

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 29 January 2026

D. Benjamin
D. O'Brien
Google LLC
B. E. Westerbaan
L. Valenta
Cloudflare
F. Valsorda
Geomys
28 July 2025

Merkle Tree Certificates
draft-davidben-tls-merkle-tree-certs-06

Abstract

This document describes Merkle Tree certificates, a new form of X.509 certificates which integrate public logging of the certificate, in the style of Certificate Transparency. The integrated design reduces logging overhead in the face of both shorter-lived certificates and large post-quantum signature algorithms, while still achieving comparable security properties to traditional X.509 and Certificate Transparency. Merkle Tree certificates additionally admit an optional signatureless optimization, which decreases the message size by avoiding signatures altogether, at the cost of only applying to up-to-date relying parties and older certificates.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at <https://davidben.github.io/merkle-tree-certs/draft-davidben-tls-merkle-tree-certs.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-davidben-tls-merkle-tree-certs/>.

Source for this draft and an issue tracker can be found at <https://github.com/davidben/merkle-tree-certs>.

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 29 January 2026.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	4
2. Conventions and Definitions	5
2.1. Terminology and Roles	6
3. Overview	7
4. Subtrees	9
4.1. Definition of a Subtree	9
4.2. Subtree Inclusion Proofs	9
4.3. Subtree Consistency Proofs	10
4.3.1. Generating a Subtree Consistency Proof	10
4.3.2. Verifying a Subtree Consistency Proof	11
4.4. Arbitrary Intervals	12
5. Issuance Logs	15
5.1. Log Parameters	15
5.2. Log IDs	16
5.3. Log Entries	16
5.4. Cosigners	17
5.4.1. Signature Format	18
5.4.2. Signature Algorithms	20
5.5. Certification Authority Cosigners	20
5.6. Publishing Logs	21

5.6.1. Log Pruning	22
6. Certificates	25
6.1. Certificate Format	25
6.2. Full Certificates	27
6.3. Signatureless Certificates	28
6.3.1. Landmark Checkpoints	28
6.3.2. Constructing Signatureless Certificates	30
6.4. Size Estimates	31
7. Relying Parties	32
7.1. Trust Anchors	32
7.2. Verifying Certificate Signatures	32
7.3. Trusted Cosigners	33
7.4. Trusted Subtrees	35
7.5. Revocation by Index	36
8. Use in TLS	37
9. ACME Extensions	38
10. Deployment Considerations	38
10.1. Operational Costs	38
10.1.1. Certification Authority Costs	39
10.1.2. Cosigner Costs	39
10.1.3. Monitor Costs	40
10.2. Choosing Cosigners	40
10.3. Log Availability	41
10.4. Certificate Renewal	42
10.5. Multiple CA Keys	43
11. Privacy Considerations	43
12. Security Considerations	44
12.1. Authenticity	44
12.2. Transparency	44
12.3. Public Key Hashes	45
12.4. Non-Repudiation	46
12.5. New Log Entry Types	46
12.6. Certificate Malleability	47
13. IANA Considerations	49
14. References	49
14.1. Normative References	49
14.2. Informative References	51
Appendix A. ASN.1 Module	53
Appendix B. Extensions to Tiled Transparency Logs (To Be Removed)	53
B.1. Pruning a Tiled Transparency Log	53
B.2. Subtree Signed Note Format	53
B.3. Requesting Subtree Signatures	54
Acknowledgements	56
Change log	56
Since draft-davidben-tls-merkle-tree-certs-00	56
Since draft-davidben-tls-merkle-tree-certs-01	57
Since draft-davidben-tls-merkle-tree-certs-02	57

Since draft-davidben-tls-merkle-tree-certs-03	57
Since draft-davidben-tls-merkle-tree-certs-04	57
Since draft-davidben-tls-merkle-tree-certs-05	58
Authors' Addresses	58

1. Introduction

Authors' Note: This is an early draft of a proposal with many parts. We expect most details will change as the proposal evolves. This document has a concrete specification of these details, but this is only intended as a starting point, and to help convey the overall idea. The name of the draft says "tls" to keep continuity with earlier iterations of this work, but the protocol itself is not TLS-specific.

In Public Key Infrastructures (PKIs) that use Certificate Transparency (CT) [RFC6962] for a public logging requirement, an authenticating party must present Signed Certificate Timestamps (SCTs) alongside certificates. CT policies often require two or more SCTs per certificate [APPLE-CT] [CHROME-CT], each of which carries a signature. These signatures are in addition to those in the certificate chain itself.

Current signature schemes can use as few as 32 bytes per key and 64 bytes per signature [RFC8032], but post-quantum replacements are much larger. For example, ML-DSA-44 [FIPS204] uses 1,312 bytes per public key and 2,420 bytes per signature. ML-DSA-65 uses 1,952 bytes per public key and 3,309 bytes per signature. Even with a directly-trusted intermediate (Section 7.5 of [I-D.ietf-tls-trust-anchor-ids]), two SCTs and a leaf certificate signature adds 7,260 bytes of authentication overhead with ML-DSA-44 and 9,927 bytes with ML-DSA-65.

This increased overhead additionally impacts CT logs themselves. Most of a log's costs scale with the total storage size of the log. Each log entry contains both a public key, and a signature from the CA. With larger public keys and signatures, the size of each log entry will grow.

Additionally, as PKIs transition to shorter-lived certificates [CABF-153] [CABF-SC081], the number of entries in the log will grow.

This document introduces Merkle Tree certificates, a new form of X.509 certificate that integrates logging with certificate issuance. Each CA maintains a log of everything it issues, signing views of the log to assert it has issued the contents. The CA signature is combined with cosignatures from other parties who verify correct operation and optionally mirror the log. These signatures, together with an inclusion proof for an individual entry, constitute a certificate.

This achieves the following:

- * Log entries do not scale with public key and signature sizes. Entries replace public keys with hashes and do not contain signatures, while preserving non-repudiability (Section 12.4).
- * To bound growth, long-expired entries can be pruned from logs and mirrors without interrupting existing clients. This allows log sizes to scale by retention policies, not the lifetime of the log, even as certificate lifetimes decrease.
- * After a processing delay, authenticating parties can obtain a second "signatureless" certificate for the same log entry. This second certificate is an optional size optimization that avoids the need for any signatures, assuming an up-to-date client that has some predistributed log information.

Section 3 gives an overview of the system. Section 4 describes a Merkle Tree primitive used by this system. Section 5 describes the log structure. Finally, Section 6 and Section 7 describes how to construct and consume a Merkle Tree certificate.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

This document additionally uses the TLS presentation language defined in Section 3 of [RFC8446].

U+ followed by four hexadecimal characters denotes a Unicode codepoint, to be encoded in UTF-8 [RFC3629]. 0x followed by two hexadecimal characters denotes a byte value in the 0-255 range.

[start, end), where start <= end, denotes the half-open interval containing integers x such that start <= x < end.

2.1. Terminology and Roles

This document discusses the following roles:

Authenticating party: The party that authenticates itself in the protocol. In TLS, this is the side sending the Certificate and CertificateVerify message.

Certification authority (CA): The service that issues certificates to the authenticating party, after performing some validation process on the certificate contents.

Relying party: The party to whom the authenticating party presents its identity. In TLS, this is the side receiving the Certificate and CertificateVerify message.

Monitor: Parties who watch logs for certificates of interest, analogous to the role in Section 8.2 of [RFC9162].

Issuance log: A log, maintained by the CA, of everything issued by that CA.

Cosigner: A service that signs views of an issuance log, to assert correct operation and other properties about the entries.

Additionally, there are several terms used throughout this document to describe this proposal. This section provides an overview. They will be further defined and discussed in detail throughout the document.

Checkpoint: A description of the complete state of the log at some time.

Entry: An individual element of the log, describing information which the CA has validated and certified.

Subtree: A smaller Merkle Tree over a portion of the log, defined by an interior node of some snapshot of the log. Subtrees can be efficiently shown to be consistent with the whole log.

Inclusion proof: A sequence of hashes that efficiently proves some entry is contained in some checkpoint or subtree.

Consistency proof: A sequence of hashes that efficiently proves a checkpoint or subtree is contained within another checkpoint.

Cosignature: A signature from either the CA or other cosigner, over some checkpoint or subtree.

Landmark checkpoint: One of an infrequent subset of checkpoints that can be used to predistribute trusted subtrees to relying parties for signatureless certificates.

Landmark subtree: A subtree determined by a landmark checkpoint. Landmark subtrees are common points of reference between relying parties and signatureless certificates.

Full certificate: A certificate containing an inclusion proof to some subtree, and several cosignatures over that subtree.

Signatureless certificate: An optimized certificate containing an inclusion proof to a landmark subtree, and no signatures.

3. Overview

Merkle Tree certificate issuance begins as in a typical CA: The CA receives an issuance request, e.g. over ACME [RFC8555], and validates it, e.g. with ACME challenges. From there, the issuance process differs.

1. The CA operates an `_issuance log_` (Section 5), which is an append-only Merkle Tree (Section 2.1 of [RFC9162]). Unlike a CT log, this log only contains entries added by the CA.
2. The CA adds a `TBSCertificateLogEntry` (Section 5.3) to its log, describing the information it is certifying.
3. The CA signs a `_checkpoint_`, which describes the current state of the log. A signed checkpoint certifies that the CA issued `_every_` entry in the Merkle Tree (Section 5.5).
4. The CA additionally signs `_subtrees_` (Section 4) that together contain certificates added since the last checkpoint (Section 4.4). This is an optimization to reduce inclusion proof sizes. A signed subtree certifies that the CA has issued `_every_` entry in the subtree.

These signatures prove to relying parties that the CA has issued the certificate.

A relying party may require additional proof of the proper operation of the issuance log by third parties, like CT logs provide additional transparency. Instead of presenting certificates to transparency logs, CAs present their log state to `_cosigners_`. Cosigners validate the log is append-only and optionally provide additional services, such as mirroring its contents. They cosign the CA's checkpoints and subtrees.

A signature from either the CA or a cosigner is known as a `_cosignature_`. A combination of cosignatures from sufficient trusted parties proves both that the CA has issued the certificate, and that the issuance satisfies any additional relying party transparency requirements.

A `_full certificate_` contains:

- * The TBSCertificate being certified
- * An inclusion proof from the TBSCertificate to some subtree
- * Sufficient cosignatures to meet the relying party's requirements

This same issuance process also produces a `_signatureless certificate_`. This is an optional optimized certificate that avoids any cosignatures:

1. Periodically, the CA's most recent checkpoint is designated as a `_landmark checkpoint_`. This determines `_landmark subtrees_`, which are common points of reference between relying parties and signatureless certificates.
2. Once some landmark checkpoint includes the TBSCertificate, the signatureless certificate is constructed with:
 - * The TBSCertificate being certified
 - * An inclusion proof from the TBSCertificate to a landmark subtree
3. In the background, landmark subtrees are predistributed to relying parties, with cosignatures checked against relying party requirements. This occurs periodically in the background, separate from the application protocol.
4. During the application protocol, such as TLS [RFC8446], if the relying party already supports the landmark subtree, the authenticating party can present the signatureless certificate. Otherwise, it presents a full certificate.

4. Subtrees

This section extends the Merkle Tree definition in Section 2.1 of [RFC9162] by defining a `_subtree_` of a Merkle Tree. A subtree is an interior node of a Merkle Tree, which can be efficiently shown consistent with the original Merkle Tree and any Merkle Tree with additional elements appended. This specification uses subtrees to reduce the size of inclusion proofs.

4.1. Definition of a Subtree

Given an ordered list of n inputs, $D_n = \{d[0], d[1], \dots, d[n-1]\}$, Section 2.1.1 of [RFC9162] defines the Merkle Tree via the Merkle Tree Hash $MTH(D_n)$.

A `_subtree_` of this Merkle Tree is defined by two integers, `start` and `end`, such that:

- * $0 \leq \text{start} < \text{end} \leq n$

- * If k is the largest power of 2 that is greater than or equal to $\text{end} - \text{start}$, `start` must be a multiple of k

Note that, if `start` is zero, the second condition is always true.

A subtree is itself a Merkle Tree, defined by $MTH(D[\text{start}:\text{end}])$. In the context of a single Merkle Tree, the subtree defined by `start` and `end` is denoted by half-open interval $[\text{start}, \text{end})$. It contains the entries whose indices are in that half-open interval.

The `_size_` of the subtree is $\text{end} - \text{start}$. If the subtree's size is a power of 2, it is said to be `_full_`, otherwise it is said to be `_partial_`.

If a subtree is full, then it is directly contained in the tree of hash operations in $MTH(D_n)$ for $n \geq \text{end}$.

If a subtree is partial, it is directly contained in $MTH(D_n)$ only if $n = \text{end}$.

4.2. Subtree Inclusion Proofs

Subtrees are Merkle Trees, so entries can be proven to be contained in the subtree. A subtree inclusion proof for entry index of the subtree $[\text{start}, \text{end})$ is a Merkle inclusion proof, as defined in Section 2.1.3.1 of [RFC9162], where m is $\text{index} - \text{start}$ and the tree inputs are $D[\text{start}:\text{end}]$. A subtree inclusion proof can be verified with the procedure in Section 2.1.3.2 of [RFC9162], where `leaf_index`

is $\text{index} - \text{start}$ and tree_size is $\text{end} - \text{start}$.

4.3. Subtree Consistency Proofs

A subtree $[\text{start}, \text{end})$ can be efficiently proven to be consistent with the full Merkle Tree. That is, given $\text{MTH}(\text{D}[\text{start}:\text{end}])$ and $\text{MTH}(\text{D}_n)$, the proof demonstrates that the input $\text{D}[\text{start}:\text{end}]$ to the subtree hash was equal to the corresponding elements of the input D_n to the Merkle Tree hash.

4.3.1. Generating a Subtree Consistency Proof

The subtree consistency proof, $\text{SUBTREE_PROOF}(\text{start}, \text{end}, \text{D}_n)$ is defined similarly to Section 2.1.4.1 of [RFC9162], in terms of a helper function that tracks whether the subtree hash is known:

```
SUBTREE_PROOF(start, end, D_n) =
    SUBTREE_SUBPROOF(start, end, D_n, true)
```

If $\text{start} = 0$ and $\text{end} = n$, the subtree is the root:

```
SUBTREE_SUBPROOF(0, n, D_n, true) = {}
SUBTREE_SUBPROOF(0, n, D_n, false) = {MTH(D_n)}
```

Otherwise, $n > 1$. Let k be the largest power of two smaller than n . The consistency proof is defined recursively as:

- * If $\text{end} \leq k$, the subtree is on the left of k . The proof proves consistency with the left child and includes the right child:

```
SUBTREE_SUBPROOF(start, end, D_n, b) =
    SUBTREE_SUBPROOF(start, end, D[0:k], b) : MTH(D[k:n])
```

- * If $k \leq \text{start}$, the subtree is on the right of k . The proof proves consistency with the right child and includes the left child.

```
SUBTREE_SUBPROOF(start, end, D_n, b) =
    SUBTREE_SUBPROOF(start - k, end - k, D[k:n], b) : MTH(D[0:k])
```

- * Otherwise, $\text{start} < k < \text{end}$, which implies $\text{start} = 0$. The proof proves consistency with the right child and includes the left child.

```
SUBTREE_SUBPROOF(0, end, D_n, b) =
    SUBTREE_SUBPROOF(0, end - k, D[k:n], false) : MTH(D[0:k])
```

When start is zero, this computes a Merkle consistency proof:

`SUBTREE_PROOF(0, end, D_n) = PROOF(end, D_n)`

When `end = start + 1`, this computes a Merkle inclusion proof:

`SUBTREE_PROOF(start, start + 1, D_n) = PATH(start, D_n)`

4.3.2. Verifying a Subtree Consistency Proof

The following procedure can be used to verify a subtree consistency proof.

Given a Merkle Tree over `n` elements, a subtree defined by [`start`, `end`), a consistency proof `proof`, a subtree hash `node_hash`, and a root hash `root_hash`:

1. If `end` is `n`, run the following:
 1. Set `fn` to `start` and `sn` to `end - 1`.
 2. Set `r` to `node_hash`.
 3. Right-shift `fn` and `sn` equally until `LSB(fn)` is set or `sn` is zero.
 4. For each value `p` in the proof array:
 1. If `sn` is 0, then stop iteration and fail the proof verification.
 2. Set `r` to `HASH(0x01, || p || r)`.
 3. If `LSB(sn)` is not set, the right-shift `sn` until either `LSB(sn)` is set or `sn` is zero.
 4. Right-shift `sn` once more.
 5. Check `sn` is 0 and `r` is `root_hash`. If either is not equal, fail the proof verification. If all are equal, accept the proof.
2. Otherwise, run the following:
 1. If `proof` is an empty array, stop and fail verification.
 2. If `end - start` is an exact power of 2, prepend `node_hash` to the proof array.
 3. Set `fn` to `start`, `sn` to `end - 1`, and `tn` to `n - 1`.

4. Right-shift `fn`, `sn`, and `tn` equally until `LSB(sn)` is not set or `fn = sn`.
5. Set both `fr` and `sr` to the first value in the proof array.
6. For each subsequent value `c` in the proof array:
 1. If `tn` is 0, then stop the iteration and fail the proof verification.
 2. If `LSB(sn)` is set, or if `sn` is equal to `tn`, then:
 1. If `fn < sn`, set `fr` to `HASH(0x01 || c || fr)`.
 2. Set `sr` to `HASH(0x01 || c || sr)`.
 3. If `LSB(sn)` is not set, then right-shift each of `fn`, `sn`, and `tn` equally until either `LSB(sn)` is set or `sn` is 0.
 3. Otherwise:
 1. Set `sr` to `HASH(0x01 || sr || c)`.
 4. Finally, right-shift each of `fn`, `sn`, and `tn` one time.
7. Check `tn` is 0, `fr` is `node_hash`, and `sr` is `root_hash`. If any are not equal, fail the proof verification. If all are equal, accept the proof.

4.4. Arbitrary Intervals

Not all `[start, end)` intervals of a Merkle tree are valid subtrees. This section describes how, for any `start < end`, to determine up to two subtrees that efficiently cover the interval. The subtrees are determined by the following Python procedure:

```
def find_subtrees(start, end):
    """ Returns a list of one or two subtrees that efficiently
    cover [start, end). """
    assert start < end
    if end - start < 2:
        return [(start, end),]
    last = end - 1
    # Find where start and last's tree paths diverge. The two
    # subtrees will be on either side of the split.
    split = (start ^ last).bit_length() - 1
    mask = (1 << split) - 1
    mid = last & ~mask
    # Maximize the left endpoint. This is just before start's
    # path leaves the right edge of its new subtree.
    left_split = (~start & mask).bit_length()
    left = start & ~(1 << left_split) - 1
    return [(left, mid), (mid, end)]
```

[[TODO: Write this up in prose too.]]

When the procedure returns a single subtree, the subtree is [start, end). When it returns two subtrees, left and right, the subtrees satisfy the following properties:

1. left.end = right.start. That is, the two subtrees cover adjacent intervals.
2. left.start <= start and end = right.end. That is, the two subtrees together cover the entire target interval, possibly with some extra entries before start left, but not after end.
3. left.end - left.start < 2 * (end - start) and right.end - right.start <= end - start. That is, the two subtrees efficiently cover the interval.
4. left is full, while right may be partial.

Figure 1 shows the subtrees which cover [5, 13) in a Merkle Tree of 13 elements. The two subtrees selected are [4, 8) and [8, 13). Note that the subtrees cover a slightly larger interval than [5, 13).

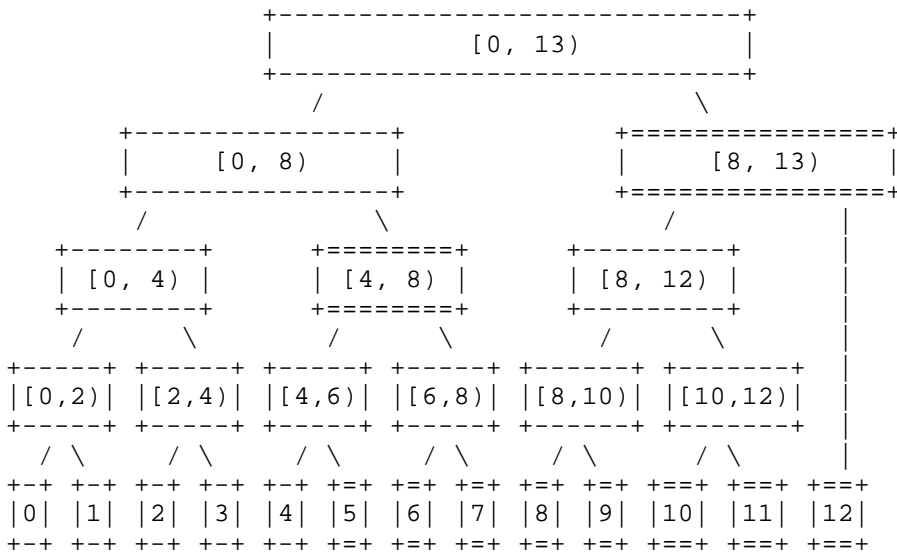


Figure 1: An example selection of subtrees to cover an interval

Two subtrees are needed because a single subtree may not be able to efficiently cover an interval. Figure 2 shows the smallest subtree that contains $[7, 9)$ in a 9-element tree. The smallest single subtree that contains the interval is $[0, 9)$ but this is the entire tree. Using two subtrees, the interval can be described by $[7, 8)$ and $[8, 9)$.




Figure 2: An example showing an inefficient choice of a single subtree

5. Issuance Logs

This section defines the structure of an `_issuance log_`.

An issuance log describes an append-only sequence of `_entries_` (Section 5.3), identified consecutively by an index value, starting from zero. Each entry is an assertion that the CA has certified. The entries in the issuance log are represented as a Merkle Tree, described in Section 2.1 of [RFC9162].

Unlike [RFC6962] and [RFC9162], an issuance log does not have a public submission interface. The log only contains entries which the log operator, i.e. the CA, chose to add. As entries are added, the Merkle Tree is updated to be computed over the new sequence.

A snapshot of the log is known as a `_checkpoint_`. A checkpoint is identified by its `_tree size_`, that is the number of elements committed to the log at the time. Its contents can be described by the Merkle Tree Hash (Section 2.1.1 of [RFC9162]) of entries zero through `tree_size - 1`.

Cosigners (Section 5.4) sign assertions about the state of the issuance log. A Merkle Tree CA operates a combination of an issuance log and one or more CA cosigners (Section 5.5) that authenticate the log state and certifies the contents. External cosigners may also be deployed to assert correct log operation or provide other services to relying parties (Section 7.3).

5.1. Log Parameters

An issuance log has the following parameters:

- * A log ID, which uniquely identifies the log. See Section 5.2.
- * A collision-resistant cryptographic hash function. SHA-256 [SHS] is RECOMMENDED.
- * A minimum index, which is the index of the first log entry which is available. See Section 5.6.1. This value changes over the lifetime of the log.

Throughout this document, the hash algorithm in use is referred to as `HASH`, and the size of its output in bytes is referred to as `HASH_SIZE`.

5.2. Log IDs

Each issuance log is identified by a `_log ID_`, which is a trust anchor ID [I-D.ietf-tls-trust-anchor-ids].

An issuance log's log ID determines an X.509 distinguished name (Section 4.1.2.4 of [RFC5280]). The distinguished name has a single relative distinguished name, which has a single attribute. The attribute has type `id-rdna-trustAnchorID`, defined below:

```
id-rdna-trustAnchorID OBJECT IDENTIFIER ::= {
    iso(1) identified-organization(3) dod(6) internet(1) security(5)
    mechanisms(5) pkix(7) rdna(TBD1) TBD2}
```

[[TODO: Fill in TBD1 from the PKIX rdna arc in
[I-D.draft-ietf-lamps-x509-alg-none], once allocated.]]

The attribute's value is a RELATIVE-OID containing the trust anchor ID's ASN.1 representation. For example, the distinguished name for a log named 32473.1 would be represented in syntax of [RFC4514] as:

```
1.3.6.1.5.5.7.TBD1.TBD2=#0d0481fd5901
```

For initial experimentation, early implementations of this design will use the OID 1.3.6.1.4.1.44363.47.1 instead of `id-rdna-trustAnchorID`.

5.3. Log Entries

Each entry in the log is a `MerkleTreeCertEntry`, defined with the TLS presentation syntax below. A `MerkleTreeCertEntry` describes certificate information that the CA has validated and certified.

```
struct {} Empty;

enum {
    null_entry(0), tbs_cert_entry(1), (2^16-1)
} MerkleTreeCertEntryType;

struct {
    MerkleTreeCertEntryType type;
    select (type) {
        case null_entry: Empty;
        case tbs_cert_entry: opaque tbs_cert_entry_data[N];
        /* May be extended with future types. */
    }
} MerkleTreeCertEntry;
```


When type is `tbs_cert_entry`, `N` is the number of bytes needed to consume the rest of the input. A `MerkleTreeCertEntry` is expected to be decoded in contexts where the total length of the entry is known.

`tbs_cert_entry_data` contains the DER [X.690] encoding of a `TBSCertificateLogEntry`, defined below:

```
TBSCertificateLogEntry ::= SEQUENCE {  
    version          [0] EXPLICIT Version DEFAULT v1,  
    issuer            Name,  
    validity          Validity,  
    subject           Name,  
    subjectPublicKeyInfoHash OCTET STRING,  
    issuerUniqueID    [1] IMPLICIT UniqueIdentifier OPTIONAL,  
    subjectUniqueID   [2] IMPLICIT UniqueIdentifier OPTIONAL,  
    extensions        [3] EXPLICIT Extensions OPTIONAL }
```

The version, issuer, validity, subject, issuerUniqueID, subjectUniqueID, and extensions fields have the corresponding semantics as in Section 4.1.2 of [RFC5280], with the exception of subjectPublicKeyInfoHash. subjectPublicKeyInfoHash contains the hash of subject's public key as a SubjectPublicKeyInfo (Section 4.1.2.7 of [RFC5280]). The hash uses the log's hash function (Section 5.1) and is computed over the SubjectPublicKeyInfo's DER [X.690] encoding. The issuer field MUST be the issuance log's log ID as an X.509 distinguished name, as described in Section 5.2.

When type is `null_entry`, the entry does not represent any information. The entry at index zero of every issuance log MUST be of type `null_entry`. Other entries MUST NOT use `null_entry`. `null_entry` exists to avoid zero serial numbers in the certificate format (Section 6.1).

`MerkleTreeCertEntry` is an extensible structure. Future documents may define new values for `MerkleTreeCertEntryType`, with corresponding semantics. See Section 5.5 and Section 12.5 for additional discussion.

5.4. Cosigners

This section defines a log `_cosigner_`. A cosigner follows some append-only view of the log and signs subtrees (Section 4) consistent with that view. The signatures generated by a cosigner are known as `_cosignatures_`. All subtrees signed by a cosigner MUST be consistent with each other. The cosigner may be external to the log, in which case it might ensure consistency by checking consistency proofs. The cosigner may be operated together with the log, in which case it can trust its log state.

A cosignature MAY implicitly make additional statements about a subtree, determined by the cosigner's role. This document defines one concrete cosigner role, a CA cosigner (Section 5.5), to authenticate the log and certify entries. Other documents and specific deployments may define other cosigner roles, to perform different functions in a PKI. For example, [TLOG-WITNESS] defines a cosigner that only checks the log is append-only, and [TLOG-MIRROR] defines a cosigner that mirrors a log.

Each cosigner has a public key and a `_cosigner ID_`, which uniquely identifies the cosigner. The cosigner ID is a trust anchor ID [I-D.ietf-tls-trust-anchor-ids]. By identifying the cosigner, the cosigner ID specifies both the public key and the additional statements made by the cosigner's signatures. If a single operator performs multiple cosigner roles in an ecosystem, each role MUST use a distinct cosigner ID and SHOULD use a distinct key.

A single cosigner, with a single cosigner ID and public key, MAY generate cosignatures for multiple logs. In this case, signed subtrees only need to be consistent with others for the same log.

5.4.1. Signature Format

A cosigner computes a cosignature for a subtree in some log by signing a `MTCSubtreeSignatureInput`, defined below using the TLS presentation language (Section 3 of [RFC8446]):

```
opaque HashValue[HASH_SIZE];

/* From Section 4.1 of draft-ietf-tls-trust-anchor-ids */
opaque TrustAnchorID<1..2^8-1>;

struct {
    TrustAnchorID log_id;
    uint64 start;
    uint64 end;
    HashValue hash;
} MTCSubtree;

struct {
    uint8 label[16] = "mtc-subtree/v1\n\0";
    TrustAnchorID cosigner_id;
    MTCSubtree subtree;
} MTCSubtreeSignatureInput;
```

`log_id` MUST be the issuance log's ID (Section 5.2), in its binary representation (Section 3 of [I-D.ietf-tls-trust-anchor-ids]). `start` and `end` MUST define a valid subtree of the log, and `hash` MUST be the

subtree's hash value in the cosigner's view of the log. The label is a fixed prefix for domain separation. Its value MUST be the string `mtc-subtree/v1`, followed by a newline (U+000A), followed by a zero byte (U+0000). `cosigner_id` MUST be the cosigner ID, in its binary representation.

The resulting signature is known as a `_subtree signature_`. When `start` is zero, the resulting signature describes the checkpoint with `tree_size` end and is also known as a `_checkpoint signature_`.

For each supported log, a cosigner retains its checkpoint signature with the largest end. This is known as the cosigner's `_current_checkpoint`. If the cosigner's current checkpoint has `tree_size`, it MUST NOT generate a signature for a subtree [`start`, `end`) if `start > 0` and `end > tree_size`. That is, a cosigner can only sign a non-checkpoint subtree if it is contained in its current checkpoint. In a correctly-operated cosigner, every signature made by the cosigner can be proven consistent with its current checkpoint with a subtree consistency proof (Section 4.3). As a consequence, a cosigner that signs a subtree is held responsible for all the entries in the tree of size matching the subtree end, even if the corresponding checkpoint is erroneously unavailable.

Before signing a subtree, the cosigner MUST ensure that hash is consistent with its log state. Different cosigner roles may obtain this assurance differently. For example, a cosigner may compute the hash from its saved log state (e.g. if it is the log operator or maintains a copy of the log) or by verifying a subtree consistency proof (Section 4.3) from its current checkpoint. When a cosigner signs a subtree, it is held responsible `_both_` for the subtree being consistent with its other signatures, `_and_` for the cosigner-specific additional statements.

Cosigners SHOULD publish their current checkpoint, along with the checkpoint signature.

[[TODO: CT and tlog put timestamps in checkpoint signatures. Do we want them here? In CT and tlog, the timestamps are monotonically increasing as the log progresses, but we also sign subtrees. We can separate subtree and checkpoint signatures, with timestamps only in the latter, but it's unclear if there is any benefit to this.]]

5.4.2. Signature Algorithms

The cosigner's public key specifies both the key material and the signature algorithm to use with the key material. In order to change key or signature parameters, a cosigner operator MUST deploy a new cosigner, with a new cosigner ID. Signature algorithms MUST fully specify the algorithm parameters, such as hash functions used. This document defines the following signature algorithms:

- * ECDSA with P-256 and SHA-256 [FIPS186-5]
- * ECDSA with P-384 and SHA-384 [FIPS186-5]
- * Ed25519 [RFC8032]
- * ML-DSA-44 [FIPS204]
- * ML-DSA-65 [FIPS204]
- * ML-DSA-87 [FIPS204]

Other documents or deployments MAY define other signature schemes and formats. Log clients that accept cosignatures from some cosigner are assumed to be configured with all parameters necessary to verify that cosigner's signatures, including the signature algorithm and version of the signature format.

5.5. Certification Authority Cosigners

A `_CA cosigner_` is a cosigner (Section 5.4) that certifies the contents of a log.

When a CA cosigner signs a subtree, it makes the additional statement that it has certified each entry in the subtree. For example, a domain-validating CA states that it has performed domain validation for each entry, at some time consistent with the entry's validity dates. CAs are held responsible for every entry in every subtree they sign. Proving an entry is included (Section 4.2) in a CA-signed subtree is sufficient to prove the CA certified it.

What it means to certify an entry depends on the entry type:

- * To certify an entry of type `null_entry` is a no-op. A CA MAY freely certify `null_entry` without being held responsible for any validation.
- * To certify an entry of type `tbs_cert_entry` is to certify the `TBSCertificateLogEntry`, as defined in Section 5.3.

Entries are extensible. Future documents MAY define type values and what it means to certify them. A CA MUST NOT sign a subtree if it contains an entry with type that it does not recognize. Doing so would certify that the CA has validated the information in some not-yet-defined entry format. Section 12.5 further discusses security implications of new formats.

A CA operator MAY operate multiple CA cosigners that all certify the same log in parallel. This may be useful when, e.g., rotating CA keys. In this case, each CA instance MUST have a distinct name. The CA operator's ACME server can return all CA cosignatures together in a single certificate, with the application protocol selecting the cosignatures to use. Section 8 describes how this is done in TLS [RFC8446].

If the CA operator additionally operates a traditional X.509 CA, that CA key MUST be distinct from any Merkle Tree CA cosigner keys.

5.6. Publishing Logs

[[NOTE: This section is written to avoid depending on a specific serving protocol. The current expectation is that a Web PKI deployment would derive from [TLOG-TILES], to match the direction of Certificate Transparency and pick up improvements made there.]]

[For now, we avoid a normative reference on [TLOG-TILES] and also capture the fact that the certificate construction is independent of the choice of protocol. Similar to how the CT ecosystem is migrating to a tiled interface, were someone to improve on [TLOG-TILES], a PKI could migrate to that new protocol without impacting certificate verification.]]

[That said, this is purely a starting point for describing the design. We expect the scope of this document, and other related documents to adapt as the work evolves across the IETF, C2SP, Certificate Transparency, and other communities.]]

Issuance logs are intended to be publicly accessible in some form, to allow monitors to detect misissued certificates.

The access method does not affect certificate interoperability, so this document does not prescribe a specific protocol. An individual issuance log MAY be published in any form, provided other parties in the PKI are able to consume it. Relying parties SHOULD define log serving requirements, including the allowed protocols and expected availability, as part of their policies on which CAs to support. See also Section 10.3.

For example, a log ecosystem could use [TLOG-TILES] to serve logs. [TLOG-TILES] improves on [RFC6962] and [RFC9162] by exposing the log as a collection of cacheable, immutable "tiles". This works well with a variety of common HTTP [RFC9110] serving architectures. It also allows log clients to request arbitrary tree nodes, so log clients can fetch the structures described in Section 4.

5.6.1. Log Pruning

Over time, an issuance log's entries will expire and likely be replaced with certificate renewals. As this happens, the total size of the log grows, even if the unexpired subset remains fixed. To mitigate this, issuance logs MAY be pruned, as described in this section.

Pruning makes some prefix of the log unavailable, without changing the tree structure. It may be used to reduce the serving cost of long-lived logs, where any entries have long expired. Section 10.3 discusses policies on when pruning may be permitted. This section discusses how it is done and the impact on log structure.

An issuance log is pruned by updating its minimum index parameter (Section 5.1). The minimum index is the index of the first log entry that the log publishes. (See Section 5.6.) It MUST be less than or equal to the tree size of the log's current checkpoint, and also satisfy any availability policies set by relying parties who trust the CA.

An entry is said to be available if its index is greater than or equal to the minimum index. A checkpoint is said to be available if its tree size is greater than the minimum index. A subtree [start, end) is said to be available if end is greater than the minimum index.

Log protocols MUST serve enough information to allow a log client to efficiently obtain the following:

- * Signatures over the latest checkpoint by the CA's cosigners (Section 5.5)
- * Any individual available log entry (Section 5.3)
- * The hash value of any available checkpoint
- * An inclusion proof (Section 2.1.3 of [RFC9162]) for any available entry to any containing checkpoint

- * A consistency proof (Section 2.1.4 of [RFC9162]) between any two available checkpoints
- * The hash value of any available subtree (Section 4)
- * A subtree inclusion proof (Section 4.2) for any available entry in any containing subtree
- * A subtree consistency proof (Section 4.3) between any available subtree to any containing checkpoint

Meeting these requirements requires a log to retain some information about pruned entries. Given a node [start, end) in the Merkle Tree, if end is less than or equal to the minimum index, the node's children MAY be discarded in favor of the node's hash.

Figure 3 shows an example pruned tree with 13 elements, where the minimum index is 7. It shows the original tree, followed by the pruned tree. The pruned tree depicts the nodes that MUST be available or computable. Note that entry 6 MAY be discarded, only the hash of entry 6 must be available.

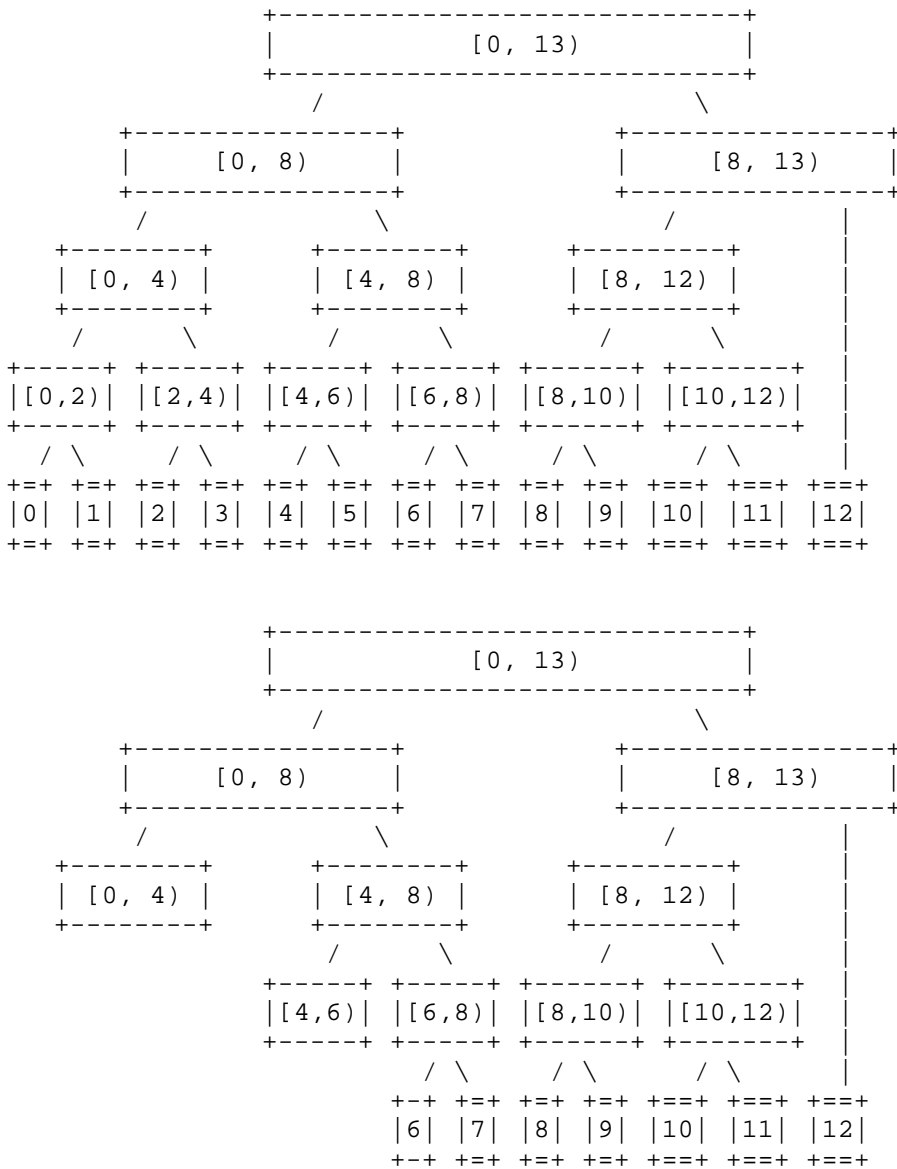


Figure 3: An example showing the minimum nodes that must be available after pruning

Logs MAY retain additional nodes, or expect log clients to compute required nodes from other nodes. For example, in Figure 3, the log's serving protocol MAY instead serve [0, 2) and [2, 4), with the log client computing [0, 4) from those values.

6. Certificates

This section defines how to construct Merkle Tree Certificates, which are X.509 Certificates [RFC5280] that assert the information in an issuance log entry. A Merkle Tree Certificate is constructed from the following:

- * A TBSCertificateLogEntry (Section 5.3) contained in the issuance log (Section 5)
- * A subject public key whose hash matches the TBSCertificateLogEntry
- * A subtree (Section 4) that contains the log entry
- * Zero or more signatures (Section 5.4) over the subtree, which together satisfy relying party requirements (Section 7.3)

For any given TBSCertificateLogEntry, there are multiple possible certificates that may prove the entry is certified by the CA and publicly logged, varying by choice of subtree and signatures. Section 6.1 defines how the certificate is constructed based on those choices. Section 6.2 and Section 6.3 define two profiles of Merkle Tree Certificates, full certificates and signatureless certificates, and how to select the subtree and signatures for them.

6.1. Certificate Format

The information is encoded in an X.509 Certificate [RFC5280] as follows:

The TBSCertificate's version, issuer, validity, subject, issuerUniqueID, subjectUniqueID, and extensions MUST match the corresponding fields of the TBSCertificateLogEntry. Per Section 5.3, this means issuer MUST be the issuance log's log ID as an X.509 distinguished name, as described in Section 5.2.

The TBSCertificate's serialNumber MUST contain the zero-based index of the TBSCertificateLogEntry in the log. Section 4.1.2.2 of [RFC5280] forbids zero as a serial number, but Section 5.3 defines a null_entry type for use in entry zero, so the index will be positive. This encoding is intended to avoid implementation errors by having the serial numbers and indices off by one.

The TBSCertificate's subjectPublicKeyInfo contains the specified public key. Its hash MUST match the TBSCertificateLogEntry's subjectPublicKeyInfoHash.

The TBSCertificate's signature and the Certificate's signatureAlgorithm MUST contain an AlgorithmIdentifier whose algorithm is id-alg-mtcProof, defined below, and whose parameters is omitted.

```
id-alg-mtcProof OBJECT IDENTIFIER ::= {
    iso(1) identified-organization(3) dod(6) internet(1) security(5)
    mechanisms(5) pkix(7) algorithms(6) TBD}
```

For initial experimentation, early implementations of this design will use the OID 1.3.6.1.4.1.44363.47.0 instead of id-alg-mtcProof.

The signatureValue contains an MTCProof structure, defined below using the TLS presentation language (Section 3 of [RFC8446]):

```
opaque HashValue[HASH_SIZE];

struct {
    TrustAnchorID cosigner_id;
    opaque signature<0..2^16-1>;
} MTCSignature;

struct {
    uint64 start;
    uint64 end;
    HashValue inclusion_proof<0..2^16-1>;
    MTCSignature signatures<0..2^16-1>;
} MTCProof;
```

start and end MUST contain the corresponding parameters of the chosen subtree. inclusion_proof MUST contain a subtree inclusion proof (Section 4.2) for the log entry and the subtree. signatures contains the chosen subtree signatures. In each signature, cosigner_id contains the cosigner ID (Section 5.4) in its binary representation (Section 3 of [I-D.ietf-tls-trust-anchor-ids]), and signature contains the signature value as described in Section 5.4.1.

The MTCProof is encoded into the signatureValue with no additional ASN.1 wrapping. The most significant bit of the first octet of the signature value SHALL become the first bit of the bit string, and so on through the least significant bit of the last octet of the signature value, which SHALL become the last bit of the bit string.

6.2. Full Certificates

A `_full certificate_` is a Merkle Tree certificate which contains sufficient signatures to allow a relying party to trust the choice of subtree, without any predistributed information beyond the cosigner(s) parameters. Full certificates can be issued without significant processing delay.

When issuing a certificate, the CA first adds the `TBSCertificateLogEntry` to its issuance log. It then schedules a job to construct a checkpoint and collect cosignatures. The job proceeds as follows:

1. The CA signs the checkpoint with its key(s) (Section 5.5).
2. Using the procedure in Section 4.4, the CA determines the two subtrees that cover the entries added between this checkpoint and the most recent checkpoint.
3. The CA signs each subtree with its key(s) (Section 5.4).
4. The CA requests sufficient checkpoint cosignatures (Section 5.4) from external cosigners to meet relying party requirements (Section 7.3).
5. The CA requests subtree cosignatures (Appendix B.3) from the cosigners above.
6. For each certificate in the interval, the CA constructs certificates (Section 6.1) using the covering subtree.

Steps 4 and 5 are analogous to requesting SCTs from CT logs in Certificate Transparency, except that a single run of this job collects signatures for many certificates at once. The CA MAY request signatures from a redundant set of cosigners and select the ones that complete first.

This document does not prescribe the specific cosigner roles, or a particular protocol for requesting cosignatures. Protocols for cosigners MAY vary depending on the needs for that cosigner. A consistency-only cosigner, such as [TLOG-WITNESS], might only require a checkpoint signature and consistency proof, while a mirroring cosigner, such as [TLOG-MIRROR] might require the full log contents.

A cosigner MAY expose a private interface for the CA, to reduce denial-of-service risk, or a cosigner MAY expose a public interface for other parties to request additional cosignatures. The latter may be useful if a relying party requires a cosigner that the CA does not

communicate with. In this case, an authenticating party MAY request cosignatures and add them to the certificate. However, it is RECOMMENDED that the CA collect cosignatures for the authenticating party. This simplifies deployment, as relying party policies change over time.

This document does not place any requirements on how frequently this job runs. More frequent runs results in lower issuance delay, but higher signing overhead. It is RECOMMENDED that CAs run at most one instance of this job at a time, starting the next instance after the previous one completes. A single run collects signatures for all entries since the most recent checkpoint, so there is little benefit to overlapping them. Less frequent runs may also aid relying parties that wish to directly audit signatures, as described in Section 5.2 of [AuditingRevisited], though this document does not define such a system.

6.3. Signatureless Certificates

A `_signatureless certificate_` is a Merkle Tree certificate which contains no signatures and instead assumes the relying party had predistributed information about which subtrees were trusted. Signatureless certificates are an optional size optimization. They require a processing delay to construct, and only work in a sufficiently up-to-date relying party. Authenticating parties thus SHOULD deploy a corresponding full certificate alongside any signatureless certificate, and use some application-protocol-specific mechanism to select between the two. Section 8 discusses such a mechanism for TLS [RFC8446].

6.3.1. Landmark Checkpoints

A signatureless certificate is constructed based on a `_landmark checkpoint sequence_`, which is a sequence of `_landmark checkpoints_`. Landmark checkpoints are agreed-upon tree sizes across the ecosystem for optimizing certificates. Landmark checkpoints SHOULD be allocated by the CA, but they can also be allocated by some other coordinating party. It is possible, but NOT RECOMMENDED, for multiple landmark checkpoint sequences to exist per CA. Landmark checkpoints are allocated to balance minimizing the delay in obtaining a signatureless certificate with minimizing the size of the relying party's predistributed state.

A landmark checkpoint sequence has the following fixed parameters:

- * `base_id`: An OID arc for trust anchor IDs of individual landmark checkpoints

- * `max_checkpoints`: A positive integer, describing the maximum number of landmark checkpoints that may contain unexpired certificates at any time
- * `checkpoint_url`: Some URL to fetch the current list of landmark checkpoints

Landmark checkpoints are numbered consecutively from zero. Each landmark checkpoint has a trust anchor ID, determined by appending the checkpoint number to `base_id`. For example, the trust anchor ID for landmark checkpoint 42 of a sequence with `base_id` of 32473.1 would be 32473.1.42.

Each landmark checkpoint specifies a tree size. The first landmark checkpoint, numbered zero, is always a tree size of zero. The sequence of tree sizes MUST be append-only and monotonically increasing.

Landmark checkpoints determine `_landmark subtrees_`: for each landmark checkpoint, other than number zero, let `tree_size` be the checkpoint's tree size and `prev_tree_size` be that of the previous landmark checkpoint. As described in Section 4.4, select the one or two subtrees that cover `[prev_tree_size, tree_size)`. Each of those subtrees is a landmark subtree.

The most recent `max_checkpoints` landmark checkpoints are said to be `_active_`. Landmark checkpoints MUST be allocated such that, at any given time, only active landmark checkpoints contain unexpired certificates. The active landmark subtrees are those determined by the active landmark checkpoints. There are at most $2 * \text{max_checkpoints}$ active landmark subtrees at any time. Every unexpired entry will be contained in one or more landmark subtree. Active landmark subtrees are predistributed to the relying party as trusted subtrees, as described in Section 7.4.

If landmark checkpoints are allocated incorrectly (e.g. past checkpoints change, or `max_checkpoints` is inaccurate), there are no security consequences, but some older certificates may fail to validate.

It is RECOMMENDED that checkpoints be allocated by picking some `time_between_checkpoints` interval, and then appending the latest checkpoint to the sequence, once per interval. If the latest checkpoint is already a landmark checkpoint, the interval is skipped. `max_checkpoints` can then be set to $\text{ceil}(\text{max_cert_lifetime} / \text{time_between_checkpoints})$, where `max_cert_lifetime` is the CA's maximum certificate lifetime. Allocations do not need to be precise, as long as `max_checkpoints` is accurate.

Relying parties will locally retain up to $2 * \text{max_checkpoints}$ hashes (Section 7.4) per CA, so max_checkpoints should be set to balance the delay between landmark checkpoints and the amount of state the relying party must maintain. Using the recommended procedure above, a CA with a maximum certificate lifetime of 7 days, allocating a landmark checkpoint every hour, will have a max_checkpoints of 168. The client state is then 336 hashes, or 10,752 bytes with SHA-256.

`checkpoint_url` MUST serve a resource with `Content-Type: text/plain; charset=utf-8` and the following lines. Each line MUST be terminated by a newline character (U+000A):

- * Two space-separated non-negative decimal integers:
`<last_checkpoint> <num_active_checkpoints>`. This line MUST satisfy the following, otherwise it is invalid:
 - `num_active_checkpoints <= max_checkpoints`
 - `num_active_checkpoints <= last_checkpoint`
- * `num_active_checkpoints + 1` lines each containing a single non-negative decimal integer, containing a tree size. Numbered from zero to `num_active_checkpoints`, line `i` contains the tree size for checkpoint `last_checkpoint - i`. The integers MUST be monotonically decreasing and lower or equal to the log's latest tree size.

6.3.2. Constructing Signatureless Certificates

Given a `TBSCertificateLogEntry` in the issuance log and a landmark checkpoint sequence, a signatureless certificate is constructed as follows:

1. Wait for the first landmark checkpoint to be allocated that contains the entry.
2. Determine the landmark checkpoint's landmark subtrees and select the one that contains the entry.
3. Construct a certificate (Section 6.1) using the selected subtree and no signatures.

Before sending this certificate, the authenticating party SHOULD obtain some application-protocol-specific signal that implies the relying party has been configured with the corresponding landmark checkpoint. (Section 7.4 defines how relying parties are configured.) The trust anchor ID of the landmark checkpoint may be used as an efficient identifier in the application protocol. Section 8 discusses how to do this in TLS [RFC8446].

6.4. Size Estimates

The inclusion proofs in full and signatureless certificates scale logarithmically with the size of the subtree. These sizes can be estimated with the CA's issuance rate. The byte counts below assume the issuance log's hash function is SHA-256.

Some organizations have published statistics which can be used to estimate this rate for the Web PKI. As of June 9th, 2025:

- * [LetsEncrypt] reported around 558,000,000 active certificates for a single CA
- * [MerkleTown] reported around 2,100,000,000 unexpired certificates in CT logs, across all CAs
- * [MerkleTown] reported an issuance rate of around 444,000 certificates per hour, across all CAs

The current issuance rate across the Web PKI may not necessarily be representative of the Web PKI after a transition to short-lived certificates. Assuming a certificate lifetime of 7 days, and that subscribers will update their certificates 75% of the way through their lifetime (see Section 10.4), every certificate will be reissued every 126 hours. This gives issuance rate estimates of around 4,400,000 certificates per hour and 17,000,000 certificates per hour, for the first two values above. Note the larger estimate is across all CAs, while subtrees would only span one CA.

Using the per-CA short lifetime estimate, if the CA mints a checkpoint every 2 seconds, full certificate subtrees will span around 2,500 certificates, leading to 12 hashes in the inclusion proof, or 384 bytes. Full certificates additionally must carry a sufficient set of signatures to meet relying party requirements.

If a new landmark checkpoint is allocated every hour, signatureless certificate subtrees will span around 4,400,000 certificates, leading to 23 hashes in the inclusion proof, giving an inclusion proof size of 736 bytes, with no signatures. This is significantly smaller than a single ML-DSA-44 signature, 2,420 bytes, and almost ten times smaller than the three ML-DSA-44 signatures necessary to include post-quantum SCTs.

The proof sizes grow logarithmically, so 32 hashes, or 1024 bytes, is sufficient for subtrees of up to 2^{32} (4,294,967,296) certificates.

7. Relying Parties

This section discusses how relying parties verify Merkle Tree Certificates.

7.1. Trust Anchors

In order to accept certificates from a Merkle Tree CA, a relying party MUST be configured with:

- * The log ID (Section 5.2)
- * A set of supported cosigners, as pairs of cosigner ID and public key
- * A policy on which combinations of cosigners to accept in a certificate (Section 7.3)
- * An optional list of trusted subtrees that are known to be consistent with the relying party's cosigner requirements (Section 7.4)
- * A list of revoked ranges of indices (Section 7.5)

[[TODO: Define some representation for this. In a trust anchor, there's a lot of room for flexibility in what the client stores. In principle, we could even encode some of this information in an X.509 intermediate certificate, if an application wishes to use this with a delegation model with intermediates, though the security story becomes more complex. Decide how/whether to do that.]]

7.2. Verifying Certificate Signatures

When verifying the signature on an X.509 certificate (Step (a)(1) of Section 6.1.3 of [RFC5280]) whose issuer is a Merkle Tree CA, the relying party performs the following procedure:

1. Check that the TBSCertificate's signature field is id-mtc-proof with omitted parameters. If either check fails, abort this process and fail verification.
2. Let index be the certificate's serial number. Check that $\text{start} \leq \text{index} < \text{end}$, and that $[\text{start}, \text{end})$ describes a subtree per Section 4.1. If either check fails, abort this process and fail verification.
3. If index is contained in one of the relying party's revoked ranges (Section 7.5), abort this process and fail verification.
4. Construct a TBSCertificateLogEntry as follows:
 1. Copy the version, issuer, validity, subject, issuerUniqueID, subjectUniqueID, and extensions fields from the TBSCertificate.
 2. Set subjectPublicKeyInfoHash to the hash of the DER encoding of subjectPublicKeyInfo.
5. Construct a MerkleTreeCertEntry of type tbs_cert_entry with contents the TBSCertificateLogEntry.
6. Evaluate inclusion_proof against the MerkleTreeCertEntry to compute the expected subtree hash
7. If the subtree is a trusted subtree (Section 7.4), check that the hash matches. Return success if it matches and failure if it does not.
8. Otherwise, check that signatures contain a sufficient set of valid signatures from cosigners to satisfy the relying party's cosigner requirements (Section 7.3). Unrecognized cosigners MUST be ignored. Signatures are verified as described in Section 5.4.1.

This procedure only replaces the signature verification portion of X.509 path validation. The relying party MUST continue to perform other checks, such as checking expiry.

7.3. Trusted Cosigners

A relying party's cosigner policy determines the sets of cosigners that must sign a view of the issuance log before it is trusted.

This document does not prescribe a particular policy, but gives general guidance. Relying parties MAY implement policies other than those described below, and MAY incorporate cosigners acting in roles not described in this document.

In picking trusted cosigners, relying party SHOULD ensure the following security properties:

Authenticity: The relying party only accepts entries certified by the CA

Transparency: The relying party only accepts entries that are publicly accessible, so that monitors, particularly the subject of the certificate, can notice any unauthorized certificates

Relying parties SHOULD ensure authenticity by requiring a signature from the most recent CA cosigner key. If the CA is transitioning from an old to new key, the relying party SHOULD accept both until certificates that predate the new key expire. This is analogous to the signature in a traditional X.509 certificate.

While a CA signature is sufficient to prove a subtree came from the CA, this is not enough to ensure the certificate is visible to monitors. A misbehaving CA might not operate the log correctly, either presenting inconsistent versions of the log to relying parties and monitors, or refuse to publish some entries.

To mitigate this, relying parties SHOULD ensure transparency by requiring a quorum of signatures from additional cosigners. At minimum, these cosigners SHOULD enforce a consistent view of the log. For example, [TLOG-WITNESS] describes a lightweight "witness" cosigner role that checks this with consistency proofs. This is not sufficient to ensure durable logging. Section 7.5 discusses mitigations for this. Alternatively, a relying party MAY require cosigners that serve a copy of the log, in addition to enforcing a consistent view. For example, [TLOG-MIRROR] describes a "mirror" cosigner role.

Relying parties MAY accept the same set of additional cosigners across issuance logs.

Cosigner roles are extensible without changes to certificate verification itself. Future specifications and individual deployments MAY define other cosigner roles to incorporate into relying party policies.

Section 10.2 discusses additional deployment considerations in cosigner selection.

7.4. Trusted Subtrees

As an optional optimization, a relying party MAY incorporate a periodically updated, predistributed list of active landmark subtrees, determined as described in Section 6.3.1. The relying party configures these as trusted subtrees, allowing it to accept signatureless certificates (Section 6.3) constructed against those subtrees.

Before configuring the subtrees as trusted, the relying party MUST obtain assurance that each subtree is consistent with checkpoints observed by a sufficient set of cosigners (see Section 5.4) to meet its cosigner requirements. It is not necessary that the cosigners have generated signatures over the specific subtrees, only that they are consistent.

This criteria can be checked given:

- * Some `_reference_checkpoint_` that contains the latest landmark checkpoint
- * For each cosigner, either:
 - A cosignature on the reference checkpoint
 - A cosigned checkpoint containing the referenced checkpoint and a valid Merkle consistency proof (Section 2.1.4 of [RFC9162]) between the two
- * For each subtree, a valid subtree consistency proof (Section 4.3) between the subtree and the reference checkpoint

[[TODO: The subtree consistency proofs have many nodes in common. It is possible to define a single "bulk consistency proof" that verifies all the hashes at once, but it's a lot more complex.]]

This document does not prescribe how relying parties obtain this information. A relying party MAY, for example, use an application-specific update service, such as the services described in [CHROMIUM] and [FIREFOX]. If the relying party considers the service sufficiently trusted (e.g. if the service provides the trust anchor list or certificate validation software), it MAY trust the update service to perform these checks.

The relying party SHOULD incorporate its trusted subtree configuration in application-protocol-specific certificate selection mechanisms, to allow an authenticating party to select a signatureless certificate. The trust anchor IDs of the landmark checkpoints may be used as efficient identifiers in the application protocol. Section 8 discusses how to do this in TLS [RFC8446].

7.5. Revocation by Index

For each supported Merkle Tree CA, the relying party maintains a list of revoked ranges of indices. This allows a relying party to efficiently revoke entries of an issuance log, even if the contents are not necessarily known. This may be used to mitigate the security consequences of misbehavior by a CA, or other parties in the ecosystem.

When a relying party is first configured to trust a CA, it SHOULD be configured to revoke all entries from zero up to but not including the first available unexpired certificate at the time. This revocation SHOULD be periodically updated as entries expire and logs are pruned (Section 5.6.1). In particular, when CAs prune entries, relying parties SHOULD be updated to revoke all newly unavailable entries. This gives assurance that, even if some unavailable entry had not yet expired, the relying party will not trust it. It also allows monitors to start monitoring a log without processing expired entries.

A misbehaving CA might correctly construct a globally consistent log, but refuse to make some entries or intermediate nodes available. Consistency proofs between checkpoints and subtrees would pass, but monitors cannot observe the entries themselves. Relying parties whose cosigner policies (Section 7.3) do not require durable logging (e.g. via [TLOG-MIRROR]) are particularly vulnerable to this. In this case, the indices of the missing entries will still be known, so relying parties can use this mechanism to revoke the unknown entries, possibly as an initial, targeted mitigation before a complete CA removal.

When a CA is found to be untrustworthy, relying parties SHOULD remove trust in that CA. To minimize the compatibility impact of this mitigation, index-based revocation can be used to only distrust entries after some index, while leaving existing entries accepted. This is analogous to the [SCTNotAfter] mechanism used in some PKIs.

8. Use in TLS

TLS implementations that authenticate with or accept Merkle Tree certificates SHOULD support trust anchor IDs ([I-D.ietf-tls-trust-anchor-ids]) for certificate selection.

A full certificate has a trust anchor ID of the corresponding log ID (Section 5.2). The authenticating party can obtain this information either by parsing the certificate's issuer field or via out-of-band information as described in Section 3.2 of [I-D.ietf-tls-trust-anchor-ids].

[[TODO: Ideally we would negotiate cosigners. <https://github.com/tlswg/tls-trust-anchor-ids/issues/54> has a sketch of how one might do this, though other designs are possible. Negotiating cosigners allows the ecosystem to manage cosigners efficiently, without needing to collect every possible cosignature and send them all at once. This is wasteful, particularly with post-quantum algorithms.]]

A full certificate MAY be used without signals about what the relying party trusts. As with other choices of default certificates, an authenticating party that does so assumes that the relying party trusts the issuing CA, e.g. because the CA is relatively ubiquitous among the relying parties that it supports.

A signatureless certificate has a trust anchor ID of the corresponding landmark checkpoint, as described in Section 6.3.1. This can be configured in the authenticating party via out-of-band information, as described in Section 3.2 of [I-D.ietf-tls-trust-anchor-ids]. A relying party that has been configured with trusted subtrees (Section 7.4) derived from a set of landmark checkpoints SHOULD be configured to support those checkpoints' trust anchor IDs. TLS certificate selection will then correctly determine whether the signatureless certificate is compatible with the relying party.

[[TODO: We can do slightly better. If the relying party supports checkpoint 42, it can be assumed to also support checkpoint 41, 40, 39, etc. <https://github.com/tlswg/tls-trust-anchor-ids/issues/62> discusses how to fit this into the trust anchor IDs framework. This allows the client to summarize its state with one ID per CA.]]

Authenticating parties SHOULD preferentially use signatureless certificates over full certificates, when both are supported by the relying party. A signatureless certificate asserts the same information as its full counterpart, but is expected to be smaller. A signatureless certificate SHOULD NOT be used without a signal that

the relying party trusts the corresponding landmark subtree. Even if the relying party is assumed to trust the issuing CA, the relying party may not have sufficiently up-to-date trusted subtrees predistributed.

9. ACME Extensions

This section describes how to issue Merkle Tree certificates using ACME [RFC8555].

When downloading the certificate (Section 7.4.2 of [RFC8555]), ACME clients supporting Merkle Tree certificates SHOULD send "application/pem-certificate-chain-with-properties" in their Accept header (Section 12.5.1 of [RFC9110]). ACME servers issuing Merkle Tree certificates SHOULD then respond with that content type and include trust anchor ID information as described in Section 6 of [I-D.ietf-tls-trust-anchor-ids]. Section 8 describes the trust anchor ID assignments for full and signatureless certificates.

When processing an order for a Merkle Tree certificate, the ACME server moves the order to the "valid" state once the corresponding entry is sequenced in the issuance log. The order's certificate URL then serves the full certificate, constructed as described in Section 6.2.

The full certificate response SHOULD additionally carry a alternate URL for the signatureless certificate, as described Section 7.4.2 of [RFC8555]. Before the signatureless certificate is available, the alternate URL SHOULD return a HTTP 503 (Service Unavailable) response, with a Retry-After header (Section 10.2.3 of [RFC9110]) estimating when the certificate will become available. Once the next landmark checkpoint is allocated, the ACME server constructs a signatureless certificate, as described in Section 6.3 and serves it from the alternate URL.

ACME clients supporting Merkle Tree certificates SHOULD support fetching alternate chains. If an alternate chain returns an HTTP 503 with a Retry-After header, as described above, the client SHOULD retry the request at the specified time.

10. Deployment Considerations

10.1. Operational Costs

10.1.1.1. Certification Authority Costs

While Merkle Tree certificates expects CAs to operate logs, the costs of these logs are expected to be much lower than a CT log from [RFC6962] or [RFC9162]:

Section 5.6 does not constrain the API to the one defined in [RFC6962] or [RFC9162]. If the PKI uses a tile-based protocol, such as [TLOG-TILES], the issuance log benefits from the improved caching properties of such designs.

Unlike a CT log, an issuance log does not have public submission APIs. Log entries are only added by the CA directly. The costs are thus expected to scale with the CA's own operations.

A CA only needs to produce a digital signature for every checkpoint, rather than for every certificate. The lower signature rate requirements could allow more secure and/or economical key storage choices.

Individual entries are kept small and do not scale with public key or signature sizes. This mitigates growth from post-quantum algorithms. Public keys in entries are replaced with fixed-sized hashes. There are no signatures in entries themselves, and only signatures on the very latest checkpoint are retained. Every new checkpoint completely subsumes the old checkpoint, so there is no need to retain older signatures. Likewise, a subtree is only signed if contained in another signed checkpoint.

Log pruning (Section 5.6.1) allows a long-lived log to serve only the more recent entries, scaling with the size of the retention window, rather than the log's total lifetime.

Mirrors of the log can also reduce CA bandwidth costs, because monitors can fetch data from mirrors instead of CAs directly. In PKIs that deploy mirrors as part of cosigner policies, relying parties could set few availability requirements on CAs, as described in Section 10.3.

10.1.1.2. Cosigner Costs

The costs of cosigners vary by cosigner role. A consistency-checking cosigner, such as [TLOG-WITNESS], requires very little state and can be run with low cost.

A mirroring cosigner, such as [TLOG-MIRROR], performs comparable roles as CT logs, but several of the cost-saving properties in Section 10.1.1 also apply: improved protocols, smaller entries, less

frequent signatures, and log pruning. While a mirror does need to accommodate another party's (the CA's) growth rate, it grows only from new issuances from that one CA. If one CA's issuance rate exceeds the mirror's capacity, that does not impact the mirror's copies of other CAs. Mirrors also do not need to defend against a client uploading a large number of existing certificates all at once. Submissions are also naturally batched and serialized.

10.1.3. Monitor Costs

In a CT-based PKI, every log carries a potentially distinct subset of active certificates, so monitors must check the contents of every CT log. At the same time, certificates are commonly synchronized between CT logs. As a result, a monitor will typically download each certificate multiple times, once for every log. In Merkle Tree Certificates, each entry appears in exactly one log. A relying party might require a log to be covered by a quorum of mirrors, but each mirror is cryptographically verified to serve the same contents. Once a monitor has obtained some entry from one mirror, it does not need to download it from the others.

In addition to downloading each entry only once, the entries themselves are smaller, as discussed in Section 10.1.1.

10.2. Choosing Cosigners

In selecting trusted cosigners and cosigner requirements (Section 7.3), relying parties navigate a number of trade-offs:

A consistency-checking cosigner, such as [TLOG-WITNESS], is very cheap to run, but does not guarantee durable logging, while a mirroring cosigner is more expensive and may take longer to cosign structures. Requiring a mirror signature provides stronger guarantees to the relying party, which in turn can reduce the requirements on CAs (see Section 10.3), however it may cause certificate issuance to take longer. That said, mirrors are comparable to CT logs, if not cheaper (see Section 10.1), so they may be appropriate in PKIs where running CT logs is already viable.

Relying parties that require larger quorums of trusted cosigners can reduce the trust placed in any individual cosigner. However, these larger quorums result in larger, more expensive full certificates. The cost of this will depend on how frequently the signatureless optimization occurs in a given PKI. Conversely, relying parties that require smaller quorums have smaller full certificates, but place more trust in their cosigners.

Relying party policies also impact monitor operation. If a relying party accepts any one of three cosigners, monitors SHOULD check the checkpoints of all three. Otherwise, a malicious CA may send different split views to different cosigners. More generally, monitors SHOULD check the checkpoints in the union of all cosigners trusted by all supported relying parties. This is an efficient check because, if the CA is operating correctly, all cosigners will observe the same tree. Thus the monitor only needs to check consistency proofs between the checkpoints, and check the log contents themselves once. Monitors MAY also rely on other parties in the transparency ecosystem to perform this check.

10.3. Log Availability

CAs and mirrors are expected to serve their log contents over HTTP. It is possible for the contents to be unavailable, either due to temporary service outage or because the log has been pruned (Section 5.6.1). If some resources are unavailable, they may not be visible to monitors.

As in CT, PKIs which deploy Merkle Tree certificates SHOULD establish availability policies, adhered to by trusted CAs and mirrors, and enforced by relying party vendors as a condition of trust. Exact availability policies for these services are out of scope for this document, but this section provides some general guidance.

Availability policies SHOULD specify how long an entry must be made available, before a CA or mirror is permitted to prune the entry. It is RECOMMENDED to define this using a `_retention period_`, which is some time after the entry has expired. In such a policy, an entry could only be pruned if it, and all preceding entries, have already expired for the retention period. Policies MAY opt to set different retention periods between CAs and mirrors. Permitting limited log retention is analogous to the CT practice of temporal sharding [CHROME-CT], except that a pruned issuance log remains compatible with older, unupdated relying parties.

Such policies impact monitors. If the retention period is, e.g. 6 months, this means that monitors are expected to check entries of interest within 6 months. It also means that a new monitor may only be aware of a 6 month history of entries issued for a particular domain.

If historical data is not available to verify the retention period, such as information in another mirror or a trusted summary of expiration dates of entries, it may not be possible to confirm correct behavior. This is mitigated by the revocation process described in Section 7.5: if a CA were to prune a forward-dated entry

and, in the 6 months when the entry was available, no monitor noticed the unusual expiry, an updated relying party would not accept it anyway.

The log pruning process simply makes some resources unavailable, so availability policies SHOULD constrain log pruning in the same way as general resource availability. That is, if it would be a policy violation for the log to fail to serve a resource, it should also be a policy violation for the log to prune such that the resource is removed, and vice versa.

PKIs that require mirror cosignatures (Section 7.3) can impose minimal to no availability requirements on CAs, all without compromising transparency goals. If a CA never makes some entry available, mirrors will be unable to update. This will prevent relying parties from accepting the undisclosed entries. However, a CA which is persistently unavailable may not offer sufficient benefit to be used by authenticating parties or trusted by relying parties.

However, if a mirror's interface becomes unavailable, monitors may be unable to check for unauthorized issuance, if the entries are not available in another mirror. This does compromise transparency goals. As such, availability policies SHOULD set availability expectations on mirrors. This can also be mitigated by using multiple mirrors, either directly enforced in cosigner requirements, or by keeping mirrors up-to-date with each other.

In PKIs that do not require mirroring cosigners, the CA's serving endpoint is more crucial for monitors. Such PKIs thus SHOULD set availability requirements on CAs.

In each of these cases, availability failures can be mitigated by revoking the unavailable entries by index, as described in Section 7.5, likely as a first step in a broader distrust.

10.4. Certificate Renewal

When an authenticating party requests a certificate, the signatureless certificate will not be available until the next landmark checkpoint is ready. From there, the signatureless certificate will not be available until relying parties receive new trusted subtrees.

To maximize coverage of the signatureless certificate optimization, authenticating parties performing routine renewal SHOULD request a new Merkle Tree certificate some time before the previous Merkle Tree certificate expires. Renewing around 75% into the previous certificate's lifetime is RECOMMENDED. Authenticating parties

additionally SHOULD retain both the new and old certificates in the certificate set until the old certificate expires. As the new subtrees are delivered to relying parties, certificate negotiation will transition relying parties to the new certificate, while retaining the old certificate for relying parties that are not yet updated.

The above also applies if the authenticating party is performing a routine key rotation alongside the routine renewal. In this case, certificate negotiation would pick the key as part of the certificate selection. This slightly increases the lifetime of the old key but maintains the size optimization continuously.

If the service is rotating keys in response to a key compromise, this option is not appropriate. Instead, the service SHOULD immediately discard the old key and request a full certificate and the revocation of the previous certificate. This will interrupt the size optimization until the new signatureless certificate is available and relying parties are updated.

10.5. Multiple CA Keys

The separation between issuance logs and CA cosigners gives CAs additional flexibility in managing keys. A CA operator wishing to rotate keys, e.g. to guard against compromise of older key material, or upgrade to newer algorithms, could retain the same issuance log and sign its checkpoints and subtrees with both keys in parallel, until relying parties are all updated. Older relying parties would verify the older signatures, while newer relying parties would verify the newer signatures. A cosignature negotiation mechanism in the application protocol (see Section 8) would avoid using extra bandwidth for the two signatures.

11. Privacy Considerations

The Privacy Considerations described in Section 9 of [I-D.ietf-tls-trust-anchor-ids] apply to its use with Merkle Tree Certificates.

In particular, relying parties that share an update process for trusted subtrees (Section 7.4) will fetch the same stream of updates. However, updates may reach different users at different times, resulting in some variation across users. This variation may contribute to a fingerprinting attack [RFC6973]. If the Merkle Tree CA trust anchors are sent unconditionally in `trust_anchors`, this variation will be passively observable. If they are sent conditionally, e.g. with the DNS mechanism, the trust anchor list will require active probing.

12. Security Considerations

12.1. Authenticity

A key security requirement of any PKI scheme is that relying parties only accept assertions that were certified by a trusted certification authority. Merkle Tree certificates achieve this by ensuring the relying party only accepts authentic subtree hashes:

- * In full certificates, the relying party's cosigner requirements (Section 7.3) are expected to include some signature by the CA's cosigner. The CA's cosigner (Section 5.5) is defined to certify the contents of every checkpoint and subtree that it signs.
- * In signatureless certificates, the cosigner requirements are checked ahead of time, when the trusted subtrees are predistributed (Section 7.4).

Given such a subtree hash, computed over entries that the CA certified, it then must be computationally infeasible to construct an entry not on this list, and some inclusion proof, such that inclusion proof verification succeeds. This requires using a collision-resistant hash in the Merkle Tree construction.

Log entries contain public key hashes, so it must additionally be computationally infeasible to compute a public key whose hash matches the entry, other than the intended public key. This also requires a collision-resistant hash.

12.2. Transparency

The transparency mechanisms in this document do not prevent a CA from issuing an unauthorized certificate. Rather, they provide comparable security properties as Certificate Transparency [RFC9162] in ensuring that all certificates are either rejected by relying parties, or visible to monitors and, in particular, the subject of the certificate.

Compared to Certificate Transparency, some of the responsibilities of a log have moved to the CA. All signatures generated by the CA in this system are assertions about some view of the CA's issuance log. However, a CA does not need to function correctly to ensure transparency properties. Relying parties are expected to require a quorum of additional cosigners, which together enforce properties of the log (Section 7.3) and prevent or detect CA misbehavior:

A CA might violate the append-only property of its log and present different views to different parties. However, each individual cosigner will only follow a single append-only view of the log history. Provided the cosigners are correctly operated, relying parties and monitors will observe consistent views between each other. Views that were not cosigned at all may not be detected, but they also will not be accepted by relying parties.

If the CA sends one view to some cosigners and another view to other cosigners, it is possible that multiple views will be accepted by relying parties. However, in that case monitors will observe that cosigners do not match each other. Relying parties can then react by revoking the inconsistent indices (Section 7.5), and likely removing the CA. If the cosigners are mirrors, the underlying entries in both views will also be visible.

A CA might correctly construct its log, but refuse to serve some unauthorized entry, e.g. by feigning an outage or pruning the log outside the retention policy (Section 10.3). If the relying party requires cosignatures from trusted mirrors, the entry will either be visible to monitors in the mirrors, or have never reached a mirror. In the latter case, the entry will not have been cosigned, so the relying party would not accept it. If the relying party accepts log views without a trusted mirror, the unauthorized entry may not be available. However, the existence of `_some_` entry at that index will be visible, so monitors will know the CA is failing to present an entry. Relying parties can then react by revoking the undisclosed entries by index (Section 7.5), and likely removing the CA.

12.3. Public Key Hashes

Unlike Certificate Transparency, the mechanisms in this document do not provide the subject public keys, only the hashed values. This is intended to reduce log serving costs, particularly with large post-quantum keys. As a result, monitors look for unrecognized hashes instead of unrecognized keys. Any unrecognized hash, even if the preimage is unknown, indicates an unauthorized certificate.

This optimization complicates studies of weak public keys, e.g. [SharedFactors]. Such studies will have to retrieve the public keys separately, such as by connecting to the TLS servers, or fetching from the CA if it retains the unhashed key. This document does not define a mechanism for doing this, or require that CAs or mirrors retain unhashed keys. The transparency mechanisms in this protocol are primarily intended to allow monitors to observe certificate issuance.

12.4. Non-Repudiation

When a monitor finds an unauthorized certificate issuance in a log or mirror, it must be possible to prove the CA indeed certified the information in the entry. However, only the latest checkpoint signature is retained by the transparency ecosystem, so it may not be possible to reconstruct the exact certificate seen by relying parties.

However, per Section 5.5, any checkpoint signature is a binding assertion by the CA that it has certified every entry in the checkpoint. Thus, given `_any_` signed checkpoint that contains the unauthorized entry, a Merkle inclusion proof (Section 2.1.3 of [RFC9162]) is sufficient to prove the CA issued the entry. This is analogous to how, in Section 3.2.1 of [RFC9162], CAs are held accountable for signed CT precertificates.

The transparency ecosystem does not retain unhashed public keys, so it also may not be possible to construct a complete certificate from the checkpoint signature and inclusion proof. However, if the log entry's `subjectPublicKeyInfoHash` does not correspond to an authorized key for the subject of the certificate, the entry is still unauthorized. A Merkle Tree CA is held responsible for all log entries it certifies, whether or not the preimage of the hash is known.

12.5. New Log Entry Types

`MerkleTreeCertEntry` (Section 5.3) is extensible and permits protocol extensions to define new formats for the CA to certify. This means older CAs, cosigners, relying parties, and monitors might interact with new entries:

Section 5.3 and Section 5.5 forbid a CA from logging or signing entries that it does not recognize. A CA cannot faithfully claim to certify information if it does not understand it. This is analogous to how a correctly-operated X.509 can never sign an unrecognized X.509 extension.

External cosigners may or may not interact with the unrecognized entries. [TLOG-MIRROR] and [TLOG-WITNESS] describe cosigners whose roles do not interpret the contents of log entries. New entry types MAY be added without updating them. If a cosigner role does interpret a log entry, it MUST define how it interacts with unknown ones.

If a relying party trusts an issuance log, but the issuance log contains an unrecognized entry, the entry will not cause it to accept an unexpected certificate. In Section 7.2, the relying party constructs the `MerkleTreeCertEntry` that it expects. The unrecognized entry will have a different type value, so the proof will never succeed, assuming the underlying hash function remains collision-resistant.

If a monitor observes an entry with unknown type, it may not be able to determine if it is of interest. For example, it may be unable to tell whether it covers some relevant DNS name. Until the monitor is updated to reflect the current state of the PKI, the monitor may be unable to detect all misissued certificates.

This situation is analogous to the addition of a new X.509 extension. When relying parties add support for log entry types or new X.509 extensions, they SHOULD coordinate with monitors to ensure the transparency ecosystem is able to monitor the new formats.

12.6. Certificate Malleability

An ASN.1 structure like X.509's `Certificate` is an abstract data type that is independent of its serialization. There are multiple encoding rules for ASN.1. Commonly, protocols use DER [X.690], such as Section 4.4.2 of [RFC8446]. This aligns with Section 4.1.1.3 of [RFC5280], which says X.509 signatures are computed over the DER-encoded `TBSCertificate`. After signature verification, applications can assume the DER-encoded `TBSCertificate` is not malleable.

While the signature verification process in Section 7.2 first transforms the `TBSCertificate` into a `TBSCertificateLogEntry`, it preserves this non-malleability. There is a unique valid DER encoding for every abstract `TBSCertificate` structure, so malleability of the DER-encoded `TBSCertificate` reduces to malleability of the `TBSCertificate` value:

- * The version, issuer, validity, subject, issuerUniqueID, subjectUniqueID, and extensions fields are copied from the `TBSCertificate` to the `TBSCertificateLogEntry` unmodified, so they are directly authenticated by the inclusion proof.
- * `serialNumber` is omitted from `TBSCertificateLogEntry`, but its value determines the inclusion proof index, which authenticates it.
- * The redundant signature field in `TBSCertificate` is omitted from `TBSCertificateLogEntry`, but Section 7.2 checks for an exact value, so no other values are possible.

- * `subjectPublicKeyInfo` is hashed as `subjectPublicKeyInfoHash` in `TBSCertificateLogEntry`. Provided the underlying hash function is collision-resistant, no other values are possible for a given log entry.

Some non-conforming X.509 implementations use a BER [X.690] parser instead of DER and then, instead of computing the canonical DER encoding per Section 4.1.1.3 of [RFC5280], verify the signature over the particular BER encoding received. Those implementations may still assume the received BER-encoded `TBSCertificate` is not malleable. However, as BER encoding is not unique, this breaks the above assumptions in Merkle Tree Certificates. To preserve non-malleability, these non-conforming implementations MUST do the following when verifying a Merkle Tree Certificate:

- * When parsing the outermost `TBSCertificate` SEQUENCE tag and length, reparse with a conforming DER parser.
- * When copying the version, issuer, validity, subject, `issuerUniqueID`, `subjectUniqueID`, and extensions fields, either reparse with a conforming DER parser, or copy over the observed BER encodings.
- * Reparse the `serialNumber` field with a conforming DER parser.
- * Reparse the signature field with a conforming DER parser. Equivalently, check for an exact match for the expected, DER-encoded value.
- * When hashing `subjectPublicKeyInfo`, either reparse with a conforming DER parser, or hash the observed BER encoding.

Section 5.3 requires that the `TBSCertificateLogEntry` in a `MerkleTreeCertEntry` be DER-encoded, so applying a stricter parser will be compatible with conforming CAs. While these existing non-conforming implementations may be unable to switch to a DER parser due to compatibility concerns, Merkle Tree Certificates is new, so there is no existing deployment of malformed BER-encoded `TBSCertificateLogEntry` structures.

The above only ensures the TBSCertificate portion is non-malleable. In Merkle Tree Certificates, similar to ECDSA X.509 signature, the signature value is malleable. Multiple MTCProof structures may prove a single TBSCertificate structure. Additionally, in all X.509-based protocols, a BER-based parser for the outer, unsigned Certificate structure will admit malleability in those portions of the encoding. Applications that derive a unique identifier from the Certificate MUST instead use the TBSCertificate, or some portion of it, for Merkle Tree Certificates.

13. IANA Considerations

TODO: Fill this in.

14. References

14.1. Normative References

[FIPS186-5]

"Digital Signature Standard (DSS)", National Institute of Standards and Technology (U.S.),
DOI 10.6028/nist.fips.186-5, February 2023,
<<https://doi.org/10.6028/nist.fips.186-5>>.

[FIPS204]

"Module-lattice-based digital signature standard",
National Institute of Standards and Technology (U.S.),
DOI 10.6028/nist.fips.204, August 2024,
<<https://doi.org/10.6028/nist.fips.204>>.

[I-D.draft-ietf-lamps-x509-alg-none]

Benjamin, D., "Unsigned X.509 Certificates", Work in Progress, Internet-Draft, draft-ietf-lamps-x509-alg-none-08, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-lamps-x509-alg-none-08>>.

[I-D.ietf-tls-trust-anchor-ids]

Beck, B., Benjamin, D., O'Brien, D., and K. Nekritz, "TLS Trust Anchor Identifiers", Work in Progress, Internet-Draft, draft-ietf-tls-trust-anchor-ids-01, 14 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-tls-trust-anchor-ids-01>>.

[RFC2119]

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

- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<https://www.rfc-editor.org/rfc/rfc3629>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/rfc/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, <<https://www.rfc-editor.org/rfc/rfc5280>>.
- [RFC8032] Josefsson, S. and I. Liusvaara, "Edwards-Curve Digital Signature Algorithm (EdDSA)", RFC 8032, DOI 10.17487/RFC8032, January 2017, <<https://www.rfc-editor.org/rfc/rfc8032>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/rfc/rfc8446>>.
- [RFC8555] Barnes, R., Hoffman-Andrews, J., McCarney, D., and J. Kasten, "Automatic Certificate Management Environment (ACME)", RFC 8555, DOI 10.17487/RFC8555, March 2019, <<https://www.rfc-editor.org/rfc/rfc8555>>.
- [RFC9110] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "HTTP Semantics", STD 97, RFC 9110, DOI 10.17487/RFC9110, June 2022, <<https://www.rfc-editor.org/rfc/rfc9110>>.
- [RFC9162] Laurie, B., Messeri, E., and R. Stradling, "Certificate Transparency Version 2.0", RFC 9162, DOI 10.17487/RFC9162, December 2021, <<https://www.rfc-editor.org/rfc/rfc9162>>.
- [RFC9612] Mirsky, G., Tantsura, J., Varlashkin, I., and M. Chen, "Bidirectional Forwarding Detection (BFD) Reverse Path for MPLS Label Switched Paths (LSPs)", RFC 9612, DOI 10.17487/RFC9612, July 2024, <<https://www.rfc-editor.org/rfc/rfc9612>>.

- [SHS] "Secure hash standard", National Institute of Standards and Technology (U.S.), DOI 10.6028/nist.fips.180-4, 2015, <<https://doi.org/10.6028/nist.fips.180-4>>.
- [X.690] ITU-T, "Information technology - ASN.1 encoding Rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", ISO/IEC 8824-1:2021 , February 2021.

14.2. Informative References

- [APPLE-CT] Apple, "Apple's Certificate Transparency policy", 5 March 2021, <<https://support.apple.com/en-us/HT205280>>.
- [AuditingRevisited] Heimberger, L., Patton, C., and B. Westerbaan, "Private SCT Auditing, Revisited", 25 April 2025, <<https://eprint.iacr.org/2025/556.pdf>>.
- [CABF-153] CA/Browser Forum, "Ballot 153 Short-Lived Certificates", 11 November 2015, <<https://cabforum.org/2015/11/11/ballot-153-short-lived-certificates/>>.
- [CABF-SC081] CA/Browser Forum, "Ballot SC081v3: Introduce Schedule of Reducing Validity and Data Reuse Periods", 11 April 2025, <<https://cabforum.org/2025/04/11/ballot-sc081v3-introduce-schedule-of-reducing-validity-and-data-reuse-periods/>>.
- [CHROME-CT] Google Chrome, "Chrome Certificate Transparency Policy", 17 March 2022, <https://googlechrome.github.io/CertificateTransparency/ct_policy.html>.
- [CHROMIUM] Chromium, "Component Updater", 3 March 2022, <https://chromium.googlesource.com/chromium/src/+/main/components/component_updater/README.md>.
- [FIREFOX] Mozilla, "Firefox Remote Settings", 20 August 2022, <<https://wiki.mozilla.org/Firefox/RemoteSettings>>.
- [LetsEncrypt] Let's Encrypt, "Let's Encrypt Stats", 7 March 2023, <<https://letsencrypt.org/stats/>>.

[MerkleTown]

Cloudflare, Inc., "Merkle Town", 7 March 2023,
<<https://ct.cloudflare.com/>>.

[RFC4514] Zeilenga, K., Ed., "Lightweight Directory Access Protocol (LDAP): String Representation of Distinguished Names", RFC 4514, DOI 10.17487/RFC4514, June 2006, <<https://www.rfc-editor.org/rfc/rfc4514>>.

[RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", RFC 6962, DOI 10.17487/RFC6962, June 2013, <<https://www.rfc-editor.org/rfc/rfc6962>>.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/rfc/rfc6973>>.

[SCTNotAfter]

Adrian, D., "How to distrust a CA without any certificate errors", March 2025, <<https://dadrian.io/blog/posts/sct-not-after/>>.

[SharedFactors]

Vage, H. F. and University of Bergen, "Finding shared RSA factors in the Certificate Transparency logs", 13 May 2022, <https://bora.uib.no/bora-xmlui/bitstream/handle/11250/3001128/Masters_thesis_for_University_of_Bergen.pdf>.

[SIGNED-NOTE]

C2SP, "Note", April 2025, <<https://c2sp.org/signed-note>>.

[STH-Discipline]

Barnes, R., "STH Discipline & Security Considerations", 3 March 2017, <<https://mailarchive.ietf.org/arch/msg/trans/Zm4NqyRc7LDsOtV56EchBIT9r4c/>>.

[TLOG-CHECKPOINT]

C2SP, "Transparency Log Checkpoints", March 2024, <<https://c2sp.org/tlog-checkpoint>>.

[TLOG-MIRROR]

C2SP, "Transparency Log Mirrors", July 2025, <<https://c2sp.org/tlog-mirror>>.

`[TLOG-TILES]`

C2SP, "Tiled Transparency Logs", June 2025,
<<https://c2sp.org/tlog-tiles>>.

`[TLOG-WITNESS]`

C2SP, "Transparency Log Witness Protocol", June 2025,
<<https://c2sp.org/tlog-witness>>.

Appendix A. ASN.1 Module

TODO: Add an ASN.1 module. Leaving a placeholder as a reminder.

Appendix B. Extensions to Tiled Transparency Logs (To Be Removed)

[[TODO: This section is expected to be removed. It is sketched here purely for illustrative purposes, until the features are defined somewhere else, e.g. in the upstream tlog documents.]]

B.1. Pruning a Tiled Transparency Log

<https://github.com/C2SP/C2SP/pull/138> describes how to support pruning in tlogs but, as of writing, is not part of [TLOG-TILES] yet.

[TLOG-MIRROR] is also updated to add a minimum index parameter to each mirror and apply the above changes to its monitoring interface. The first step of the recommended update procedure in [TLOG-MIRROR] is updated to say, "Let `mirror_size` be the tree size of the mirror checkpoint, or the mirror's minimum index if the mirror has not copied the log yet."

A mirror MAY use the origin log's current minimum index to initialize its starting minimum index, but it is not required to reflect subsequent pruning. A mirror MAY initialize its minimum index to some higher value than the origin, if it does not intend to log some entries. It MAY initialize its minimum index to a lower value, if the pruned resources are available from some other source, e.g., another mirror.

B.2. Subtree Signed Note Format

A subtree, with signatures, can be represented as a signed note [SIGNED-NOTE]. Trust anchor IDs can be converted into log origins and cosigner names by concatenating the ASCII string `oid/1.3.6.1.4.1.` and the ASCII representation of the trust anchor ID. For example, the checkpoint origin for a log named `32473.1` would be `oid/1.3.6.1.4.1.32473.1.`

The note body is a sequence of the following lines, each terminated by a newline character (U+000A):

- * The log origin
- * Two space-separated, non-negative decimal integers, <start> <end>
- * The subtree hash, as single hash encoded in base64

Each note signature has a key name of the cosigner name. The signature's key ID is computed using the reserved signature type in [SIGNED-NOTE], and a fixed string, as follows:

```
key ID = SHA-256(key name || 0x0A || 0xFF || "mtc-subtree/v1")[:4]
```

A subtree whose start is zero can also be represented as a checkpoint [TLOG-CHECKPOINT]. A corresponding subtree signature can be represented as a note signature using a key ID computed as follows:

```
key ID = SHA-256(key name || 0x0A || 0xFF || "mtc-checkpoint/v1")[:4]
```

The only difference between the two forms is the implicit transformation from the signed note text to the MTCSubtree structure.

B.3. Requesting Subtree Signatures

This section defines the sign-subtree cosigner HTTP endpoint for clients to obtain subtree signatures from non-CA cosigners, such as mirrors and witnesses. It may be used by the CA when assembling a certificate, or by an authenticating party to add a cosignature to a certificate that the CA did not themselves obtain.

The cosigner MAY expose this endpoint publicly to general authenticating parties, or privately to the CA. The latter is sufficient if the CA is known to automatically request cosignatures from this cosigner when constructing certificates. If private, authenticating the CA is out of scope for this document.

Clients call this endpoint as POST <prefix>/sign-subtree, where prefix is some URL prefix. For a mirror or witness, the URL prefix is the submission prefix. The client's request body MUST be a sequence of:

- * The requested subtree as a signed note (Appendix B.2), with zero or more signatures. The endpoint MAY require signatures from the CA as a DoS mitigation, as described below.
- * A blank line

- * A checkpoint, signed by the requested cosigner. The checkpoint's tree size must be at least end.
- * A blank line
- * Zero or more subtree consistency proof (Section 4.3) lines. Each line MUST encode a single hash in base64 [RFC4648]. The client MUST NOT send more than 63 consistency proof lines.

Each line MUST terminate in a newline character (U+000A).

The cosigner performs the following steps:

1. Check that the checkpoint contains signatures from itself
2. Check that the subtree consistency proof proves consistency between the subtree hash and the checkpoint
3. If all checks pass, cosign the subtree, as described in Section 5.4

On success, the response body MUST be a sequence of one or more note signature lines [SIGNED-NOTE], each starting with an em dash character (U+2014) and ending with a newline character (U+000A). The signatures MUST be cosignatures from the cosigner key(s) on the subtree.

Instead of statelessly validating checkpoints by signature, the cosigner MAY statefully check the requested checkpoint against internal witness or mirror state. In this case, if the cosigner needs a newer checkpoint, it responds with a "409 Conflict" with its latest signed checkpoint. In this case, the subtree cosigning SHOULD remember and accept the last few signed checkpoints, to minimize conflicts.

If operating statefully, the subtree cosigner process only needs read access to the mirror or witness state and can freely operate on stale state without violating any invariants.

Mirrors MAY choose to check subtree hashes by querying their log state, instead of evaluating proofs.

Publicly-exposed subtree cosigning endpoints MAY mitigate DoS in a variety of techniques:

- * Only cosigning recent subtrees, as old subtrees do not need to be co-signed

- * Caching subtree signatures
- * Requiring a CA signature on the subtree; CAs are only expected to sign two subtrees (Section 4.4) for each checkpoint
- * Rate-limiting requests

Acknowledgements

This document stands on the shoulders of giants and builds upon decades of work in TLS authentication, X.509, and Certificate Transparency. The authors would like to thank all those who have contributed over the history of these protocols.

The authors additionally thank Bob Beck, Ryan Dickson, Aaron Gable, Nick Harper, Dennis Jackson, Chris Patton, Ryan Sleevi, and Emily Stark for many valuable discussions and insights which led to this document. We wish to thank Mia Celeste in particular, whose implementation of an earlier draft revealed several pitfalls.

The idea to mint tree heads infrequently was originally described by Richard Barnes in [STH-Discipline]. The size optimization in Merkle Tree Certificates is an application of this idea to the certificate itself.

Change log

RFC Editor's Note: Please remove this section prior to publication of a final version of this document.

Since draft-davidben-tls-merkle-tree-certs-00

- * Simplify hashing by removing the internal padding to align with block size. #72
- * Avoid the temptation of floating points. #66
- * Require lifetime to be a multiple of batch_duration. #65
- * Rename window to validity window. #21
- * Split Assertion into Assertion and AbridgedAssertion. The latter is used in the Merkle Tree and HTTP interface. It replaces subject_info by a hash, to save space by not serving large post-quantum public keys. The original Assertion is used everywhere else, including BikeshedCertificate. #6
- * Add proper context to every node in the Merkle tree. #32

- * Clarify we use a single CertificateEntry. #11
- * Clarify we use POSIX time. #1
- * Elaborate on CA public key and signature format. #27
- * Miscellaneous changes.

Since draft-davidben-tls-merkle-tree-certs-01

- * Minor editorial changes

Since draft-davidben-tls-merkle-tree-certs-02

- * Replace the negotiation mechanism with TLS Trust Anchor Identifiers.

Since draft-davidben-tls-merkle-tree-certs-03

- * Switch terminology from "subscriber" to "authenticating party".
- * Use <1..2²⁴-1> encoding for all certificate types in the CertificateEntry TLS message
- * Clarify discussion and roles in transparency ecosystem
- * Update references

Since draft-davidben-tls-merkle-tree-certs-04

Substantially reworked the design. The old design was essentially the landmark checkpoint and CA-built logs ideas, but targeting only the optimized and slow issuance path, and with a more bespoke tree structure:

In both draft-04 and draft-05, a CA looks like today' s CAs except that they run some software to publish what they issue and sign tree heads to certify certificates in bulk.

In draft-04, the CA software publishes certificates in a bunch of independent Merkle trees. This is very easy to do as a collection of highly cacheable, immutable static files because each tree is constructed independently, and never appended to after being built. In draft-05, the certificates are published in a single Merkle tree. The [TLOG-TILES] interface allows such trees to also use highly cacheable, immutable static files.

In draft-04, there only are hourly tree heads. Clients are provisioned with tree heads ahead of time so we can make small, inclusion-proof-only certificates. In draft-05, the ecosystem must coordinate on defining "landmark" checkpoints. Clients are provisioned with subtrees describing landmark checkpoints ahead of time so we can make small, inclusion-proof-only certificates.

In draft-04, each tree head is independent. In draft-05, each landmark checkpoint contains all the previous checkpoints.

In draft-04, the independent tree heads were easily prunable. In draft-05, we define how to prune a Merkle tree.

In draft-04, there is no fast issuance mode. In draft-05, frequent, non-landmark checkpoints can be combined with inclusion proofs and witness signatures for fast issuance. This is essentially an STH and inclusion proof in CT.

Since draft-davidben-tls-merkle-tree-certs-05

- * Add some discussion on malleability
- * Discuss the monitoring impacts of the responsibility shift from CA with log quorum to CA+log with mirror quorum
- * Sketch out a more concrete initial ACME extension

Authors' Addresses

David Benjamin
Google LLC
Email: davidben@google.com

Devon O'Brien
Google LLC
Email: asymmetric@google.com

Bas Westerbaan
Cloudflare
Email: bas@cloudflare.com

Luke Valenta
Cloudflare
Email: lvalenta@cloudflare.com

Filippo Valsorda
Geomys
Email: ietf@filippo.io