

TRANS
Internet-Draft SECOM CO., LTD., Intelligent Systems Laboratory T. Ito
Intended status: Informational October 30, 2017
Expires: May 3, 2018

Use of Name Redaction for Mass Devices
draft-ito-yet-another-name-redaction-00

Abstract

This document describes mechanisms to allow CT log submitters to submit name redacted certificates. While public Certificate Transparency (CT) logs allow anyone to observe server certificates and make confident to trust Certificate Authorities (CAs), there are some problems scaling to mass devices. This document presents some use cases, and describes a use of name redacted certificates that retains most of the security benefits that gained from using Certificate Transparency.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 3, 2018.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Requirements Language	3
3. Terminology	3
4. Redacted CT submission mechanism	3
4.1. Using Wildcard Certificates	4
4.2. Using a Name-Constrained Intermediate CA	4
4.2.1. Presenting SCTs, Inclusion Proofs and STHs	5
4.2.2. Matching an SCT to the Correct Certificate	5
4.3. Redacting Labels in Precertificates	6
4.3.1. redactedSubjectAltName Certificate Extension	6
4.3.2. Verifying the redactedSubjectAltName extension	7
4.3.3. Reconstructing the TBSCertificate	7
5. IANA Considerations	8
6. Security Considerations	8
7. Acknowledgements	8
8. References	8
8.1. Normative References	8
8.2. Informative References	9
Author's Address	9

1. Introduction

*****Scope/position of this document will be discussed at IETF 100 Singapore*****

Many devices communicate with TLS. These devices include surveillance cameras and Network Attached Storage. Such devices use server certificates to communicate with other devices such as smart phones. The number of these TLS-communicating devices is expected to grow exponentially. In contrast, searchable mass devices may assist attackers (typically, to construct a botnet). In this document, I describe needs of name redaction for those devices' certificates. Their certificates are typically issued by an intermediate certificate authority, which is tied to the device vendor or service provider.

On the other hand, there are some organizations who issue certificates only for their own domain space (with global IP address). For that case, CA/BForum defines "technical constrained intermediate certificate authority", and allows organizations to moderate portions of the audit process CA/BForum BR1.5.1

[EV.Certificate.Guidelines], according to limitation of influence in case of miss issuance.

However, Certificate Transparency v1 [RFC6962] and current v2 I-D.ietf-trans-rfc6962-bis26 [I-D.ietf-trans-rfc6962-bis] describe protocols for publicly logging all TLS server certificates issued by publicly trusted CAs. CT log server also store certificates with above uses, and can end up assisting attacker in hijacking massive numbers of devices. In addition, it would increase burden of CT log server near future, by exponential increase of mass devices.

I-D.draft-strad-trans-redaction-01

[I-D.draft-strad-trans-redaction-01] focused on end-entity's privacy with name redaction. This document focuses on other aspects, such as avoiding lack of scalability, or prohibiting use on large scale Botnet. I believe this document will reinforce discussion of I-D.draft-strad-trans-redaction-01 [I-D.draft-strad-trans-redaction-01].

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 RFC 2119 [RFC2119].

3. Terminology

This document relies on terminology and data structures defined in [RFC-6962-BIS-26], including STH, SCT.

4. Redacted CT submission mechanism

The technical part of this section refers to I-D.draft-strad-trans-redaction-01 [I-D.draft-strad-trans-redaction-01], since its mechanisms are directly applicable to this document.

I describe the device scalability and security for three name redaction mechanisms, in order of increasing implementation complexity:

- o Using wildcard certificates (Section 4.1) is the simplest option, but is not suitable for use with massive numbers of devices. Devices with a common wildcard certificate would need to share a private key, which would increase risk of key leakage dramatically.
- o Logging a name-constrained intermediate CA certificate in place of the end-entity certificate (Section 4.2) covers more, and is

suitable for mass devices, according to scalability of Log server. However it requires some non-scalable operations for CA (i.e. issuing new intermediate certificate.).

- o Domain label redaction mechanism (Section 4.3) reduces the burden on CA's operation. In addition, flexible operation of mass devices will become possible for CAs. Furthermore, geometric or geographic information is very useful for managing mass device, and service providers may want or try to use that information with certificates. Therefore, if it were without a "name constrained intermediate" mechanism, this mechanism might be needed to prevent large-scale physical attacks on devices with geometric or geographic information. However, this option increases the implementation complexity considerably.

*****the rest of this section is parts of I-D.draft-strad-trans-redaction-01 [I-D.draft-strad-trans-redaction-01] *****

4.1. Using Wildcard Certificates

A certificate containing a DNS-ID [RFC6125] of "*.example.com" could be used to secure the devices under some domain "topsecret.example.com"

4.2. Using a Name-Constrained Intermediate CA

An intermediate CA certificate or intermediate CA precertificate that contains the Name Constraints [RFC5280] extension MAY be logged in place of end-entity certificates issued by that intermediate CA, as long as all of the following conditions are met:

- o there MUST be a non-critical extension (OID 1.3.101.76, whose extnValue OCTET STRING contains ASN.1 NULL data (0x05 0x00)). This extension is an explicit indication that it is acceptable to not log certificates issued by this intermediate CA.
- o there MUST be a Name Constraints extension, in which:
 - * permittedSubtrees MUST specify one or more dNSNames.
 - * excludedSubtrees MUST specify the entire IPv4 and IPv6 address ranges.

Below is an example Name Constraints extension that meets these conditions:

```

SEQUENCE {
  OBJECT IDENTIFIER '2 5 29 30'
  BOOLEAN TRUE
  OCTET STRING, encapsulates {
    SEQUENCE {
      [0] {
        SEQUENCE {
          [2] 'example.com'
        }
      }
      [1] {
        SEQUENCE {
          [7] 00 00 00 00 00 00 00 00
        }
        SEQUENCE {
          [7]
            00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
            00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
        }
      }
    }
  }
}

```

4.2.1. Presenting SCTs, Inclusion Proofs and STHs

Each SCT (and optional corresponding inclusion proof and STH) presented by a TLS server, or included by a certification authority in a Transparency Information X.509v3 extension in the "singleExtensions" of a "SingleResponse" in an OCSP response MAY correspond to an intermediate CA certificate or intermediate CA precertificate (to which the server certificate chains) that meets the requirements in Section 4.2. This extends section TBD of CT v2 [I-D.ietf-trans-rfc6962-bis], which specifies that SCT of intermediate Certificate always corresponds to the server certificates or to a precertificates that corresponds to that certificate.

4.2.2. Matching an SCT to the Correct Certificate

Before considering any SCT to be invalid, a TLS client MUST attempt to validate it against the server certificate and against each of the zero or more suitable name-constrained intermediates in the chain. These certificates may be evaluated in the order they appear in the chain, or indeed, in any order.

4.3. Redacting Labels in Precertificates

When creating a precertificate, the CA MAY include a `redactedSubjectAltName` (Section 4.3.1) 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".

4.3.1. `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 4.3).

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 4.3.3). Hashing the redacted labels allows the legitimate domain owner to identify whether or not each redacted label correlates to a label they know of.

Only DNS-ID labels can be redacted using this mechanism. However, CAs can use the Section 4.2 mechanism to allow DNS domain name labels in other `subjectAltName` entries to not appear in logs.

4.3.2. 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:

- o The two entries are identical; or
- o 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 4.3).

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

4.3.3. 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 4.3.1) is present in the certificate, TLS clients MUST also:

- o Verify the redactedSubjectAltName extension against the subjectAltName extension according to Section 4.3.2.
- o Once verified, remove the subjectAltName extension from the TBSCertificate.

5. IANA Considerations

TBD

6. Security Considerations

TODO: describe how CA can get assurance for domain owner's control over underlying domain. It should contain some management mechanism, and need further discuss.

7. Acknowledgements

Portions of this text were unabashedly borrowed from I-D.draft-strad-trans-redaction-01 [I-D.draft-strad-trans-redaction-01].

8. References

8.1. Normative References

- [I-D.draft-strad-trans-redaction-01]
Stradling, R. and E. Messeri, "Certificate Transparency: Domain Label Redaction", draft-strad-trans-redaction-01 (work in progress), January 2017.
- [I-D.ietf-trans-rfc6962-bis]
Laurie, B., Langley, A., Kasper, E., Messeri, E., and R. Stradling, "Certificate Transparency Version 2.0", draft-ietf-trans-rfc6962-bis-24 (work in progress), December 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>>.
- [RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", RFC 6962, DOI 10.17487/RFC6962, June 2013, <<http://www.rfc-editor.org/info/rfc6962>>.

8.2. Informative References

- [EV.Certificate.Guidelines]
CA/Browser Forum, "Guidelines For The Issuance And Management Of Extended Validation Certificates Version 1.6.6", 2017, <https://cabforum.org/wp-content/uploads/EV-V1_6_6.pdf>.

Author's Address

Tadahiko Ito
SECOM CO., LTD., Intelligent Systems Laboratory
Mitaka, Tokyo
Japan

Phone: +81 422 76 2111
Email: tadahi-ito@secom.co.jp