Fall 2023 CA/B Forum F2F

Validation Sub-committee

Progress since Summer 2023 F2F

- Domain validation threat modeling tiger team led by Michael Slaughter of Amazon Trust Services produced a comprehensive threat model for domain validation
- Will be primary topic of discussion today

Progress since Summer 2023 F2F

- MPDV/MPIC effort led by Ryan and Chris of the Chrome Root program making steady progress
  - Draft ballot text being refined
- Will be second topic of discussion today

Progress since Summer 2023 F2F

- Completed review of items that were identified during “Applicant” and “Applicant Representative” analysis of the TLS BRs
  - Several items are incorporated in Ben and Dustin’s Subscriber Agreement improvement ballot
  - Some items were moved up to servercert-wg
  - Others are in Github issues
  - and a few won’t be addressed

Progress since Summer 2023 F2F

- Guest presentation by Q Misell on the use of ACME for certificates that contain Onion Domain Names
- Special focus on CAA checking for Onion Domain Names
- Robust discussion, several participants are now involved in the discussions surrounding standardization of the proposal at IETF

Agenda

1. Michael Slaughter’s presentation on domain validation threat modeling
2. Ryan Dickson and Chris Clement’s presentation on the MPDV/MPIC draft ballot

Delegation of Domain Validation to the CA

Threat Model Overview

Michael Slaughter
Amazon Trust Services

Background

- At F2F 59 (July 23’), the Validation Subcommittee of the Server Certificate WG presented the following conclusions on the practice of the **Delegation of Domain Validation to the CA**:
  - More clarity is needed around the practice
  - Applicants generally delegate the performance of many aspects of operating a website.
  - If done correctly, allowing Applicants to delegate the placement of the Random Value/ Request Token boosts agility and automation.
  - There are reasonable interpretations of the BRs that such delegation is already allowed today.

- Following F2F59, The Delegated Domain Validation Tiger Team was formed

- Emphasis on the practice of delegation of domain validation to the CA under Method 7 (DNS Change)

Tiger Team Participants and Practices

- Formed in July 2023
  - Worked on shared Threat Model Doc
  - Held (4) Syncs and regular updates to validation subcommittee

- Followed OWASP Threat Modeling process using a STRIDE Threat List
  - Decompose the Application
  - Determine and Rank Threats
  - Determine Countermeasures and Mitigations

Participants

Name	Company
Michael Slaughter	Amazon
Corey Bonnell	DigiCert
Clint Wilson	Apple
Martijn Katerbarg	Sectigo

Goals of the Tiger Team

- Analyze and document the use cases, threats and mitigations for Method 7 when delegated to the CA
- Propose improvements to current Method 7 to clarify the practice and mitigate the threats identified
- Identify areas that new automation-centric domain validation method should explore in future efforts

Delegated Domain Validation Roadmap

1. Threat Model Tiger Team
   a. Form Threat Model Tiger Team
   b. Draft a Threat Model w/ an emphasis on Method 7
   c. Review Threat Model with broader community

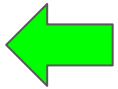

2. Propose a ballot with clarifications for the practice with Method 7

3. Propose a ballot with a new DCV method focused on automation

Threat Model Overview

Threat Model Assumptions

1. Network is reliable and secure.
2. DNS lookup results can be relied upon.
3. Applicants/Subscribers will not share login credentials or private key material with unauthorized parties.
4. All CA software systems and hardware are trustworthy and function as designed.
5. If a CA has an account separation mechanism it can be relied upon

Use Cases, Assets and Trust Levels

Use Cases

- First-time Domain Control Verification (DCV) under Method 7
- Subsequent Domain Control Verification (DCV) Request under Method 7

Assets

- Ability to (re)issue publicly-trusted certificate for a given (sub)domain name

Trust Levels

- Applicant with provable control of a domain
- Applicant without control of a domain

Entry Points

- CA Domain Control Verification System
  - A web API operated by a Certificate Authority that performs Domain Control Verification (DCV).

Current Method 7

- Confirming the Applicant's control over the FQDN by confirming the presence of a **Random Value or Request Token** for either in a DNS **CNAME**, **TXT** or **CAA** record for either
  - 1) an **Authorization Domain Name**; or
  - 2) an Authorization Domain Name that is prefixed with a Domain Label that begins with an underscore character.

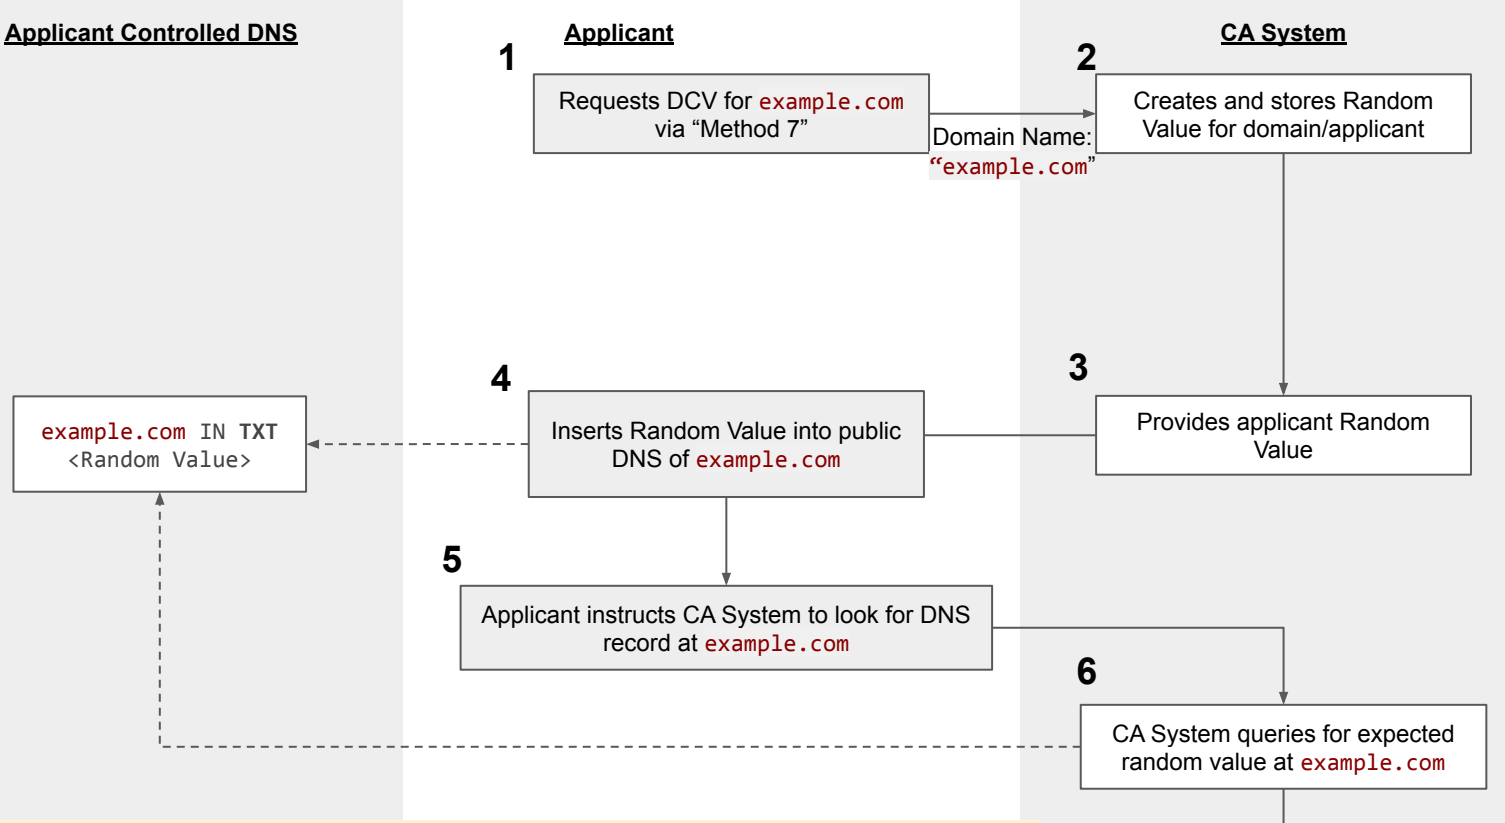

Traditional DNS Validation (First-time and Subsequent Requests)

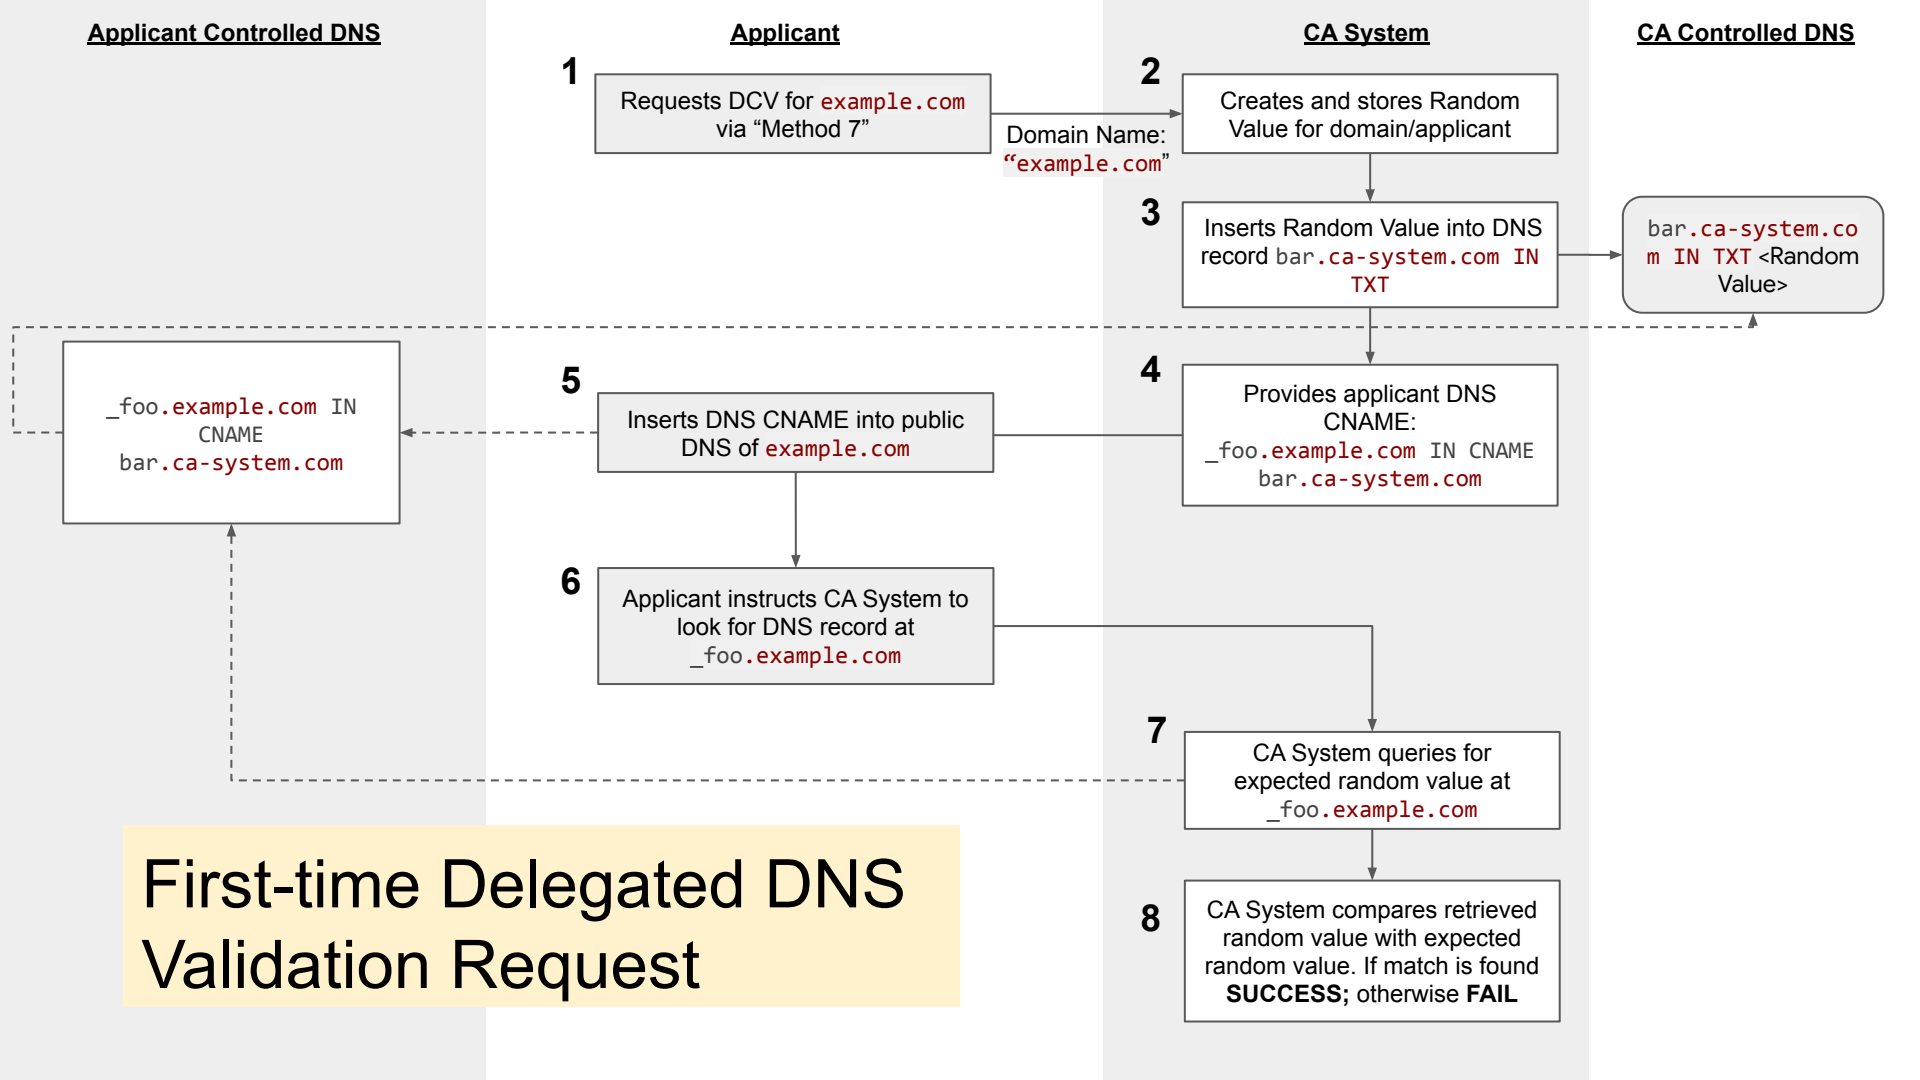
Applicant Controlled DNS
Applicant
CA System
CA Controlled DNS
1
Requests DCV for example.com via “Method 7”
Domain Name: “example.com”
2
Creates and stores Random Value for domain/applicant
3
Inserts Random Value into DNS record bar.ca-system.com IN TXT
bar.ca-system.com IN TXT <Random Value>
4
Provides applicant DNS CNAME: _foo.example.com IN CNAME bar.ca-system.com
5
Inserts DNS CNAME into public DNS of example.com
_foo.example.com IN CNAME bar.ca-system.com
6
Applicant instructs CA System to look for DNS record at _foo.example.com
7
CA System queries for expected random value at _foo.example.com
8
CA System compares retrieved random value with expected random value. If match is found SUCCESS; otherwise FAIL
First-time Delegated DNS Validation Request

Subsequent Delegated Domain Control Verification Requests

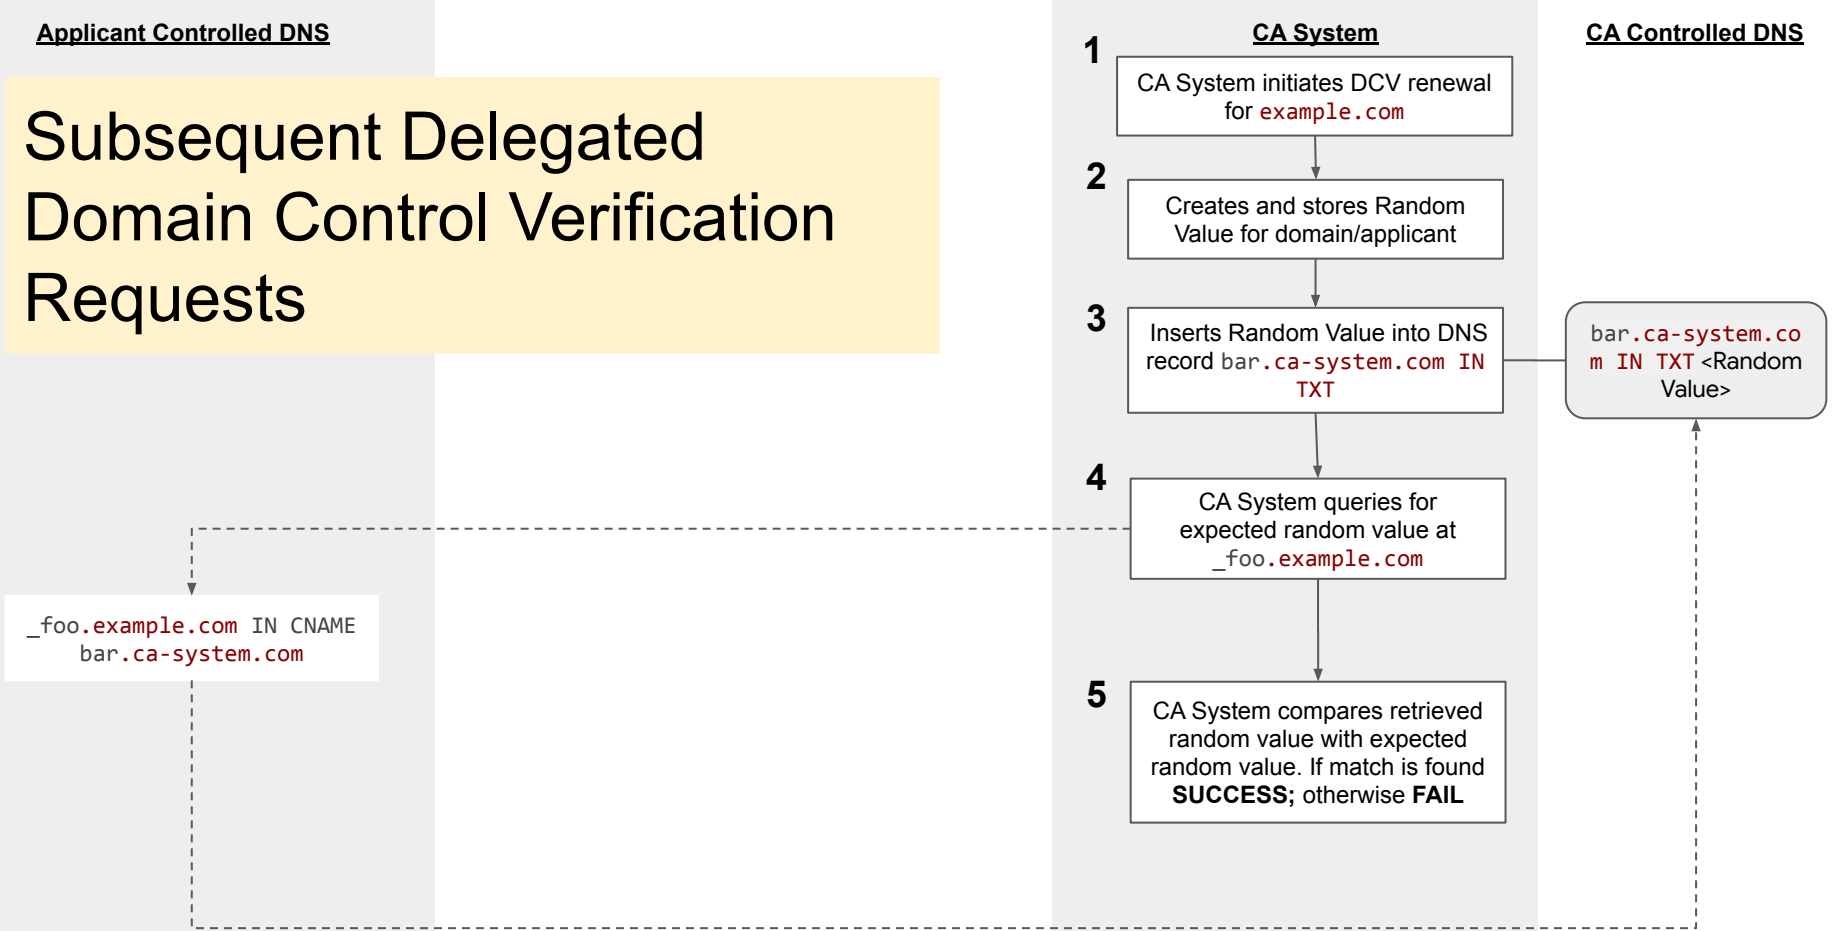

Summary of Threats

ID	Name	Threat	STRIDE Type
1	CNAME is Not Unique to Applicant	A malicious user performs DNS validation using a CNAME intended for a different applicant to attain a certificate for a domain name they do not control.	Spoofing
2	Approval for Expired Domain Registration	An malicious user issues a certificate for an expired domain with an active delegated CNAME in DNS due to long lived caches / TTLs. A malicious user buys a domain that expired that a certificate is still active for.	Spoofing
3	Approval for Domain Transferred to different party	A malicious domain owner sets a long-lived TTL for the delegated CNAME that lives beyond their ownership of the domain	Spoofing

Summary of Threats

ID		Threat	STRIDE Type
4	Domain Owner Grants Overly Broad Permission	A malicious user requests a domain owner configure a CNAME record _foo.example.com pointing to bar.ca-system.com. The domain owner mistakenly believes the action approves the issuance of a single certificate rather than multiple.	Elevation of Privilege
5	Improper DNS Zone Usage	CA uses DNS zone for other purposes unrelated to Domain Validation.	Elevation of Privilege

Overview of Proposed Improvements to Method 7

- Recommendations we should explicitly enforce by rule:
  - Unique Account Binding
    - **Mitigates**: CNAME is Not Unique to Applicant
  - Lookup Freshness
    - **Mitigates**: Approval for Domain Transferred to different party, Approval for Expired Domain Registration
  - Restrictions on the DNS Resource Record Types allowed in the CA DNS Zone
    - **Mitigates**: Improper DNS Zone Usage

- Other Recommendations / Considerations:
  - Applicant education and expectation setting
    - **Mitigates**: **Improper DNS Zone Usage**
  - Limits to CNAME chain length

Propose a ballot for an Automation-Centric DCV method

- Recommend this method consider
  - Significantly reducing the validation reuse period
  - Removing the use of random value or request tokens
  - Utilizing CAA records

- Work may impact multiple domain validation methods
  - File-Based, CAA Based, DNS Change

Roadmap

1. Threat Model Tiger Team
   a. Form Threat Model Tiger Team
   b. Draft a Threat Model w/ an emphasis on Method 7
   c. Review Threat Model with broader community

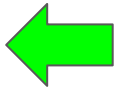

2. Propose a ballot with modifications to Method 7

3. Propose a ballot with a new DCV method focused on automation

Appendix

Proposed Improvements to Method 7 - Part 1

- [**ADD**] CAs **MAY** operate domains for the purpose of assisting customers with this validation, and **MAY** instruct customers to add a CNAME redirect from an Authorization Domain Name to such a domain.

- **Rationale:**
  - Make explicit in the language that CAs are ALLOWED to operate a domain in support domain validation using a CNAME redirect using an authorized domain name
  - Addresses the lack of clarity and room for interpretation currently in the BRs around whether or not this practice is allowed.

Proposed Improvements to Method 7 - Part 2

- [**ADD**] If the CA does so, the CA **SHALL** ensure that each domain name is used for a unique Applicant, and not shared across multiple Applicants.

- **Mitigated Threat(s):**
  - Violation of Authenticity / Spoofing

- **Rationale:**
  - Adds guardrails around the method to prevent implementations that do not ensure applicant unique automated approvals.

Proposed Improvements to Method 7 - Part 3

- [**ADD**] The CA **SHALL** treat the TTL of CNAME records as being the TTL of the CNAME record or 8 hours, whichever is lesser.

- **Mitigated Threat(s):**
  - Violation of Authenticity / Spoofing

- **Rationale:**
  - Adds guardrails and protections around stale DNS records and changes in domain registration.

Proposed Improvements to Method 7 - Part 4

- [**ADD**] If a DNS zone is utilized for the purposes of delegated DNS validation, the CA **SHALL** limit the record types in that zone records explicitly required for domain validation.

- **Mitigated Threat(s):**
  - Violation of Integrity / Tampering

- **Rationale:**
  - Adds guardrails and protections around the operation and utilization of the DNS hosted zone used by a CA to delegate domain validation.

Assets

ID	Name	Description	Trust Levels
1	Ability to (re)issue publicly-trusted certificate for a given (sub)domain name	The ability to attain a certificate signed by the private key of a publicly-trusted certificate authority that contains a given domain name. This covers both initial issuance and renewal use cases.	(2) Applicant with provable control over the DNS zone of a (sub)domain name
2	Ability to revoke publicly-trusted certificates for a given (sub)domain name	The ability to revoke a certificate signed by the private key of a publicly-trusted certificate authority that contains a given domain name. CA Systems may determine authorization of the revocation of a certificate with a given domain name by performing DCV on that domain name.	(1) Anonymous Web User (2) Applicant with provable control over the DNS zone of a (sub)domain name (3) Applicant without control over the DNS zone of a (sub)domain name

Strengthening domain validation using Multi-Perspective Issuance Corroboration (MPIC)

An introductory proposal to update the TLS Baseline Requirements to require MPIC.

Presented to the Server Certificate Working Group
at Face-to-Face 60 (October 4, 2023)

Refresher

Face-to-Face #58 (February 2023)

- Henry Birge-Lee (Princeton University) participated as a guest speaker
- During his session, Henry summarized Internet routing vulnerabilities that affect Web PKI domain validation methods defined by the BRs
    - specific concern: the relative ease with which an attacker can misroute Internet traffic to obtain a fraudulently-issued TLS certificate, enabling future abuse
- Beyond just telling us about these vulnerabilities, he demonstrated their exploitation in a live demonstration.
- If you missed the demo, here's a similar recorded version.

Commitment to Action

- Following the presentation, the Chrome Root Program Team volunteered to lead a Work Team of other interested parties focused on drafting a set of requirements that could be added to the Baseline Requirements to reduce risk related to the vulnerabilities presented by the Princeton Team.

- A Work Team was formed. Members collaborated on draft requirements that are now ready for broader community feedback.

- This presentation is intended to help collect feedback and identify concerns in preparation for Balloting.

Learn More

- Related research:
  - On MPIC:
    - How Effective is Multiple-Vantage-Point Domain Control Validation?
    - Experiences Deploying Multi-Vantage-Point Domain Validation at Let's Encrypt.
    - Let's Downgrade Let's Encrypt
    - Domain Validation++ For MitM-Resilient PKI
  - On the problem space:
    - Attackers exploit fundamental flaw in the web's security to steal $2 million in cryptocurrency
    - Celer Bridge incident analysis
    - Bamboozling Certificate Authorities with BGP
    - Face-to-Face #58 (Presentation from Princeton Team)
    - Securing Internet Applications from Routing Attacks

MPIC Work Team, Artifacts, and Outputs

Thank You, Work Team Participants!

Name	Organization	Name	Organization
Aaron Gable	Let's Encrypt	**Henry Birge-Lee**	Princeton University
Aaron Poulsen	Amazon Trust Services	**Liang Wang**	Princeton University
Antonios Eleftheriadis	HARICA	**Michael Slaughter**	Amazon Trust Services
Ben Wilson	Mozilla	**Prateek Mittal**	Princeton University
Chris Clements	Google Chrome	**Rollin Yu**	TrustAsia
Clint Wilson	Apple	**Ryan Dickson**	Google Chrome
Corey Bonnell	DigiCert	**Tim Crawford**	BDO
David Kluge	Google Trust Services	**Tim Hollebeek**	DigiCert
Dimitris Zacharopoulos	HARICA	**Tobias Josefowitz**	Opera
Ellie Lu	TrustAsia	**Trevoli Ponds-White**	Amazon Trust Services
Grace Cimaszewski	Princeton University	**Wayne Thayer**	Fastly
Gurleen Grewal	Google Trust Services		

Work Team Artifacts & Outputs

- Project Plan (*note: many discussions carried out in Comments that are now closed, but still accessible in the document history*)

- Meeting Minutes
  - Kick-off
  - Follow-up

- Google Group (*limited use, it took some time to get this set up*)

Work Team Artifacts & Outputs

- Draft Requirements
  - Version 1
    - Markdown (*clean*)
    - Compare (*to "Ballot SC-063: Make OCSP Optional"*)
  - Version 2 (current)
    - Markdown (*clean*)
    - Compare to:
      - Version 1
      - "Ballot SC-063: Make OCSP Optional"

Background and Motivation

Borrowing a metaphor

*If **Border Gateway Protocol (BGP) is the "postal service" of the Internet responsible for delivering data through the most efficient routes**, then Autonomous Systems (AS) are individual post office branches that represent an Internet network run by a single organization.*

***Sometimes network-level adversaries advertise false routes over BGP to steal traffic**, especially if that traffic contains something important, like a domain's certificate."*

- Cloudflare Blog, 2019

Types of BGP Attack

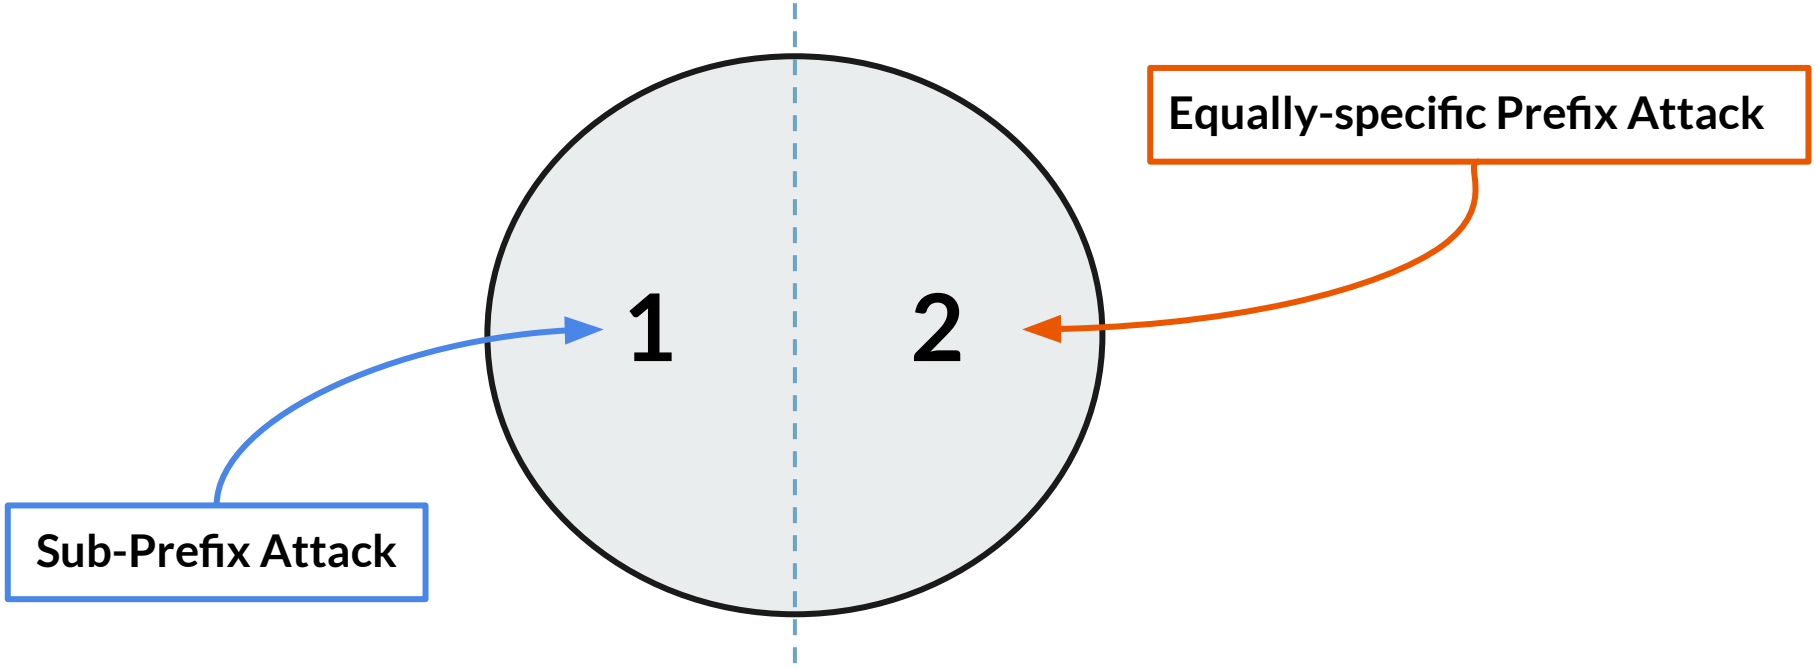

Sub-Prefix Attack [Type 1]

- Routers prefer to follow more specific routes (e.g., 2001:DB8:1000::/**48** is preferred over 2001:DB8::/**32**, as it's **more** specific.)

- By announcing a more specific prefix, an attacker can capture the victim's traffic.

- Deployed Resource Public Key Infrastructure (RPKI) today mitigates this problem, and **further RPKI adoption can solve this problem completely.**

Equally-specific Prefix Attack [Type 2]

- In an equally-specific prefix attack, the attacker announces the same prefix as the victim.
- ASes choose path based on properties (like topological proximity to the attacker) meaning an attacker can influence AS decision making.
- Only a portion of internet traffic is intercepted in this style of attack.
- **RPKI does not address this problem.**

Specific Problem Statement

- Make it more difficult for adversaries to launch **equally-specific** prefix attacks against the domain validation processes described in the TLS BRs.
    - *Note: While other sets of BRs might benefit from similar work, they were not in scope of this effort.*

Understanding the attack (illustrative)

- **Step 1: Attacker...**
  - Selects a victim domain and a victim CA
  - Launches a BGP attack on victim domain affecting victim CA, waits for route to propagate
    - *Note: Traffic routed to the victim domain will now be re-routed to an attacker controlled server for the affected parts of the Internet*
  - Requests a certificate representing the victim domain
- **Step 2: CA...** generates and shares a challenge Request Token with the attacker

Understanding the attack (illustrative, continued)

- **Step 3: Attacker...** posts the challenge as requested by the CA
- **Step 4: CA...**
  - Verifies the challenge was posted as requested (*by the attacker*)
  - Issues the requested certificate (*to the attacker*)
- **Step 5: Attacker...** abuse (impersonation, interception, $other)
  - *Note: the abuse is not necessarily limited to the same BGP attack "audience" or time window as the attack required to issue a certificate to the victim domain*

How often do BGP attack/hijacks take place?

- Mutually Agreed Norms for Routing Security (MANRS) is a global initiative, supported by the Internet Society, that studies and provides crucial fixes to reduce the most common routing threats.
    - August 2023: MANRS reported 1,634 incidents from 1,092 culprits.

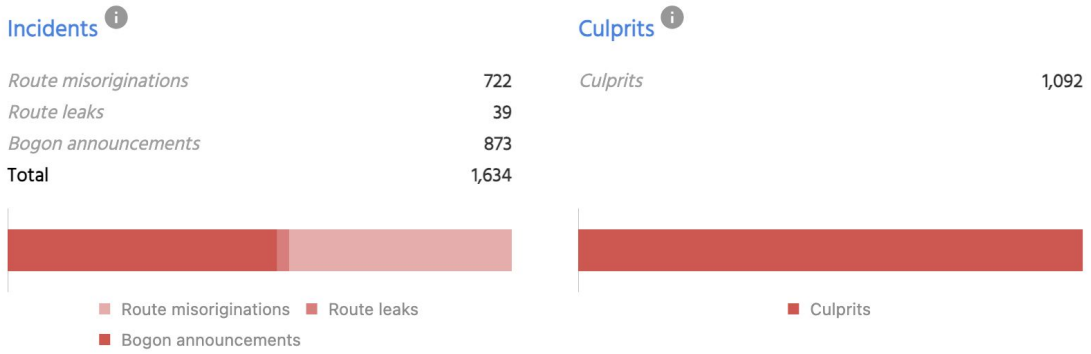

Is it possible to fully estimate the # of attack/hijacks?

- No.
- Challenges:
  - the Internet is a distributed system, and there is no ground truth on what proper BGP announcements are, or even who is connected to who.
  - Most Internet-scale attack detection ultimately boils down to using historical data.
  - Challenges with transparency, detection, and attribution.

Illustrative impact of one of these attacks

- KLAYswap Attack
  - **When:** February 3, 2022
  - **What:** attackers stole approximately $2 million worth of cryptocurrency from users of the Korean crypto exchange KLAYswap
  - **How:** attackers served a malicious javascript file over TLS using a publicly-trusted certificate representing a partner domain, obtained and exploited by using a BGP attack
    - KLAYswap customers were connecting to the legimtiate KLAYswap website, but were connecting to attacker-controlled subresources.
      - How would they have known any different?

Background: Subresources

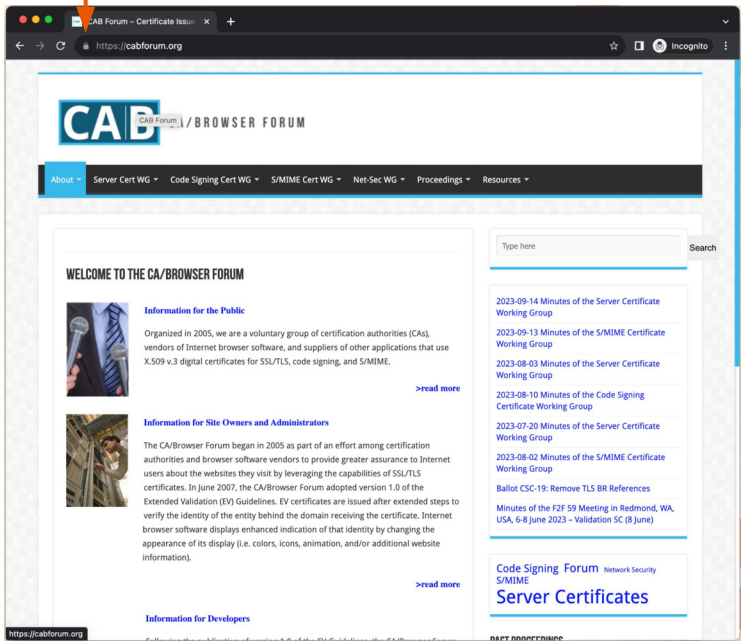

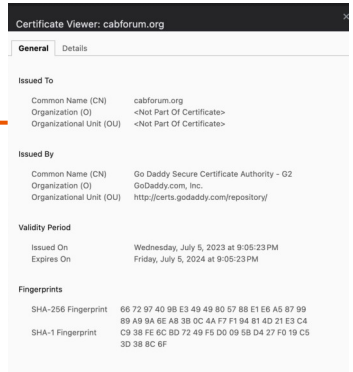

ID	Source Type	Description
351872	CERT_VERIFIER_JOB	cabforum.org
-2147482560	CERT_VERIFIER_TASK	cabforum.org
351942	CERT_VERIFIER_JOB	fonts.googleapis.com
-2147482559	CERT_VERIFIER_TASK	fonts.googleapis.com
351944	CERT_VERIFIER_JOB	ssl.gstatic.com
-2147482558	CERT_VERIFIER_TASK	ssl.gstatic.com
-2147482557	CERT_VERIFIER_TASK	fonts.gstatic.com
352007	CERT_VERIFIER_JOB	fonts.gstatic.com
-2147482556	CERT_VERIFIER_TASK	img1.wsimg.com
352009	CERT_VERIFIER_JOB	img1.wsimg.com
-2147482555	CERT_VERIFIER_TASK	img6.wsimg.com
352027	CERT_VERIFIER_JOB	img6.wsimg.com
-2147482554	CERT_VERIFIER_TASK	content-autofill.googleapis.com
352073	CERT_VERIFIER_JOB	content-autofill.googleapis.com
352075	CERT_VERIFIER_JOB	events.api.secureserver.net
-2147482553	CERT_VERIFIER_TASK	events.api.secureserver.net

Doesn't $other_thing solve this problem?

- **Resource Public Key Infrastructure (RPKI)** -> **No.**
  - Does not protect against equally-specific prefix attacks.
- **Certificate Transparency** -> **No.**
  - Monitoring CT would allow a victim to learn they've been subject of an attack... most likely after it's finished.
- **Certification Authority Authorization** -> **No.**
  - An attacker can forge DNS records, including CAA.
- **Domain Name System Security Extensions (DNSSEC)** -> **No.**
  - DNSSEC doesn't directly prevent BGP attack/hijacks.

How MPIC Helps

- **Increases attack friction** by requiring an attacker to succeed at launching a BGP attack/hijack at a global scale.
    - Doing so would be complex, have low viability, and consequently, is unlikely.

Why standardize this in the BRs?

- We’re only as strong as the weakest link.
- So long as any one CA remains vulnerable to this style of attack, so is any domain on the Internet.

Approach

Approach

1. Define key terms related to MPIC
    - Multi-Perspective Issuance Corroboration
    - Network Perspective
    - Primary Network Perspective

2. Identify the validation methods that must rely on MPIC

Approach (continued)

3. Define MPIC requirements
    - Describe requirements for Network Perspectives
    - Describe what it means for a Network Perspective to “corroborate” the primary determination
    - Define corroborating quorum requirements
4. Define logging requirements
5. Define implementation timeline

Proposal

[PROPOSED] Key Terms

- **Multi-Perspective Issuance Corroboration:** A process by which the determinations made during domain validation and CAA checking by the Primary Network Perspective are corroborated from other Network Perspectives before Subscriber Certificate issuance.

- **Network Perspective:** Related to Multi-Perspective Issuance Corroboration. A system for sending outbound Internet traffic associated with a domain control validation method and CAA check. The location of a Network Perspective is determined by the point where unencapsulated outbound Internet traffic is first handed off to the network infrastructure providing Internet connectivity to that perspective.

[PROPOSED] Key Terms

- **Primary Network Perspective:** The Network Perspective used by the CA to determine 1) its authority to issue a Certificate for the requested domain(s) or IP address(es) and 2) the Applicant's authority and/or domain authorization or control of the requested domain(s) or IP address(es).

[PROPOSED] Applicable Validation Methods

TLS BR Section	Method Name
3.2.2.4.7	DNS Change
3.2.2.4.8	IP Address
3.2.2.4.13	Email to DNS CAA Contact
3.2.2.4.14	Email to DNS TXT Contact
3.2.2.4.17	Phone Contact with DNS CAA Phone Contact
3.2.2.4.18	Agreed-Upon Change to Website v2
3.2.2.4.19	Agreed-Upon Change to Website - ACME
3.2.2.4.20	TLS Using ALPN
3.2.2.5.1	Agreed-Upon Change to Website
3.2.2.5.2	Email, Fax, SMS, or Postal Mail to IP Address Contact
3.2.2.5.6	ACME “http-01" method for IP Addresses
3.2.2.5.7	ACME “tls-alpn-01" method for IP Addresses

[PROPOSED] Requirements for Network Perspectives

- Network Perspectives MUST:
  - be unique and sufficiently diverse:
    - straight-line distance between perspective States, Provinces, or Countries MUST be > 500km
    - spread across two (2) distinct regional Internet registries (after December 15, 2025)
  - independently verify:
    - the presence of the expected 1) Random Value, 2) Request Token, 3) IP Address, or 4) Contact Address, as required by the relied upon validation method specified in Sections 3.2.2.4 and 3.2.2.5, and
    - the CA's authority to issue to the requested domain(s) or IP address(es), as specified in Section 3.2.2.8

Why does Network Perspective “diversity" matter?

- Reduce likelihood of localized or regional BGP attack/hijacks from circumventing these controls
    - *Imagine a BGP attack/hijack targeting an asset in Portsmouth, NH, with a corroborating perspective in Boston, MA (85 km away).*
        - *It’s likely that both locations would be affected by the same attack.*

[PROPOSED] Quorum Requirements

# of Distinct Network Perspectives Used	# of Allowed non-Corroborations
2-5	1
6+	2

- **Why these?**
  - Promote resilience
  - Allow for the unexpected (failures, outages, DNS caching, etc.)

[PROPOSED] Perspective System Security Requirements

- **MUST:**
  - Forward all Internet traffic via a network or set of networks that filter all RPKI-invalid BGP routes as defined by RFC 6811

- **SHOULD:**
  - Generally follow best practices derived from the NetSec requirements related to:
    - Facility & Service Provider Requirements
    - Network Hardening
    - System Hardening
    - Vulnerability Detection and Patch Management

[PROPOSED] Logging Requirements

- General goals:
  - Allow sufficient evidence to demonstrate MPIC is working as intended
  - Support forensic evaluation when it's not

- Specific requirements were added to 5.4.1.
  - *"an identifier that uniquely identifies the perspective used"*
  - *"the attempted domain name or IP address"*
  - *"the result of the attempt (i.e., "DCV pass/fail, CAA allow/disallow")"*
  - *"quorum results for each attempted domain name or IP address represented in a Certificate request"*

[PROPOSED] Implementation Timeline

- Phased implementation approach:
  - **Effective June 15, 2024:** CAs SHOULD implement MPIC.
  - **Effective December 15, 2024:** CAs must have implemented MPIC, but are not required to block in the absence of corroboration.
  - **Effective June 15, 2025:** CAs must implement “blocking” MPIC.

- The minimum quorum requirements strengthen over time:
  - From 2 (until **12/15/2025**) to >=5 (after **12/15/2026**)

[PROPOSED] Implementation Timeline (continued)

- **Why this approach?**
  - offer flexibility to CA owners
  - allow sufficient time to implement and fine-tune a new process that has the potential to block certificate issuance (e.g., tune for false-positives)

Opportunities to make MPIC more accessible

- **APIs and Services:**
  - **Cloudflare:** Multipath DCV service API
  - **Princeton Open Source Project:** working on an open-source project that can be easily run from a CA owner's preferred cloud service provider, possibly as early as January 2024
- **Standardization:**
  - IETF RFC to promote consistent implementations?

Next Steps

- We'd appreciate your feedback!
  - **Preference:** Add *suggested edits* directly to the existing Pull Request, please share motivation for the suggestion.

- If you're interested in endorsing a future ballot, let us know!
  - **Proposer:** Chrome Root Program (Ryan and Chris)
  - **Endorsers:** Let's Encrypt (Aaron) and _____

Discussion