TRANS (Public Notary Transparency) R. Stradling
Internet-Draft Comodo CA, Ltd.
Intended status: Experimental E. Messeri
Expires: March 4, 2017 Google UK Ltd.
August 31, 2016

Certificate Transparency: Domain Label Redaction
draft-strad-trans-redaction-00

Abstract

We define a mechanism to allow DNS domain name labels that are considered to be private to not appear in public Certificate Transparency (CT) logs, while still retaining most of the security benefits that accrue from using Certificate Transparency mechanisms.

Table of Contents

1. Introduction

Some domain owners regard certain DNS domain name labels within their registered domain space as private and security sensitive. Even though these domains are often only accessible within the domain owner's private network, it's common for them to be secured using publicly trusted Transport Layer Security (TLS) server certificates.

Certificate Transparency [I-D.ietf-trans-rfc6962-bis] describes a protocol for publicly logging the existence of TLS server certificates as they are issued or observed. Since each TLS server certificate lists the domain names that it is intended to secure, private domain name labels within registered domain space could end up appearing in CT logs, especially as TLS clients develop policies that mandate CT compliance. This seems like an unfortunate and potentially unnecessary privacy leak, because it's the registered domain names in each certificate that are of primary interest when using CT to look for suspect certificates.

TODO: Highlight better the differences between registered domains and subdomains, referencing the relevant DNS RFCs.

Section TBD of [I-D.ietf-trans-rfc6962-bis] proposes two mechanisms for dealing with this conundrum: wildcard certificates and name-constrained intermediate CAs. However, these mechanisms are insufficient to cover all use cases.

TODO(eranm): Expand on when each of the other mechanisms is suitable and when this mechanism may be suitable.

We define a domain label redaction mechanism that covers all use cases, at the cost of increased implementation complexity. CAs and domain owners should note that there are privacy considerations (Section 8) and that TLS clients may apply additional requirements (relating to the use of this redaction mechanism) for a certificate to be considered compliant.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Redacting Labels in Precertificates

When creating a precertificate, the CA MAY include a redactedSubjectAltName (Section 4) extension that contains, in a redacted form, the same entries that will be included in the certificate's subjectAltName extension. When the redactedSubjectAltName extension is present in a precertificate, the subjectAltName extension MUST be omitted (even though it MUST be present in the corresponding certificate).

Wildcard "*" labels MUST NOT be redacted, but one or more non-wildcard labels in each DNS-ID [RFC6125] can each be replaced with a redacted label as follows:

```
  REDACT(label) = prefix || BASE32(index || _label_hash)
    _label_hash = LABELHASH(keyid_len || keyid || label_len ||
label)
```

"label" is the case-sensitive label to be redacted.

"prefix" is the "?" character (ASCII value 63).

"index" is the 1 byte index of a hash function in the CT hash algorithm registry (section TBD of [I-D.ietf-trans-rfc6962-bis]). The value 255 is reserved.

"keyid_len" is the 1 byte length of the "keyid".

"keyid" is the keyIdentifier from the Subject Key Identifier extension (section 4.2.1.2 of [RFC5280]), excluding the ASN.1 OCTET STRING tag and length bytes.

"label_len" is the 1 byte length of the "label".

"||" denotes concatenation.

"BASE32" is the Base 32 Encoding function (section 6 of [RFC4648]). Pad characters MUST NOT be appended to the encoded data.

"LABELHASH" is the hash function identified by "index".

4. redactedSubjectAltName Certificate Extension

The redactedSubjectAltName extension is a non-critical extension (OID 1.3.101.77) that is identical in structure to the subjectAltName extension, except that DNS-IDs MAY contain redacted labels (Section 3).

When used, the redactedSubjectAltName extension MUST be present in both the precertificate and the corresponding certificate.

This extension informs TLS clients of the DNS-ID labels that were redacted and the degree of redaction, while minimizing the complexity of TBSCertificate reconstruction (Section 6). Hashing the redacted labels allows the legitimate domain owner to identify whether or not each redacted label correlates to a label they know of.

TODO: Consider the pros and cons of this 'un'redaction feature. If the cons outweigh the pros, switch to using Andrew Ayer's alternative proposal of hashing a random salt and including that salt in an extension in the certificate (and not including the salt in the precertificate).

Only DNS-ID labels can be redacted using this mechanism. However, CAs can use Name Constraints (section TBD of [I-D.ietf-trans-rfc6962-bis]) to allow DNS domain name labels in other subjectAltName entries to not appear in logs.

TODO: Should we support redaction of SRV-IDs and URI-IDs using this mechanism?

5. Verifying the redactedSubjectAltName extension

If the redactedSubjectAltName extension is present, TLS clients MUST check that the subjectAltName extension is present, that the subjectAltName extension contains the same number of entries as the redactedSubjectAltName extension, and that each entry in the subjectAltName extension has a matching entry at the same position in the redactedSubjectAltName extension. Two entries are matching if either:

- The two entries are identical; or

- Both entries are DNS-IDs, have the same number of labels, and each label in the subjectAltName entry has a matching label at the same position in the redactedSubjectAltName entry. Two labels are matching if either:
  - The two labels are identical; or,
  - Neither label is "*" and the label from the redactedSubjectAltName entry is equal to REDACT(label from subjectAltName entry) (Section 3).

If any of these checks fail, the certificate MUST NOT be considered compliant.

6. Reconstructing the TBSCertificate

Section TBD of [I-D.ietf-trans-rfc6962-bis] describes how TLS clients can reconstruct the TBSCertificate component of a precertificate from a certificate, so that associated SCTs may be verified.

If the redactedSubjectAltName extension (Section 4) is present in the certificate, TLS clients MUST also:

- Verify the redactedSubjectAltName extension against the subjectAltName extension according to Section 5.
- Once verified, remove the subjectAltName extension from the TBSCertificate.

7. Security Considerations

7.1. Avoiding Overly Redacting Domain Name Labels

Redaction of domain name labels carries the same risks as the use of wildcards (e.g., section 7.2 of [RFC6125]). If the entirety of the domain space below the unredacted part of a domain name is not registered by a single domain owner (e.g., REDACT(label).com, REDACT(label).co.uk and other [Public.Suffix.List] entries), then the domain name may be considered by clients to be overly redacted.

CAs should take care to avoid overly redacting domain names in precertificates. It is expected that monitors will treat precertificates that contain overly redacted domain names as potentially misissued. TLS clients MAY consider a certificate to be non-compliant if the reconstructed TBSCertificate (Section 6) contains any overly redacted domain names.

8. Privacy Considerations

8.1. Ensuring Effective Redaction

Although the domain label redaction mechanism removes the need for private labels to appear in logs, it does not guarantee that this will never happen. Anyone who encounters a certificate could choose to submit it to one or more logs, thereby rendering the redaction futile.

Domain owners are advised to take the following steps to minimize the likelihood that their private labels will become known outside their closed communities:

- Avoid registering private labels in public DNS.
- Avoid using private labels that are predictable (e.g., "www", labels consisting only of numerical digits, etc). If a label has insufficient entropy then redaction will only provide a thin layer of obfuscation, because it will be feasible to recover the label via a brute-force attack.
- Avoid using publicly trusted certificates to secure private domain space.

CAs are advised to carefully consider each request to redact a label. When a CA believes that redacting a particular label would be futile, we advise rejecting the redaction request. TLS clients may have policies that forbid redaction, so redaction should only be used when it's absolutely necessary and likely to be effective.

9. Acknowledgements

The authors would like to thank Andrew Ayer and TBD for their valuable contributions.

A big thank you to Symantec for kindly donating the OID from the 1.3.101 arc that is used in this document.

10. References

10.1. Normative References

[I-D.ietf-trans-rfc6962-bis]
Laurie, B., Langley, A., Kasper, E., Messeri, E., and R. Stradling, "Certificate Transparency", draft-ietf-trans-

rfc6962-bis-18 (work in progress), July 2016.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <http://www.rfc-editor.org/info/rfc4648>.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <http://www.rfc-editor.org/info/rfc5280>.

[RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, DOI 10.17487/RFC6125, March 2011, <http://www.rfc-editor.org/info/rfc6125>.

10.2. Informative References

[Public.Suffix.List] Mozilla Foundation, "Public Suffix List", 2016, <https://publicsuffix.org>.

Authors' Addresses

Rob Stradling
Comodo CA, Ltd.

Email: rob.stradling@comodo.com

Eran Messeri
Google UK Ltd.

Email: eranm@google.com