Using BGP to Acquire Bogus TLS Certificates

Henry Birge-Lee[1], Yixin Sun[1], Annie Edmundson[1], Jennifer Rexford[1], and Prateek Mittal[1]
{birgelee, yixins, annee, jrex, pmittal}@princeton.edu

[1]Princeton University

ABSTRACT

Digital certificates play an important role in secure and private communication using TLS. Thus, vulnerabilities in the process of issuing digital certificates (identity verification) can have devastating consequences for the security and privacy of online communications. In this talk, we explore the impact of BGP hijack and interception attacks on the domain verification process of obtaining a certificate. These attacks allow adversaries to obtain fake certificates for a victim's domain. While these attacks have been outlined in recent work, no study has yet to measure the effectiveness of these attacks on real-world certificate authorities. In this paper we perform these BGP interception attacks and measure the responses of some of the top certificate authorities. We also propose a new BGP attack this is more effective than those previously studied. Our results show that *none* of these certificate authorities have measures in place to prevent issuing certificates using intercepted routes which allows an attacker to obtain a certificate for a domain it does not control. In addition, this study presents two countermeasures (with reference implementations) and performs a detailed analysis of the false-positive rate of these countermeasures. Our results show that with a 0.3% false-positive rate the vast majority of attacks can be prevented.

1. INTRODUCTION

1.1 Domain validation

Upon receiving a Certificate Signing Request (CSR), a Certificate Authority CA must verify that the party submitting the CSR actually has control over the domains that are covered by that CSR. This process is known as domain control verification and is a core part of the Public Key Infrastructure (PKI) because it is the process which gives certificates the authority to identify certain domains on the web [1]. If adversaries can get a certificate for a domain they do not control, they can start a man-in-the-middle attack that tricks web clients that want to visit that domain into handing over their sensitive data to the adversary.

Thus, in this study, we will focus on HTTP domain verification: the common method of domain control verification that involves the CA requiring a user to upload a string (specified by the CA) to a specific HTTP URL at the domain. Fundamentally, in order for the CA to contact the correct web server (as opposed to a malicious server controlled by an adversary), two levels of identifiers must be resolved. First, the DNS name must be resolved into an IP address, and second the IP address must be routed to the correct server. While the ongoing deployment of a secure DNS infrastructure helps with the DNS resolution process [5], routing successfully to the resolved IP address remains a problem. Although prior work has shown how an adversary could use BGP hijacking to get a fake certificate [1], it did not perform real-world measurements of certificate authorities or develop solutions that could strengthen the domain verification process. This study performs these critical steps and exposes a larger attack surface than previously understood.

Using the BGP protocol, the most obvious method to capture traffic to a victim's domain is with a sub-prefix hijack where an adversary announces a route to a more specific IP prefix containing the victim's IP. This captures all internet traffic because in BGP more-specific routes are always preferred over more general ones. However, this attack is visible to the entire internet making it easy for network administrators to detect if exposed for too long. Adversaries can also use BGP to more stealthily hijack only part of the internet [1], but these attacks are also limited in that they require the adversary to have a specific location in the internet topology. Our research shows how an adversary in *any* location can perform a similarly stealthy attack.

2. BGP ATTACKS

2.1 AS Path Poisoning for a Stealthy Attack

Here we present what we see as the most effective attack in this space: the sub-prefix interception attack. While this attack has been outlined before [3], it has never been considered for obtaining fake certificates. The attack which involves an adversary announcing a sub-prefix would normally spread to the entire internet. However, the adversary uses a technique known as AS path poisoning (where certain AS numbers are prepended to the announcement) to prevent select ASs from importing the route due to loop detection. This can be used to maintain a path to the legitimate origin. With this path to the origin, an AS can perform a global interception attack that would be harder to detect than a hijack attack (in an interception attack, traffic to the prefix can remain uninterrupted while in a hijack attack many users will lose connectivity). Another important use of AS path poisoning is to hide the route announcement so it cannot be detected. In an extreme case, AS path poisoning could be used to make an announcement that would

[1]In this paper we will focus on Domain Validation (DV) certificates even though domain control verification is an important part of Extended Validation and Organization Validation certificates as well.

only propagate on the path between the adversary and the certificate authority by poisoning all ASs adjacent to this path that would normally propagate the announcement. In this situation an AS *anywhere* in the internet topology could make an announcement that was seen by *very few* ASs (allowing it to evade detection) and maintain *full* connectivity to the victim's domain from all parts of the internet. Thus such an attack could potentially go completely unnoticed by the community.

2.2 Executing a Real World Attack

To verify that CAs will issue certificates using hijacked routes, we performed real-world attacks (using our own IP prefixes and domains as to not affect any operating web clients or domains) against the major CAs Let's Encrypt, and Symantec. Using the PEERING framework [4] to make real BGP announcements, we ran a website in a /23 IP prefix that we controlled. We then hijacked the IP address of the website by announcing a more specific /24 prefix from second "adversarial" AS, but used AS path poisoning to forward traffic to the original website and not interrupt user connectivity. The only traffic that was answered by the adversary as opposed to being forwarded was traffic from the CA we were attacking. Using this setup we were able to obtain a certificate for the website from both CAs. We then concluded the attack by using the newly issued certificates to begin intercepting HTTPS traffic that had previously been connecting to the real website. Both of these attacks were able to begin intercepting the HTTPS connections in under 12 minutes. We also tested Comodo and GoDaddy with traditional sub-prefix hijacks and found that they were also vulnerable. These real-world attacks confirm that an adversary could indeed use a short-term BGP hijack to get a certificate for a domain they do not control. We also noted that each CA only contacted the domain from one IP address which leaves them vulnerable to a local hijack like the one previously outlined [1].

3. COUNTERMEASURES

Here we propose two countermeasures that can be implemented to strengthen the domain control verification process. The first countermeasure (5.1) forces an adversary to announce a malicious route globally to the entire internet. The second countermeasure (5.2) requires the adversary to announce that same route for a significant duration. Together, these two countermeasures eliminate the stealthy angle of this attack and should lead to the route being seen and blocked by network administrators. [2]

3.1 Multiple Vantage Point Verification

CAs are currently vulnerable to stealthy local hijacks that are not visible to the whole internet because they only verify the domain from their own IP address. To defend against this, a CA should perform the domain control verification from multiple vantage points and only consider the verification a success if all the vantage points see it [3]. The number and location of these vantage points are crucial. While only one vantage point is required to detect the most localized hijacks, an adversary could easily design an attack that hijacks both the CA itself and its vantage point. A more robust approach is to have a set of vantage points that each use the route of a different tier 1 or tier 2 provider. This way, for an adversary to hijack all the vantage points, they would have to hijack a large portion of the internet eliminating the stealthiness of the equal-prefix-length attacks.

3.2 Route Age Heuristic

In addition to using multiple vantage points, CAs need to verify the authenticity of the routes they use through control plane data to prevent sub-prefix-hijacks that affect the entire internet. CAs must do this verification on demand for any prefix with no prior contact with the prefix owner. This makes many BGP monitoring systems not applicable. Based on recent work [6], we developed a framework CAs can use to determine if a route is suspect by looking at how long ago the last routing update for that prefix has been heard. Here, in order for a route to be trusted, it can not be based on a BGP update that is more recent than a given threshold.

In addition, we performed a study of the effectiveness of this metric. We used the certificate transparency project to find out whenever a CA signed a new certificate and cross referenced this with public BGP data to get the age of the route used during the domain verification process. Our preliminary results show that with a 1/1000 false-positive rate we could require an adversary to announce a route for at least 24 hours. This would give network administrators enough time to detect the hijack (in addition, multiple vantage point verification forces the adversary's route announcement to be globally visible, making it easier to detect).

4. CONCLUSION

In this work we are able to show that it is much easier to perform a stealthy attack against a CA than previously anticipated. In addition, current CAs do not appear to have any measures preventing such attacks. We then propose and evaluate countermeasures that would require an attack to be extremely visible for a long time before using it to issue a certificate.

5. REFERENCES

[1] A. Gavrichenkov. Breaking HTTPS with BGP hijacking. Black Hat USA Briefings, 2015.

[2] D. Holmes. Should you be worried about BGP hijacking your HTTPS? *Security Week*, September 2015.

[3] A. Pilosov and T. Kapela. Stealing the internet: An internet-scale man in the middle attack. Defcon, 2008.

[4] B. Schlinker, K. Zarifis, I. Cunha, N. Feamster, and E. Katz-Bassett. Peering: An AS for Us. In *Proceedings of the ACM Workshop on Hot Topics in Networks*, pages 18:1–18:7, 2014.

[5] C. Shar. State of DNSSEC deployment 2016. Technical report, Internet Society, Reston, VA, December 2016.

[6] Y. Sun, A. Edmundson, N. Feamster, M. Chiang, and P. Mittal. Counter-RAPTOR: Safeguarding Tor Against Active Routing Attacks. *IEEE S&P 2017*, Apr. 2017. arXiv: 1704.00843.

[2]The countermeasures we outline are implemented in the Let's Encrypt code base at `https://github.com/birgelee/boulder`.

[3]Although some have brought up multiple vantage point verification [2], it is not implemented in any of the CAs we have tested and more importantly is not part of the CA and Browser Forum Baseline Requirements for obtaining a certificate.