period (Sep 3-22, 2020), with detailed runtime data on each validation, including the domain name (e.g., www.example.com) being validated, the challenge type in use (HTTP-01, TLS-ALPN-01, or DNS-01), the ACME account ID that initiated the domain validation, the reason for the failure, and the validation result of each remote VA.

Benign failures caused by multiVA. The data from Let’s Encrypt shows that the primary VAs have a validation failure rate of more than 65%. This high failure rate is due mainly to repeated failed requests for a small portion of the domains. If the primary VA succeeds and *any remote VA* fails (e.g., returning no/incorrect/inconsistent responses), we say there is a *domain validation differential* or differential for short. Only 1.2% of the validations caused differentials (i.e., failed due to remote VAs) in the 20-day period.

Reason	% Differentials
DNS	52.6
Connection	24.3
Unauthorized	21.1
ServerInternal	2.0
Total	100.0

Table 7: Differentials by reason for validation failure.

Causes of differentials. Let’s Encrypt has contacted certain domain owners based on the ACME account information, including hosting providers serving millions of customer domains whose certificates have caused differentials, to understand the potential causes. Some issues may cause validation failures in both the primary and remote VAs (e.g., misconfiguration of web servers). We only focus on those issues specific to multiVA in this section.

Based on the certificate validation dataset, the causes of differentials fall into four categories: DNS-related issues, connection-related issues, HTTP unauthorized errors, and server internal errors, as summarized in Table 7.

- *DNS.* For 52.8% of the differentials, the remote VAs failed to resolve the validated domain names to IP addresses. Most commonly we saw that the multiple authoritative DNS servers responsible for a domain were not synchronized to serve the same content, causing different VAs to receive different answers. A domain owner cannot always determine the synchronization state of their DNS zone globally (e.g., if the DNS provider uses anycast routing and does not provide an API to query the state of a zone across all DNS servers). Thus, when the owner requests a domain name for a new website around the same time as asking the CA for a certificate, the VAs may perform validations before the DNS synchronization has completed, and thus see a different DNS zone state from what the ACME client sees.

- *Connection.* We find that 24.3% of the differentials were caused by connection-related issues (i.e., the HTTP connection for validation gets blocked or dropped). We found that one cause is overly sensitive DDoS mitigation. By its nature multiVA generates a burst of traffic from multiple distinct sources—similar to a small-scale DDoS attack. If a domain owner uses a web hosting or DNS hosting service that employs DDoS mitigation tuned at a low threshold, multiVA traffic could be classified as a DDoS attack and dropped. If the number of dropped requests exceeded our configured threshold, our quorum policy would report a validation failure.

Similarly we found that some domain owners employed firewall policies that block traffic not sent by the Let’s Encrypt data center’s IP address block. Thus, validation requests from the remote VAs were blocked, affecting the quorum and preventing domain issuance. Let’s Encrypt has never published a list of source IP addresses for use in firewall policies, to avoid constraining operational agility and the roll-

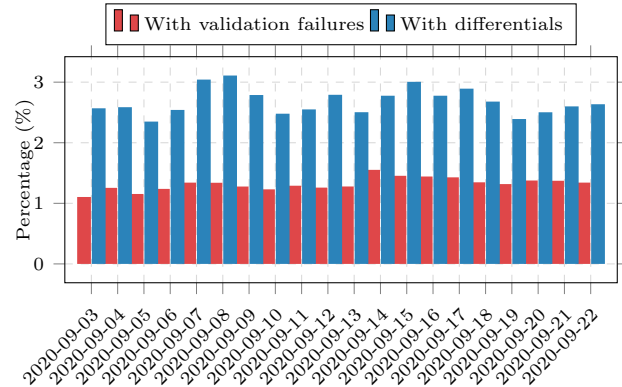


Figure 8: The fractions of certificates that were associated with differentials and certificate validation failures over time.

out of new features. Despite consistent advice from Let’s Encrypt to prefer DNS-based challenges in environments that require strict source IP address filtering for inbound HTTP/HTTPS, some domain owners have assembled their own (incomplete) static lists of Let’s Encrypt IP addresses.

We also saw that increasing the number of validation requests would increase domain validation failure rates for some domains hosted on underprovisioned, low-end shared hosting platforms. In these cases even making several concurrent challenge requests could overwhelm the servers responsible for replying to the requests. This would frequently manifest as the first one or two validation requests completing successfully, with subsequent requests resulting in a timeout or error page response.

- *HTTP unauthorized.* For about 20.9% of the differentials, the remote VAs successfully communicated with the target web/DNS server but did not receive the expected response (e.g., the validation document had not been uploaded). One commonly occurring explanation is some customers may migrate to a new hosting provider that has the proper content for validation but, once again because of DNS propagation delay, the remote VAs are given DNS records for the old hosting provider where the domain control validation content has not been uploaded.

- *Server internal errors.* A small fraction (2.0%) of the differentials were caused by RPC call failures between the primary and remote VAs.

Effectiveness of quorum policy on reducing certificate validation failures. We further find the certificates that caused the differentials at remote VAs. Due to the use of quorum policy, differentials may not necessarily cause a *certificate validation failure* (i.e., certificate rejection). For instance, when employing the 2-3 quorum policy as Let’s Encrypt, even if one remote VA fails to validate every domain belonging to a certificate, multiVA would still consider the certificate validation a success (if the primary VA has completed the validation successfully) and issue

the certificate. To demonstrate the effectiveness of quorum policy on reducing certificate validation failures, we show the fractions of the validated certificates that were associated with differentials and certificate validation failures on each day in Figure 8. The quorum policy has reduced the daily certificate validation failure rate by 50% on average.

Though each day about 1% of certificates were temporarily rejected due to differentials, most domain owners have retried validation and eventually had their certificates signed by Let’s Encrypt after retrying by the end of Sep 22, 2020. Overall, only 0.65% of all certificates that have been validated in 20 days (36.8M) failed to get signed because of differentials.

Considering the high failure rate of the primary VAs (over 65% without multiVA), the remote VAs have minor impact on the benign failure rate. Most of the failures caused by the remote VAs can be resolved by retrying or having domain owners whitelist the IP addresses of remote VAs.

5 Evaluating the Let’s Encrypt MultiVA Deployment Against Real-World BGP Attacks

In this section, we demonstrate the ability of multiVA to mitigate *ethically* launched real-world BGP attacks. This methodology allows us to understand how multiVA interacts with real Internet routing. Since we can only launch BGP attacks from limited number of locations, we complement this methodology with simulated attacks in Section 6.3 which yield concurring results.

First, we verify that multiVA properly distinguishes between the primary and remote VAs, ensuring that all certificates signed under multiVA are a subset of what would be signed without multiVA in place. Second, these measurements show that an AWS-only deployment detects the vast majority of attacks (up to 94% of attacks) even though all remote vantage points are hosted by the same cloud provider. The security improvements are more significant for certain domains. Furthermore, the small number of BGP attacks that are able to fool many vantage points are highly visible and, as such, can be mitigated by other BGP attack-prevention methods. We discuss how to further enhance security by adding more vantage points (see §6.3).

5.1 Launching Ethical Attacks

We launched attacks using the PEERING platform [38], which allows us to make real-world BGP announcements. Our experimental setup consisted of one adversary server and one victim server. Each server was connected to a PEERING mux [38] through a VPN tunnel which it used to forward packets and make BGP announcements. Our experiments consisted of first the victim announcing its prefix through a designated PEERING mux. Next, the adversary subsequently hijacks that prefix from a different PEERING mux (see Figure 9) and requests a certificate from Let’s Encrypt for the




Figure 9: The experimental setup with the PEERING testbed. The victim makes BGP announcements via the University of Wisconsin - Madison (wisc01) mux, while the adversary hijacks the prefix using the North Eastern University (neu01) mux. With this setup, we measured that the Salt Lake City and Denver data centers (along with the Frankfurt remote vantage point) routed validation requests to the adversary, but the Oregon and Ohio remote vantage points were able to reach the legitimate domain and, hence, thwart the attack.

victim’s domain (that has an A DNS record within the victim’s prefix). Our experiments focus on BGP attacks with equally-specific prefixes (where the victim and adversary announce the same IP prefix) because sub-prefix attacks are not viable against all domains (like those running on a /24), are easily noticed by route monitors due to global distribution [10, 11], and are mitigated by deployed BGP security proposals [20].

Ethical considerations. We took several steps to ensure our attacks were ethical. The domains we used had *no real users* and were pointed to IP addresses in the prefixes allocated to us by the PEERING framework (that similarly ran no network services other than those needed for the experiment). We only requested certificates for these domains that were registered for the express purpose of conducting our experiments. We also followed all of the policies and guidelines of the PEERING framework including not announcing any prefixes other than the prefixes allocated to us and not spoofing packets from IP prefixes outside of the PEERING framework range.

5.2 Example of Successful Attack Mitigation

To demonstrate how multiVA mitigates real-world BGP attacks, we used the University of Wisconsin - Madison (wisc01) mux for the victim’s domain and North Eastern University (neu01) mux as the adversary, as shown in Figure 9. We started by making a BGP announcement from the victim’s server for an IP prefix we controlled. We then had the adversary server hijack the victim’s IP prefix by announcing it to neu01 in an equally-specific prefix hijack attack. Then, we used the adversary server to request a certificate for

the victim’s domain. **Let’s Encrypt did not authorize the certificate because the remote vantage points caught the attack.** Additionally, system logs from the adversary server indicated that the primary data centers were routing data to the adversary and this attack would have succeeded without multi-vantage-point domain control validation.

5.3 Characterization of Attack Mitigation

Considering other potential adversaries. To expand beyond a single (victim, adversary) pair, we considered the set of all possible adversaries (i.e., other PEERING muxes) attacking the victim domain hosted by wisc01 to find how many attacks were prevented by multiVA. **When Let’s Encrypt only uses a single data center (i.e., without the multiVA approach), the victim’s domain was only resilient to one in six attacks (17%) we launched against it. With multiVA in place, the victim’s domain was resilient to five of the six attacks (85%).** These additional attacks were detected by Let’s Encrypt’s AWS-hosted remote vantage points using the k - n quorum policy.

Considering a broader set of domains. Next, we further varied our attacks by selecting different available PEERING muxes for the victim and the adversary. Overall, we launched attacks from 62 different experiment configurations, and analyzed multiVA security with different quorum policies.

With Let’s Encrypt’s quorum policy, 67% of the attacks we launched failed to obtain certificates (i.e., the domain was resilient to attack). This is significantly higher than the 48% of attacks mitigated using Let’s Encrypt’s previous single-data-center configuration. We note that the multiVA deployment has a larger impact on some of the weakest domains which are most vulnerable to BGP attacks (like the victim domain in wisc01 explored above that saw a fivefold improvement). In fact, ten of the attacks detected by multiVA under Let’s Encrypt’s quorum policy were against three highly impacted victim domains that on average saw fraction of attacks mitigated increase from 20% to 72%. We explore this trend in greater depth in Section 6.3.

While Let’s Encrypt’s current quorum policy offers substantial improvement for the most vulnerable domains, a full quorum policy is more effective at protecting the average domain. **58 out of 62 attacks (94%) were detected by multiVA with full quorum policy of $k=3$.** This result demonstrates that, even if a single cloud provider is used to host all remote VAs for ease of management, multiVA can significantly reduce the attack surface of BGP attacks.

The four attacks where the Let’s Encrypt vantage points were unable to reach the victim were attacks where the adversary significantly overpowered the victim from the perspective of network connectivity. For example, two of these attacks used the PEERING mux in Amsterdam as the adversary. This location has substantially richer connectivity

than any other PEERING mux (it is the only mux with two providers and it has substantially more peers than any other mux). While these results shed light on the limits of multiVA, in that certain AS-level adversaries may still succeed, overall, our approach significantly reduces the number of viable BGP attacks against domain validation. Furthermore, the number of attacks that still succeed can be reduced by simply deploying a small number of additional vantage points (as discussed in Section 6.3).

6 Quantifying the Security of multiVA

While our evaluation in Section 5 used real-world BGP attacks, we were limited to considering PEERING muxes as target domains and adversary locations (which may not be representative). In this section we evaluate the security of the multiVA deployment with respect to (i) real-world domains served by Let’s Encrypt and (ii) any AS-level adversary. We use a combination of Internet-scale traceroutes and Internet topology simulations to evaluate system security. We focus our security analysis on the following questions:

- How effective is Let’s Encrypt’s current quorum policy and set of vantage points at catching localized BGP attacks on domain control validation?
- How many more attacks would be caught under a full/strict quorum policy that uses results from all three remote VAs?
- How much would additional vantage points enhance security, and where should these remote VAs be located?

6.1 Evaluation Methodology

We first introduce the primary dataset we collected to facilitate our analysis, and discuss the analysis techniques we used.

Domain dataset. By parsing log data shared with us by Let’s Encrypt, we collected 47 million domains seen in the Subject Alternative Name (SAN) field of certificates issued by Let’s Encrypt between April 13th and May 13th, 2018. For each domain, we also collected the IP address Let’s Encrypt used for HTTP and TLS-based domain control validation. 43 million domains used HTTP and TLS-based domain control validation to IPv4 and were used in our final data set.

An overview of evaluation methodology. To strike a balance between capturing the dynamics of real Internet routing and the flexibility of simulations, we used two different techniques, traceroutes and Internet-topology simulations, to evaluate the security (i.e., the fraction of attacks detected by the system) provided by Let’s Encrypt’s vantage points:

- **Traceroutes:** We ran traceroutes to the IP addresses of the domains in the data set from the three AWS locations

used by Let’s Encrypt. We recorded the percentage of AS-level paths that were identical across the data centers. These measurements allow us to study the routes to real domains and weigh our results based on the number of domains that are reached via a given route.

- **Internet-topology simulations:** We ran simulations of BGP hijacks against domains on an inferred Internet topology, augmented with information about differences in connectivity across AWS data centers. This allowed us to better understand the location of Let’s Encrypt’s vantage points relative to both the locations of domains using the PKI and potential sources of BGP attacks.

The traceroute experiments allow us to study real Internet routing to real domains (albeit in the absence of BGP attacks), while the simulation experiments allow us to study the effects of BGP attacks (albeit under an inferred Internet topology). Together, these results help us understand the effectiveness of multiVA across a diverse set of source vantage points, destination victim domains, and possible adversaries.

6.2 Traceroute Path Diversity of Remote VAs

Next, we assess the path diversity of vantage points with respect to the real-world distribution of domains. We perform traceroutes from Let’s Encrypt’s vantage points towards the domains in the dataset and use traceroute similarity across vantage points as a metric for path diversity.

Of the 43M domains in our domain dataset, we performed traceroutes to a randomly chosen 250K (0.6%) sample. For a given domain, we traced the route to the domain’s IP address used by Let’s Encrypt. There were 67K unique IP addresses in total. We performed these traceroutes from the three Let’s Encrypt remote vantage points (Oregon, Ohio, and Frankfurt). We resolved each IP address seen in the traceroute results to an ASN using 1) the originating AS of that IP address in the current global BGP routing table (compiled from RIPE NCC RIS [10] and Routeviews [11]), followed by 2) the originating AS listed in the whois record for that IP address (for IP addresses that were not currently routeable). We filtered the traceroute results to exclude any traceroute that did not have any IP addresses resolve to an ASN other than the ASN of Let’s Encrypt’s vantage points. We were left with valid traceroutes to 192K domains.

For each domain, we measured traceroute similarity between the different vantage points by seeing if the traceroutes from the different vantage points had the same AS-level forwarding path. We classified each domain as either having similar or different AS-level paths.

A single cloud provider can offer path diversity. Table 10 outlines our results. We find that for 80% of domains with traceroute information were reached via different AS-level forwarding paths from different vantage points. Even for

VP comb.	#Valid domain	Distinct paths	Same path
All three	191,653	153,048 (80%)	38,605 (20%)
Ohio-Oregon	197,054	119,914 (61%)	77,140 (39%)
Oregon-Frankfurt	196,576	138,202 (70%)	58,374 (30%)
Ohio-Frankfurt	196,369	140,815 (72%)	55,554 (28%)

Table 10: The number of domains with different levels of traceroute similarities for each pair/trio of vantage points. Percentages are taken in comparison to number of domains with full traceroute info (i.e., “Valid domain” in the table).

the 20% of domains that have the same paths, intra-AS routing differences might still allow vantage points to route independently in the event of a BGP hijack.

Ohio and Oregon had significantly more domains using similar paths than Ohio and Frankfurt or Oregon and Frankfurt. This result supports the importance of geographic diversity in vantage points and explains the security improvements seen in §5.3 for the full quorum policy.

6.3 Simulating BGP Attacks on Domains

Next, we performed simulations of BGP hijacks using the CAIDA AS-Relationship data set [2] to measure the impact of multi-vantage-point domain control validation on preventing BGP attacks against domains. In contrast to our real-world attacks (that had limited locations for domains and adversaries), simulations let us consider attacks from 1000 different randomly-sampled adversary ASes against real Let’s Encrypt domains. In addition, we consider alternate vantage points that have not been deployed by Let’s Encrypt.

6.3.1 Effective resilience

We use *effective domain resilience*, in contrast to traditional AS-level resilience [30], to measure the resilience of a domain against domain validation attacks launched from a set of potential adversary ASes.

For a given domain name d whose IP address is i , we assume a set of adversaries \mathcal{A} that each control a single AS and aim to obtain a bogus certificate for d . For each adversary a in \mathcal{A} , we perform an Internet-topology simulation of an equally-specific BGP hijack by a against the IP prefix containing i . We use the result of this simulation to compute $\alpha(a, d, v)$, which indicates whether a is capable of launching a successful BGP attack to hijack traffic from a given VA v (selected from a VA set \mathcal{V}) to d . We define

$$\alpha(a, d, v) = \begin{cases} 0, & \text{If the BGP attack launched from } a \text{ fails to} \\ & \text{hijack Internet traffic from } v \text{ to } i \\ 1, & \text{otherwise} \end{cases}$$

Next, we take quorum policy into consideration. The quorum policy q is a function that takes the subset of vantage points an adversary can hijack traffic from (which is a subset of \mathcal{V}) as an input and outputs either a 1 or a 0 depending on whether this subset of vantage points is sufficient to sign a certificate. An output of 1 implies the adversary hijacked traffic from enough vantage points and the validation request was successful and a 0 implies the attack did not lead to a mis-issued certificate because validation did not succeed at enough vantage points.

CAs may instantiate the quorum policy in different ways. We primarily consider Let’s Encrypt’s quorum policy, which can be expressed as:

$$q(\mathcal{W}) = \begin{cases} 1, & \text{If (primary VA} \in \mathcal{W}) \\ \text{AND } (|\{\text{remote VAs}\} \cap \mathcal{W}| \geq 2) \\ 0, & \text{If otherwise} \end{cases}$$

We use $\alpha^+(a, d, q)$ to denote whether a is capable of launching successful BGP attacks against the CA under the quorum policy.

Then we have

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

Finally, we define the effective resilience for a domain d as

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

The effective resilience of a domain is affected by the quorum policy, and the number and location of adversary ASes and remote VAs. In our simulation, we vary each parameter to understand how each factor affects the resilience to shed light on the possible directions for improving multiVA to be more robust against domain validation attacks.

6.3.2 Novel prefix-level simulations of BGP attacks

Our simulations are based on modeling equally-specific prefix attacks with the Gao-Rexford model of AS routing preferences [26] and are ostensibly similar to previous simulation work [30] but with several significant improvements. First, unlike previous approaches which model a cloud provider with multiple data centers as a homogeneous entity, our simulations are at the finer granularity of IP prefixes. Second, we augment the CAIDA AS relationship data [2] with AWS’s upstream links inferred from the bdrmap tool [33]. Finally, we populated victim domains based on the real-world domain dataset.

AS-level simulation fails to capture routing diversity. Prior work on Internet topology simulation simulates routing at the granularity of ASes, and considers each of the geographically-distributed cloud (or content) providers as a

single homogeneous network [17, 30]. However, such coarse-grained simulation and oversimplification fails to capture the path diversity of cloud providers. Many major cloud providers use a single AS number for all of their globally-distributed data centers, even though each of those data centers has a different set of neighboring ASes. For example, the routes to all AWS data centers are announced via AS 16509. Furthermore, AWS’s documentation explains that not all AWS IP prefixes are announced at all points of presence (instead only local IP prefixes are announced in each region) [5]. CloudFlare also has a similar setup with AS 13335.

Measuring the diversity of different data centers within the same AS is crucial as Let’s Encrypt’s remote vantage points are all deployed in AWS. AS-level simulations would inaccurately count Let’s Encrypt’s three AWS vantage points as if they were a single location, and thus fail to capture the resulting security benefits of multiVA.

Improving simulation accuracy with prefix-level Internet-topology simulation. To address the issues in the AS-level simulation, we develop a novel finer-grained simulation framework that operates at the granularity of IP prefixes and can more accurately model the routing behaviors of cloud providers.

For a cloud provider that has a similar routing setup to AWS, we use a combination of BGP data [10, 11] and the bdrmap tool [33] to construct a unique list of peers/providers for each of its data centers. Then, we use BGP data to observe which providers’ AS numbers are being used for specific prefixes, and simulate those prefixes as only being announced through those providers. This allows us to capture how different data centers (of the same cloud provider) select different BGP routes for the same destination prefix.

We also consider AS-path prepending [34] in our simulations. Recent work has shown that AS-path prepending, where an AS intentionally lengthens the AS-path it announces to certain neighbors, has a substantial negative impact on the resilience of IP prefixes against real-world BGP hijacks [34]. To capture the AS-path prepending behaviors, we measure how many times the origin AS for a given IP prefix prepends its announcement to each provider, and apply the same provider-specific prepending when simulating the announcements of each IP prefix.

We applied these prefix-level simulations to the prefixes for the IP addresses seen in our domain data set to model BGP hijacks against Let’s Encrypt customer domains.

6.3.3 Security evaluation results

We compare the distribution of domain resilience under multiVA to several different alternative deployments and a single-vantage-point deployment under a range of simulation parameters. Our analysis shows that Let’s Encrypt’s current multiVA deployment greatly improves the security of the CA over the status-quo single-VA system. **The current**




Figure 11: The effective resilience of domains (ordered by percentile) under 1) no multiVA, 2) the current deployment, 3) two additional vantage points 4) a full quorum policy 5) both two additional vantage points and a full quorum policy. The effect on the median and 10th percentiles is marked.

system achieves the goal of making the vast majority of ASes on the Internet (>90%) topologically incapable of launching BGP attacks against the majority of domains. Furthermore, if we consider the domains most vulnerable to BGP attacks (i.e., the bottom 10th percentile), multiVA shows a five-fold improvement in resilience. An expansion of multiVA that uses one additional vantage point with the existing $k = n - 1$ quorum policy brings the median resilience up to .98 meaning that **under these proposed modifications, the median domain is resilient to attacks from 98% of ASes on the Internet.**

Comparing multiVA to Let’s Encrypt’s previous deployment. Comparing the resilience of domains against BGP attacks from randomly chosen ASes, we find that multiVA increases the median domain resilience to 0.94 (up from 0.62) when compared to Let’s Encrypt’s previous deployment (see Figure 11). This improvement is even more significant for the domains most vulnerable to BGP attacks. Under Let’s Encrypt’s previous deployment, 10% of domains had a resilience of only 0.10 or less. With multiVA, the 10th percentile is brought up five fold to a resilience of 0.59.

While this improvement is substantial, multiVA has the potential to further improve domain resilience with some relatively small modifications. We discuss two additional ways to improve multiVA.

Impact of a full quorum policy. While a 0.94 median domain resilience is a significant improvement over the status quo, resilience can be further improved by strengthening the quorum policy (which has no impact on operating cost). Moving from an 2-3 quorum to a full quorum (3-3) further enhances the median resilience to 0.98. While worth considering, a stricter quorum policy also comes at the price of higher benign failures which could potentially outweigh the security benefits.

Improving resilience by adding vantage points. An alternative way of improving the security of multiVA is to add additional vantage points while maintaining the quorum policy ($k = n - 1$). We considered four different AWS data centers for the potential locations of additional vantage points—London, Paris, Tokyo, and Singapore—and computed the effective resilience of domains under Let’s Encrypt’s quorum policy with these additional vantage points.

Compared to Let’s Encrypt’s current deployment, adding a vantage point in Paris (the optimal location among the potential vantage points we studied) increased median resilience to 0.98 (meaning **98% of ASes on the Internet are topologically incapable of launching attacks against the median domain**). We further experimented with adding an additional vantage point in Singapore (the optimal location we found for a second vantage point after Paris) and found the median resilience to only increase to 0.986. A similar story is found with the 10th percentile domain: adding Paris improves resilience from 0.59 to .67, but further adding Singapore only improves resilience to 0.71.

With diminishing security returns and a constant cost increase associated with adding an additional vantage point, we recommend adding one additional vantage point which offers a comparable resilience improvement to the full quorum policy while maintaining the operational advantages—lower latency and a lower benign failure rate—of the current (looser) quorum policy.

Additionally, if maximum security is needed, operating one additional vantage point with a full quorum policy brings the median resilience to .99 (offering resilience against attacks from 99% of adversaries) and improves the 10th percentile resilience **seven fold** to 0.71.

7 Related Work

Routing attacks on critical applications. It is well known that attackers can exploit the insecurity of Internet routing (BGP) to hijack or intercept communications [14, 28, 35]. In fact, numerous routing attacks occur in the wild, and these attacks are getting more widespread and sophisticated [36, 37, 40]. However, most prior works analyzed these attacks from the viewpoint of availability and surveillance of unencrypted communications. A recent line of work has shown that routing attacks can compromise the security of important Internet infrastructure such as certificate authorities. Birge-Lee *et al.* [17] systematically analyzed the threat of routing attacks against the domain control validation protocol, demonstrating the ease of fraudulently obtaining certificates for a target victim domain from major certificate authorities. Gavrichenkov [27] also explored the use of BGP attacks to fraudulently obtain valid TLS certificates. These works motivate our deployment of multi-vantage-point domain validation, which substantially reduces the attack surface of

BGP attacks against CAs. In similar spirit to BGP attacks on certificate authorities, Sun *et al.* [42] and Apostolaki *et al.* [13] demonstrated routing attacks against critical infrastructure such as the Tor anonymity network and the Bitcoin crypto-currency network. These networks can also benefit from the concept of multiple vantage points.

Defenses against routing attacks. There have been substantial efforts in the industry and research community to defend against routing attacks [3, 4, 9, 20, 23, 32, 41], but unfortunately, the current status quo leaves CAs vulnerable to attacks. First, defenses based on BGP monitoring [3, 4] monitor the control plane of Internet routing to check for suspicious announcements. However, it is very challenging to accurately classify BGP announcements as legitimate or illegitimate. Furthermore, such approaches merely aid in attack detection, and cannot prevent attacks and the resulting issuance of fraudulent TLS certificates. Second, defenses based on route filtering, such as MANRS [9] and peer locking [41], use out-of-band information about the Internet topology to filter bogus BGP announcements. However, the deployment of filtering-based solutions is not widespread, and this approach does not provide a bulletproof security solution due to the difficulty of scaling out-of-band information sharing. Third, cryptographic mechanisms like RPKI [20] and BGPSEC [32] have been proposed to fully authenticate BGP announcements. While such cryptographic techniques could eliminate the threat of BGP attacks, RPKI is only partially deployed, and BGPsec has not seen any deployment. Note that RPKI only prevents an adversary from claiming ownership of an IP prefix, but does not prevent an adversary from advertising a bogus path to the prefix owner [28]. Finally, new Internet architectures like SCION [46] have been designed from the ground up to eliminate the threat of routing attacks. While SCION has made great strides in adoption, its use is still not widespread. We hope that our work on securing CAs against routing attacks provides much-needed momentum for fixing the insecurity of Internet routing.

Enhancing security of CAs. Recent work has made significant improvements in standardizing and securing the process of issuing TLS certificates [12, 16, 17, 19, 43–45]. Birge-Lee *et al.* [17] discussed the idea of multi-vantage-point domain control validation, which served as a motivation for our work and deployment. A similar idea was also explored by Brandt *et al.* [19]. Concurrent with our efforts, CloudFlare has also released an experimental API for performing domain validation using multiple CloudFlare vantage points [29]. To the best of our knowledge, their API is not in use by any CA, and our work is the first to demonstrate the feasibility of multi-vantage-point domain validation at Internet scale with successful issuance of over 300 million TLS certificates. Another thread of research has focused on transparency frameworks like Certificate Transparency [31] which aim to provide global visibility into TLS certificates issued by

CAs. Certificate Transparency logs allow domain owners to detect that fraudulent TLS certificates were issued for their domain, but user communications remain vulnerable until those certificates are revoked (a process that is itself error prone). In contrast, our approach of multi-vantage-point validation aims to prevent the issuance of bogus certificates.

8 Conclusion

We explored the design space of multi-vantage-point domain validation and showed the feasibility of balancing multiple objectives such as security, manageability, performance, and benign failures. Our deployment at Let’s Encrypt, which has secured the issuance of over 300 million TLS certificates, demonstrates the viability of multi-vantage-point domain validation at Internet scale. We make the following concluding recommendations:

- Industry-wide adoption. We have released our multi-vantage-point domain validation implementation as open-source software, and urge all certificate authorities to adopt the approach to secure TLS certificate issuance.
- Testing multiVA deployments. As other CAs start to adopt this technology, we recommend using our evaluation methodology (such as our open-source BGP simulation framework) to guide the selection of sufficiently diverse vantage points and validate the overall deployment.

References

- [1] China will block VPN access for individuals, companies must register with the Government. <https://www.neowin.net/news/china-will-block-vpn-access-for-individuals-companies-must-register-with-the-government/>, Jul 2017.
- [2] The CAIDA AS relationships dataset (March) 2020. <http://www.caida.org/data/as-relationships/>, 2019.
- [3] ThousandEyes: Network intelligence software. <https://www.thousandeyes.com/>, 2019.
- [4] Use BGPmon to monitor your prefixes and assess the risks to your network. <https://bgpmon.net/>, 2019.
- [5] Amazon peeringdb. <https://www.peeringdb.com/net/1418>, Jul 2020.
- [6] Boulder. <https://github.com/letsencrypt/boulder>, Feb 2020.
- [7] Certbot. <https://certbot.eff.org/>, Aug 2020.
- [8] gRPC: A high-performance, open source universal RPC framework. <https://grpc.io/>, Aug 2020.
- [9] MANRS Project Homepage. <https://www.manrs.org/>, 2020.

- [10] RIS raw data – RIPE network coordination centre. <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data>, 2020.
- [11] University of Oregon Route Views Project. <http://www.routeviews.org/routeviews/>, 2020.
- [12] M. Alicherry and A. D. Keromytis. DoubleCheck: Multipath verification against man-in-the-middle attacks. In *IEEE Symposium on Computers and Communications*, pages 557–563. IEEE, 2009.
- [13] M. Apostolaki, A. Zohar, and L. Vanbever. Hijacking Bitcoin: Routing attacks on cryptocurrencies. In *IEEE Symposium on Security and Privacy*. IEEE, 2017.
- [14] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the internet. In *ACM SIGCOMM*, pages 265–276, 2007.
- [15] R. Barnes, J. Hoffman-Andrews, D. McCarney, and J. Kasten. Automatic certificate management environment (ACME). RFC 8555, RFC Editor, March 2019.
- [16] D. Basin, C. Cremers, T. H.-J. Kim, A. Perrig, R. Sasse, and P. Szalachowski. Design, analysis, and implementation of ARPKI: An attack-resilient public-key infrastructure. *IEEE Transactions on Dependable and Secure Computing*, 15(3):393–408, 2016.
- [17] H. Birge-Lee, Y. Sun, A. Edmundson, J. Rexford, and P. Mittal. Bamboozling certificate authorities with BGP. In *USENIX Security Symposium*, pages 833–849, Baltimore, MD, 2018. USENIX Association.
- [18] R. Broman. Hackers emptied Ethereum wallets by breaking the basic infrastructure of the internet. <https://www.theverge.com/2018/4/24/17275982/myetherwallet-hack-bgp-dns-hijacking-stolen-ethereum>, Apr. 2018.
- [19] M. Brandt, T. Dai, A. Klein, H. Shulman, and M. Waidner. Domain validation++ for MitM-resilient PKI. In *ACM SIGSAC Conference on Computer and Communications Security*, pages 2060–2076, 2018.
- [20] R. Bush and R. Austein. The resource public key infrastructure (RPKI) to router protocol. RFC 6810, RFC Editor, January 2013.
- [21] CA/Browser Forum. *Baseline Requirements for the Issuance and Management of Publicly-Trusted Certificates*, v.1.5.4, Oct 2017.
- [22] Y. S. Chris C. Demchak. China’s maxim – leave no access point unexploited: The hidden story of China Telecom’s BGP hijacking. *Military Cyber Affairs*, 3, 2018.
- [23] J. Durand, I. Pepelnjak, and G. Doering. BGP operations and security. *IETF RFC 7454-Best Current Practice*, 2015.
- [24] L. Encrypt. Let’s Encrypt Certificates Issued Per Day. <https://letsencrypt.org/stats/#daily-issuance>, 2020.
- [25] Flexential. Flexential. <https://www.flexential.com/>, 2020.
- [26] L. Gao and J. Rexford. Stable Internet routing without global coordination. *IEEE/ACM Transactions on Networking*, 9(6):681–692, Dec 2001.
- [27] A. Gavrichenkov. Breaking HTTPS with BGP hijacking. *Black Hat USA Briefings*, 2015.
- [28] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How secure are secure interdomain routing protocols? In *ACM SIGCOMM*, pages 87–98, 2010.
- [29] D. Kozlov and G. Fisher. Securing Certificate Issuance using Multipath Domain Control Validation. <https://blog.cloudflare.com/secure-certificate-issuance/>, 2019.
- [30] M. Lad, R. Oliveira, B. Zhang, and L. Zhang. Understanding resiliency of Internet topology against prefix hijack attacks. In *IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 368–377, 2007.
- [31] B. Laurie. Certificate transparency. *Communications of the ACM*, 57(10):40–46, 2014.
- [32] M. Lepinski and K. Sriram. BGPsec protocol specification. RFC 8205, RFC Editor, September 2017.
- [33] M. Luckie, A. Dhamdhere, B. Huffaker, D. Clark, and k. claffy. bdrmap: Inference of Borders Between IP Networks. In *ACM Internet Measurement Conference*, pages 381–396, Nov 2016.
- [34] D. Madory. Excessive BGP AS Path prepending is a self-inflicted vulnerability. Oracle Internet Intelligence, <https://blogs.oracle.com/internetintelligence/excessive-as-path-prepend-is-a-self-inflicted-vulnerability>, Jul 2019.
- [35] O. Nordström and C. Dovrolis. Beware of BGP attacks. *ACM SIGCOMM Computer Communication Review*, 34(2):1–8, 2004.
- [36] Proton. Statement regarding BGP hijacking on September 29. <https://protonmail.com/blog/bgp-hijacking-september-2020/>, 2020.
- [37] A. Robachevsky. 14,000 Incidents: A 2017 Routing Security Year in Review. <https://www.internetsociety.org/blog/2018/01/14000-incidents-2017-routing-security-year-review/>, 2018.
- [38] B. Schlinker, T. Arnold, I. Cunha, and E. Katz-Bassett. PEERING: Virtualizing BGP at the edge for research. In *ACM SIGCOMM CoNEXT Conference*, 2019.
- [39] R. Shoemaker. Automated Certificate Management Environment (ACME) TLS Application-Layer Protocol Negotiation (ALPN) Challenge Extension. RFC 8737, RFC Editor, February 2020.
- [40] A. Siddiqui. Not just another BGP Hijack. <https://www.manrs.org/2020/04/not-just-another-bgp-hijack/>, 2020.
- [41] J. Snijders. Practical everyday BGP filtering with AS_PATH filters: Peer locking. *NANOG-67*, 2016.
- [42] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal. RAPTOR: Routing attacks on privacy in Tor. In *USENIX Security Symposium*, pages 271–286, 2015.
- [43] E. Syta, I. Tamas, D. Visher, D. I. Wolinsky, and B. Ford. Decentralizing authorities into scalable strongest-link coauthorities. *CoRR*, abs/1503.08768, 2015.

- [44] P. Szalachowski, S. Matsumoto, and A. Perrig. PoliCert: Secure and flexible TLS certificate management. In *ACM SIGSAC Conference on Computer and Communications Security*, pages 406–417, 2014.
- [45] D. Wendlandt, D. G. Andersen, and A. Perrig. Perspectives: Improving SSH-style host authentication with multi-path probing. In *USENIX Annual Technical Conference*, volume 8, pages 321–334, 2008.
- [46] X. Zhang, H. C. Hsiao, G. Hasker, H. Chan, A. Perrig, and D. G. Andersen. SCION: Scalability, control, and isolation on next-generation networks. In *IEEE Symposium on Security and Privacy*, pages 212–227, May 2011.