

anima Working Group
Internet-Draft
Intended status: Standards Track
Expires: 28 January 2021

M. Richardson
Sandelman Software Works
J. Yang
Huawei Technologies Co., Ltd.
27 July 2020

Security and Operational considerations for manufacturer installed keys
and anchors
draft-richardson-secdispatch-idevid-considerations-02

Abstract

This document provides a nomenclature to describe ways in which manufacturers secure private keys and public trust anchors in devices.

RFCEditor: please remove this paragraph. This work is occurring in <https://github.com/mcr/idevid-security-considerations>

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 28 January 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	4
2. Applicability Model	5
2.1. A reference manufacturing/boot process	6
3. Types of Trust Anchors	7
3.1. Secured First Boot Trust Anchor	8
3.2. Software Update Trust Anchor	8
3.3. Trusted Application Manager anchor	8
3.4. Public WebPKI anchors	9
3.5. DNSSEC root	9
3.6. What else?	9
4. Types of Identities	9
4.1. Manufacturer installed IDevID certificates	10
4.1.1. Operational Considerations for Manufacturer IDevID Public Key Infrastructure	10
4.1.2. Key Generation process	11
5. Public Key Infrastructures (PKI)	14
5.1. Number of levels of certification authorities	14
5.2. Protection of CA private keys	15
5.3. Supporting provisioned anchors in devices	16
6. Evaluation Questions	16
6.1. Integrity and Privacy of on-device data	16
6.2. Integrity and Privacy of device identify infrastructure	17
6.3. Integrity and Privacy of included trust anchors	17
7. Privacy Considerations	18
8. Security Considerations	18
9. IANA Considerations	18
10. Acknowledgements	18
11. Changelog	18
12. References	18
12.1. Normative References	18
12.2. Informative References	19
Authors' Addresses	23

1. Introduction

An increasing number of protocols derive a significant part of their security by using trust anchors that are installed by manufacturers. Disclosure of the list of trust anchors does not usually cause a problem, but changing them in any way does. This includes adding, replacing or deleting anchors.

Many protocols also leverage manufacturer installed identities. These identities are usually in the form of [ieee802-1AR] Initial Device Identity certificates (IDevID). The identity has two components: a private key that must remain under the strict control of a trusted part of the device, and a public part (the certificate), which (ignoring, for the moment, personal privacy concerns) may be freely disclosed.

There also situations where identities which are tied up in the provision of symmetric shared secret. A common example is the SIM card ([_3GPP.51.011]), which now comes as a virtual SIM, but which is usually not provisioned at the factory. The provision of an initial, per-device default password also falls into the category of symmetric shared secret.

It is further not unusual for many devices (particularly smartphones) to also have one or more group identity keys. This is used in, for instance, in [fidotechnote] to make claims about being a particular model of phone (see [I-D.richardson-rats-usecases]). The keypair that does this is loaded into large batches of phones for privacy reasons.

The trust anchors are used for a variety of purposes. The following uses are specifically called out:

- * to validate the signature on a software update (as per [I-D.ietf-suit-architecture]).
- * to verify the end of a TLS Server Certificate, such as when setting up an HTTPS connection.
- * to verify the [RFC8366] format voucher that provides proof of an ownership change

Device identity keys are used when performing enrollment requests (in [I-D.ietf-anima-bootstrapping-keyinfra], and in some uses of [I-D.ietf-emu-eap-noob]). The device identity certificate is also used to sign Evidence by an Attesting Environment (see [I-D.ietf-rats-architecture]).

These security artifacts are used to anchor other chains of information: an EAT Claim as to the version of software/firmware running on a device (XXX and [I-D.birkholz-suit-coswid-manifest]), an EAT claim about legitimate network activity (via [I-D.birkholz-rats-mud], or embedded in the IDevID in [RFC8520]). Known software versions lead directly to vendor/distributor signed Software Bill of Materials (SBOM), such as those described by [I-D.ietf-sacm-coswid] and the NTIA/SBOM work [ntiasbom] and CISQ/OMG SBOM work underway [cisqsbom].

In order to manage risks and assess vulnerabilities in a Supply Chain, it is necessary to determine a degree of trustworthiness in each device. A device may mislead audit systems as to its provenance, about its software load or even about what kind of device it is. (see [RFC7168] for a humorous example). In order to properly assess the security of a Supply Chain it is necessary to understand the kinds and severity of the threats which a device has been designed to resist. To do this, it is necessary to understand the ways in which the different trust anchors and identities are initially provisioned, are protected, and are updated.

To do this, this document details the different trust anchors (TrA) and identities (IDs) found in typical devices. The privacy and integrity of the TAs and IDs is often provided by a different, superior artifacts. This relationship is examined.

While many might desire to assign numerical values to different mitigation techniques in order to be able to rank them, this document does not attempt to do, as there are too many other (mostly human) factors that would come into play. Such an effort is more properly in the purview of a formal ISO9001 process such as ISO14001.

1.1. Terminology

This document is not a standards track document, and it does not make use of formal requirements language.

This section will be expanded to include needed terminology as required.

The words Trust Anchor are contracted to TrA rather than TA, in order not to confuse with [I-D.ietf-teep-architecture]'s "Trusted Application".

This document defines a number of hyphenated terms, and they are summarized here:

device-generated : a private or symmetric key which is generated on

the device

infrastructure-generated : a private or symmetric key which is generated by some system, likely located at factory that built the device

mechanically-installed : when a key or certificate is programmed into flash by out-of-band mechanism like JTAG

mechanically-transferred : when a key or certificate is transferred into a system via private interface, such as serial console, JTAG managed mailbox, or other physically private interface

network-transferred : when a key or certificate is transferred into a system using a network interface which would be available after the device has shipped. This applies even if the network is physically attached using a bed-of-nails.

device/infrastructure-co-generated : when a private or symmetric key is derived from a secret previously synchronized between the silicon vendor and the factory using a common algorithm.

2. Applicability Model

There is a wide variety of devices to which this analysis can apply. (See [I-D.bormann-lwig-7228bis]) This document will use a J-group class C13 as a sample. This class is sufficiently large to experience complex issues among multiple CPUs, packages and operating systems, but at the same time, small enough that this class is often deployed in single-purpose IoT-like uses. Devices in this class often have Secure Enclaves (such as the "Grapeboard"), and can include silicon manufacturer controlled processors in the boot process (the Raspberry PI boots under control of the GPU).

Almost all larger systems (servers, laptops, desktops) include a Baseboard Management Controller (BMC), which ranges from a M-Group Class 3 MCU, to a J-Group Class 10 CPU (see, for instance [openbmc] which uses a Linux kernel and system inside the BMC). As the BMC usually has complete access to the main CPU's memory, I/O hardware and disk, the boot path security of such a system needs to be understood first as being about the security of the BMC.

2.1. A reference manufacturing/boot process

In order to provide for immutability and privacy of the critical TANs and IDs, many CPU manufacturers will provide for some kind of private memory area which is only accessible when the CPU is in certain privileged states. See the Terminology section of [I-D.ietf-teep-architecture], notably TEE, REE, and TAM, and also section 4, Architecture.

The private memory that is important is usually non-volatile and rather small. It may be located inside the CPU silicon die, or it may be located externally. If the memory is external, then it is usually encrypted by a hardware mechanism on the CPU, with only the key kept inside the CPU.

The entire mechanism may be external to the CPU in the form of a hardware-TPM module, or it may be entirely internal to the CPU in the form of a firmware-TPM. It may use a custom interface to the rest of the system, or it may implement the TPM 1.2 or TPM 2.0 specifications. Those details are important to performing a full evaluation, but do not matter that to this model (see initial-enclave-location below).

During the manufacturing process, once the components have been soldered to the board, the system is usually put through a system-level test. This is often done on as a "bed-of-nails" test [BedOfNails], where the board has key points attached mechanically to a test system. A [JTAG] process tests the System Under Test, and then initializes some firmware into the still empty flash storage. It is now common for a factory test image to be loaded first: this image will include code to initialize the private memory key described above, and will include a first-stage bootloader and some kind of (primitive) Trusted Application Manager (TAM). Embedded in the stage one bootloader will be a Trust Anchor that is able to verify the second-stage bootloader image.

After the system has undergone testing, the factory test image is erased, leaving the first-stage bootloader. One or more second-stage bootloader images is installed. The production image may be installed at that time, or if the second-stage bootloader is able to install it over the network, it may be done that way instead.

There are many variations of the above process, and this section is not attempting to be prescriptive, but to provide enough illustration to motivate subsequent terminology.

There process may be entirely automated, or it may be entirely driven by humans working in the factory.

Or a combination of the above.

These steps may all occur on an access-controlled assembly line, or the system boards may be shipped from one place to another (maybe another country) before undergoing testing.

Some systems are intended to be shipped in a tamper-proof state, but it is usually not desirable that bed-of-nails testing be possible without tampering, so the initialization process is usually done prior to rendering the system tamper-proof.

Quality control testing may be done prior to as well as after the application of tamper-proofing, as systems which do not pass inspection may be reworked to fix flaws, and this should ideally be impossible once the system has been made tamper-proof.

3. Types of Trust Anchors

Trust Anchors are fundamentally public keys. They are used to validate other digitally signed artifacts. Typically, these are chains of PKIX certificates leading to an End-Entity certificate (EE).

The chains are usually presented as part of an externally provided object, with the term "externally" to be understood as being as close as untrusted flash, to as far as objects retrieved over a network.

There is no requirement that there be any chain at all: the trust anchor can be used to validate a signature over a target object directly.

The trust anchors are often stored in the form of self-signed certificates. The self-signature does not offer any cryptographic assurance, but it does provide a form of error detection, providing verification against non-malicious forms of data corruption. If storage is at a premium (such as inside-CPU non-volatile storage) then only the public key itself need to be stored. For a 256-bit ECDSA key, this is 32-bytes of space.

When evaluating the degree of trust for each trust anchor there are four aspects that need to be determined:

- * can the trust anchor be replaced or modified?
- * can additional trust anchors be added?
- * can trust anchors be removed?

* how is the private key associated with the trust anchor stored?

The first three things are device specific properties of how the integrity of the trust anchor is maintained.

The fourth property has nothing to do with the device, but has to do with the reputation and care of the entity that maintains the private key.

Different anchors have different purposes.

These are:

3.1. Secured First Boot Trust Anchor

This anchor is part of the first-stage boot loader, and it is used to validate a second-stage bootloader which may be stored in external flash.

3.2. Software Update Trust Anchor

This anchor is used to validate the main application (or operating system) load for the device.

It can be stored in a number of places. First, it may be identical to the Secure Boot Trust Anchor.

Second, it may be stored in the second-stage bootloader, and therefore it's integrity is protected by the Secured First Boot Trust Anchor.

Third, it may be stored in the application code itself, where the application validates updates to the application directly (update in place), or via a double-buffer arrangement. The initial (factory) load of the application code initializes the trust arrangement.

In this situation the application code is not in a secured boot situation, as the second-stage bootloader does not validate the application/operating system before starting it, but it may still provide measured boot mechanism.

3.3. Trusted Application Manager anchor

This anchor is part of a [I-D.ietf-teep-architecture] Trusted Application Manager. Code which is signed by this anchor will be given execution privileges as described by the manifest which accompanies the code. This privilege may include updating anchors.

3.4. Public WebPKI anchors

These anchors are used to verify HTTPS certificates from web sites. These anchors are typically distributed as part of desktop browsers, and via desktop operating systems.

The exact set of these anchors is not precisely defined: it is usually determined by the browser vendor (e.g., Mozilla, Google, Apple, Safari, Microsoft), or the operating system vendor (e.g., Apple, Google, Microsoft, Ubuntu). In most cases these vendors look to the CA/Browser Forum ([CABFORUM]) for inclusion criteria.

3.5. DNSSEC root

This anchor is part of the DNS Security extensions. It provides an anchor for securing DNS lookups. Secure DNS lookups may be important in order to get access to software updates. This anchor is now scheduled to change approximately every 3 years, with the new key announced several years before it is used, making it possible to embed a keys that will be valid for up to five years.

This trust anchor is typically part of the application/operating system code and is usually updated by the manufacturer when they do updates. However, a system which is connected to the Internet may update the DNSSEC anchor itself through the mechanism described in [RFC5011].

There are concerns that there may be a chicken and egg situation for devices have remained in a powered off state (or disconnected from the Internet) for some period of years. That upon being reconnected, that the device would be unable to do DNSSEC validation. This failure would result in them being unable to obtain operating system updates that would then include the updates to the DNSSEC key.

3.6. What else?

TBD?

4. Types of Identities

Identities are installed during manufacturing time for a variety of purposes.

Identities require some private component. Asymmetric identities (e.g., RSA, ECDSA, EdDSA systems) require a co-responding public component, usually in the form of a certificate signed by a trusted third party.

The process of making this coordinated key pair and then installing it into the device is called identity provisioning.

4.1. Manufacturer installed IDevID certificates

[ieee802-1AR] defines a category of certificates that are to be installed by the manufacturer, which contain at the least, a device unique serial number.

A number of protocols depend upon this certificate.

- * [I-D.ietf-anima-bootstrapping-keyinfra] introduces a mechanism for new devices (called pledges) to be onboarded into a network without intervention from an expert operator. A number of derived protocols such as {{I-D.
- * [I-D.ietf-rats-architecture] depends upon a key provisioned into the Attesting Environment to sign Evidence.
- * [I-D.ietf-suit-architecture] may depend upon a key provisioned into the device in order to decrypt software updates.
- * TBD

4.1.1. Operational Considerations for Manufacturer IDevID Public Key Infrastructure

The manufacturer has the responsibility to provision a keypair into each device as part of the manufacturing process. There are a variety of mechanisms to accomplish this, which this document will overview.

There are three fundamental ways to generate IDevID certificates for devices:

1. generating a private key on the device, creating a Certificate Signing Request (or equivalent), and then returning a certificate to the device.
2. generating a private key outside the device, signing the certificate, and then installing both into the device.
3. deriving the private key from a previously installed secret seed, that is shared with only the manufacturer

There is a fourth situation where the IDevID is provided as part of a Trusted Platform Module (TPM), in which case the TPM vendor may be making the same tradeoffs.

The document [I-D.moskowitz-ecdsa-pki] provides some practical instructions on setting up a reference implementation for ECDSA keys using a three-tier mechanism.

This document recommends the use of ECDSA keys for the root and intermediate CAs, but there may be operational reasons why an RSA intermediate CA will be required for some legacy TPM equipment.

4.1.2. Key Generation process

4.1.2.1. On-device private key generation

Generating the key on-device has the advantage that the private key never leaves the device. The disadvantage is that the device may not have a verified random number generator. [factoringrsa] is an example of this scenario!

There are a number of options of how to get the public key securely from the device to the certification authority.

This transmission must be done in an integral manner, and must be securely associated with the assigned serial number. The serial number goes into the certificate, and the resulting certificate needs to be loaded into the manufacturer's asset database. This asset database needs to be shared with the MASA.

One way to do the transmission is during a factory Bed of Nails test (see [BedOfNails]) or Boundary Scan. When done via a physical connection like this, then this is referred to as a _device-generated_ / _mechanically-transferred_ .

There are other ways that could be used where a certificate signing request is sent over a special network channel when the device is powered up in the factory. This is referred to as the _device-generated_ / _network-transferred_ method.

Regardless of how the certificate signing request is sent from the device to the factory, and the certificate is returned to the device, a concern from production line managers is that the assembly line may have to wait for the certification authority to respond with the certificate.

After the key generation, the device needs to set a flag such that it no longer generates a new key, or will accept a new IDevID via the factory connection. This may be a software setting, or could be as dramatic as blowing a fuse.

The risk is that if an attacker with physical access is able to put the device back into an unconfigured mode, then the attacker may be able to substitute a new certificate into the device. It is difficult to construct a rationale for doing this, unless the network initialization also permits an attacker to load or replace trust anchors at the same time.

Because the key is generated inside the device, it is assumed that the device can never be convinced to disclose the private key.

4.1.2.2. Off-device private key generation

Generating the key off-device has the advantage that the randomness of the private key can be better analyzed. As the private key is available to the manufacturing infrastructure, the authenticity of the public key is well known ahead of time.

If the device does not come with a serial number in silicon, then one should be assigned and placed into a certificate. The private key and certificate could be programmed into the device along with the initial bootloader firmware in a single step.

Aside from the change of origin for the randomness, a major advantage of this mechanism is that it can be done with a single write to the flash. The entire firmware of the device, including configuration of trust anchors and private keys can be loaded in a single write pass. Given some pipelining of the generation of the keys and the creation of certificates, it may be possible to install unique identities without taking any additional time.

The major downside to generating the private key off-device is that it could be seen by the manufacturing infrastructure. It could be compromised by humans in the factory, or the equipment could be compromised. The use of this method increases the value of attacking the manufacturing infrastructure.

If keys are generated by the manufacturing plant, and are immediately installed, but never stored, then the window in which an attacker can gain access to the private key is immensely reduced.

As in the previous case, the transfer may be done via physical interfaces such as bed-of-nails, giving the `_infrastructure-generated_` / `_mechanically-transferred_` method.

There is also the possibility of having a `_infrastructure-generated_` / `_network-transferred_` key. There is support for "server-generated" keys in [RFC7030], [I-D.gutmann-scep], and [RFC4210]. All methods strongly recommend encrypting the private key for transfer.

This is difficult to comply with as there is not yet any private key material in the device, so in many cases it will not be possible to encrypt the private key.

4.1.2.3. Key setup based on 256-bit secret seed

A hybrid of the previous two methods leverages a symmetric key that is often provided by a silicon vendor to OEM manufacturers.

Each CPU (or a Trusted Execution Environment [I-D.ietf-tee-architecture], or a TPM) is provisioned at fabrication time with a unique, secret seed, usually at least 256-bits in size.

This value is revealed to the OEM board manufacturer only via a secure channel. Upon first boot, the system (probably within a TEE, or within a TPM) will generate a key pair using the seed to initialize a Pseudo-Random-Number-Generator (PRNG). The OEM, in a separate system, will initialize the same PRNG and generate the same key pair. The OEM then derives the public key part, signs it and turns it into a certificate. The private part is then destroyed, ideally never stored or seen by anyone. The certificate (being public information) is placed into a database, in some cases it is loaded by the device as its IDevID certificate, in other cases, it is retrieved during the onboarding process based upon a unique serial number asserted by the device.

This method appears to have all of the downsides of the previous two methods: the device must correctly derive its own private key, and the OEM has access to the private key, making it also vulnerable. The secret seed must be created in a secure way and it must also be communicated securely.

There are some advantages to the OEM however: the major one is that the problem of securely communicating with the device is outsourced to the silicon vendor. The private keys and certificates may be calculated by the OEM asynchronously to the manufacturing process, either done in batches in advance of actual manufacturing, or on demand when an IDevID is demanded. Doing the processing in this way permits the key derivation system to be completely disconnected from any network, and requires placing very little trust in the system assembly factory. Operational security such as often incorrectly presented fictionalized stories of a "mainframe" system to which only physical access is permitted begins to become realistic. That trust has been replaced with a heightened trust placed in the silicon (integrated circuit) fabrication facility.

The downsides of this method to the OEM are: they must be supplied by a trusted silicon fabrication system, which must communicate the set of secrets seeds to the OEM in batches, and they OEM must store and care for these keys very carefully. There are some operational advantages to keeping the secret seeds around in some form, as the same secret seed could be used for other things. There are some significant downsides to keeping that secret seed around.

5. Public Key Infrastructures (PKI)

[RFC5280] describes the format for certificates, and numerous mechanisms for doing enrollment have been defined (including: EST: [RFC7030], CMP: [RFC4210], SCEP [I-D.gutmann-scep]).

[RFC5280] provides mechanisms to deal with multi-level certification authorities, but it is not always clear what operating rules apply.

PKIs can suffer two kinds of failures: 1. disclosure of a private key. 2. loss of a private key.

A PKI which discloses one or more private certification authority keys is no longer secure. An attacker can create new identities, and forge certificates connecting existing identities to attacker controlled public/private keypairs. This can permit the attacker to impersonate the specific device.

If the PKI uses Certificate Revocation Lists (CRL)s, then an attacker can also revoke existing identities.

In the other direction, a PKI which loses access to a private key can no longer function. Existing identities may continue to function, unless CRLs or OCSP is in use, in which case, the inability to sign a fresh CRL or OCSP response will result in all identities becoming invalid.

This section details some nomenclature about the structure of certification authorities.

5.1. Number of levels of certification authorities

The certification authority (CA) starts with a Trust Anchor (TA). This is counted as the zeroth level of the authority.

In the degenerate case of a self-signed certificate, then this a level zero PKI.

```
.-----<-. |Issuer= X | | |Subject=X |- ' '-----'
```

The Trust Anchor signs one or more certificates. When this first level authority trusts only End-Entity (EE) certificates, then this is a level one PKI.

```
.-----<-. |Issuer= X | |Subject=X |- ' '-----' | -----\ v | .
---EE---. .---EE---. |Issuer= X | |Issuer=
X | |Subject=Y1| |Subject=Y2| '-----' '-----'
```

When this first level authority signs intermediate certification authorities, and those certification authorities sign End-Entity certificates, then this is a level two PKI.

```
.-----<-. |Issuer= X | |Subject=X |- ' '-----' |
-----\ v | .-----. .-----. |Issuer= X | |Issuer=
X | |Subject=Y1| |Subject=Y2| '-----' '-----' | -----\ | -----\
V | V | .---EE---. .---EE---. .---EE---. .---EE---. |Issuer=
Y1 | |Issuer= Y1| |Issuer= Y2| |Issuer=
Y2 | |Subject=Z1| |Subject=Z1| |Subject=Z3| |Subject=Z4| '-----'
'-----' '-----' '-----'
```

In general, when arranged as a tree, with the End-Entity certificates at the bottom, and the Trust Anchor at the top, then the level is where the deepest EE certificates are, counting from zero.

It is quite common to have a two-level PKI, where the root of the CA is stored in a Hardware Security Module, while the level one intermediate CA is available online.

5.2. Protection of CA private keys

The private key for the certification authorities must be protected from disclosure. The strongest protection is afforded by keeping them in a offline device, passing Certificate Signing Requests (CSR)s to the offline device by human process.

For examples of extreme measures, see [kskceremony]. There is however a wide spectrum of needs, as exemplified in [rootkeyceremony]. The SAS70 audit standard is usually used as a basis for the Ceremony, see [keyceremony2].

This is inconvenient, and may involve latencies of days, possibly even weeks to months if the offline device is kept in a locked environment that requires multiple keys to be present.

There is therefore a tension between protection and convenience. This is often accomplished by having some levels of the PKI be offline, and some levels of the PKI be online.

There is usually a need to maintain backup copies of the critical keys. It is often appropriate to use secret splitting technology such as Shamir Secret Sharing among a number of parties [shamir79]. This mechanism can be setup such that some threshold k (less than the total n) of shares are needed in order to recover the secret.

5.3. Supporting provisioned anchors in devices

IDevID-type Identity (or Birth) Certificates which are provisioned into devices need to be signed by a certification authority maintained by the manufacturer. The manufacturer needs to maintain availability of this PKI.

Trust anchors which are provisioned in the devices will have corresponding private keys maintained by the manufacturer. The trust anchors will often anchor a PKI which is going to be used for a particular purpose. There will be End-Entity (EE) certificates of this PKI which will be used to sign particular artifacts (such as software updates), or communications protocols (such as TLS connections). The private key associated with these EE certificates are not stored in the device, but are maintained by the manufacturer. These need even more care than the private keys stored in the devices, as compromise of the software update key compromises all of the devices, not just a single device.

6. Evaluation Questions

This section recaps the set of questions that may need to be answered. This document does not assign valuation to the answers.

6.1. Integrity and Privacy of on-device data

`initial-enclave-location` : Is the location of the initial software trust anchor internal to the CPU package?

`initial-enclave-integrity-key` : If the first-stage bootloader is external to the CPU, and it is integrity protected, where is the key used to check the integrity?

`initial-enclave-privacy-key` : If the first-stage data is external to the CPU, is it encrypted?

`first-stage-initialization` : The number of people involved in the

first stage initialization. An entirely automated system would have a number zero. A factory with three 8 hour shifts might have a number that is a multiple of three. A system with humans involved may be subject to bribery attacks, while a system with no humans may be subject to attacks on the system which are hard to notice.

first-second-stage-gap : If a board is initialized with a first-stage bootloader in one location (factory), and then shipped to another location, there may situations where the device can not be locked down until the second step.

6.2. Integrity and Privacy of device identify infrastructure

For IDevID provisioning, which includes a private key and matching certificate installed into the device, the associated public key infrastructure that anchors this identity must be maintained by the manufacturer.

identity-pki-level : how deep are the IDevID certificates that are issued?

identity-time-limits-per-intermediate : how long is each intermediate CA maintained before before a new intermediate CA key is generated? There may be no time limit, only a device count limit.

identity-number-per-intermediate : how many identities are signed by a particular intermediate CA before it is retired? There may be no numeric limit, only a time limit.

identity-anchor-storage : how is the root CA key stored? How many people are needed to recover it?

6.3. Integrity and Privacy of included trust anchors

For each trust anchor (public key) stored in the device, there will be an associated PKI. For each of those PKI the following questions need to be answered.

pki-level : how deep is the EE that will be evaluated

pki-level-locked : (a boolean) is the level where the EE cert will be found locked by the device, or can levels be added or deleted by the PKI operator without code changes to the device.

pki-breadth : how many different EE certificates exist in this PKI

pki-lock-policy : can any EE certificate be used with this trust anchor to sign? Or, is there some kind of policy OID or Subject restriction? Are specific intermediate CAs needed that lead to the EE?

pki-anchor-storage: how is the root CA stored? How many people are needed to recover it?

7. Privacy Considerations

many yet to be detailed

8. Security Considerations

This entire document is a security considerations.

9. IANA Considerations

This document makes no IANA requests.

10. Acknowledgements

Robert Martin of MITRE provides some guidance about citing the SBOM efforts.

11. Changelog

12. References

12.1. Normative References

- [BCP14] 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/info/rfc8174>>.
- [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/info/rfc5280>>.
- [I-D.moskowitz-ecdsa-pki] Moskowitz, R., Birkholz, H., Xia, L., and M. Richardson, "Guide for building an ECC pki", Work in Progress, Internet-Draft, draft-moskowitz-ecdsa-pki-08, 14 February 2020, <<http://www.ietf.org/internet-drafts/draft-moskowitz-ecdsa-pki-08.txt>>.

[ieee802-1AR]

IEEE Standard, ., "IEEE 802.1AR Secure Device Identifier", 2009, <<http://standards.ieee.org/findstds/standard/802.1AR-2009.html>>.

12.2. Informative References

[I-D.richardson-anima-masa-considerations]

Richardson, M. and W. Pan, "Operational Considerations for Voucher infrastructure for BRSKI MASA", Work in Progress, Internet-Draft, draft-richardson-anima-masa-considerations-04, 9 June 2020, <<http://www.ietf.org/internet-drafts/draft-richardson-anima-masa-considerations-04.txt>>.

[I-D.ietf-anima-bootstrapping-keyinfra]

Pritikin, M., Richardson, M., Eckert, T., Behringer, M., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", Work in Progress, Internet-Draft, draft-ietf-anima-bootstrapping-keyinfra-41, 8 April 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-bootstrapping-keyinfra-41.txt>>.

[RFC5011] StJohns, M., "Automated Updates of DNS Security (DNSSEC) Trust Anchors", STD 74, RFC 5011, DOI 10.17487/RFC5011, September 2007, <<https://www.rfc-editor.org/info/rfc5011>>.

[RFC8366] Watsen, K., Richardson, M., Pritikin, M., and T. Eckert, "A Voucher Artifact for Bootstrapping Protocols", RFC 8366, DOI 10.17487/RFC8366, May 2018, <<https://www.rfc-editor.org/info/rfc8366>>.

[RFC7030] Pritikin, M., Ed., Yee, P., Ed., and D. Harkins, Ed., "Enrollment over Secure Transport", RFC 7030, DOI 10.17487/RFC7030, October 2013, <<https://www.rfc-editor.org/info/rfc7030>>.

[I-D.gutmann-scep]

Gutmann, P., "Simple Certificate Enrollment Protocol", Work in Progress, Internet-Draft, draft-gutmann-scep-16, 27 March 2020, <<http://www.ietf.org/internet-drafts/draft-gutmann-scep-16.txt>>.

[RFC4210] Adams, C., Farrell, S., Kause, T., and T. Mononen, "Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)", RFC 4210, DOI 10.17487/RFC4210, September 2005, <<https://www.rfc-editor.org/info/rfc4210>>.

- [_3GPP.51.011]
3GPP, "Specification of the Subscriber Identity Module - Mobile Equipment (SIM-ME) interface", 3GPP TS 51.011 4.15.0, 15 June 2005, <<http://www.3gpp.org/ftp/Specs/html-info/51011.htm>>.
- [BedOfNails]
Wikipedia, ., "Bed of nails tester", July 2020, <https://en.wikipedia.org/wiki/In-circuit_test#Bed_of_nails_tester>.
- [pelionfcu]
ARM Pelion, "Factory provisioning overview", 28 June 2020, <<https://www.pelion.com/docs/device-management-provision/1.2/introduction/index.html>>.
- [factoringrsa]
"Factoring RSA keys from certified smart cards: Coppersmith in the wild", 16 September 2013, <<https://core.ac.uk/download/pdf/204886987.pdf>>.
- [RambusCryptoManager]
Qualcomm press release, "Qualcomm Licenses Rambus CryptoManager Key and Feature Management Security Solution", 2014, <<https://www.rambus.com/qualcomm-licenses-rambus-cryptomanager-key-and-feature-management-security-solution/>>.
- [kskceremony]
Verisign, "DNSSEC Practice Statement for the Root Zone ZSK Operator", 2017, <<https://www.iana.org/dnssec/dps/zsk-operator/dps-zsk-operator-v2.0.pdf>>.
- [rootkeyceremony]
Community, "Root Key Ceremony, Cryptography Wiki", April 2020, <https://cryptography.fandom.com/wiki/Root_Key_Ceremony>.
- [keyceremony2]
Digi-Sign, "SAS 70 Key Ceremony", April 2020, <<http://www.digi-sign.com/compliance/key%20ceremony/index>>.
- [shamir79]
Shamir, A., "How to share a secret.", 1979, <<https://www.cs.jhu.edu/~sdoshi/crypto/papers/shamirturing.pdf>>.

- [nistsp800-57]
NIST, "SP 800-57 Part 1 Rev. 4 Recommendation for Key Management, Part 1: General", 1 January 2016, <<https://csrc.nist.gov/publications/detail/sp/800-57-part-1/rev-4/final>>.
- [fidotechnote]
FIDO Alliance, ., "FIDO TechNotes: The Truth about Attestation", July 2018, <<https://fidoalliance.org/fido-technotes-the-truth-about-attestation/>>.
- [ntiasbom] NTIA, ., "NTIA Software Component Transparency", n.d., <<https://www.ntia.doc.gov/SoftwareTransparency>>.
- [cisqsbom] CISQ/Object Management Group, ., "TOOL-TO-TOOL SOFTWARE BILL OF MATERIALS EXCHANGE", July 2020, <<https://www.it-cisq.org/software-bill-of-materials/index.htm>>.
- [openbmc] Linux Foundation/OpenBMC Group, ., "Defining a Standard Baseboard Management Controller Firmware Stack", July 2020, <<https://www.openbmc.org/>>.
- [JTAG] IEEE Standard, ., "1149.7-2009 - IEEE Standard for Reduced-Pin and Enhanced-Functionality Test Access Port and Boundary-Scan Architecture", 2009, <<https://ieeexplore.ieee.org/document/5412866>>.
- [rootkeyrollover]
ICANN, ., "Proposal for Future Root Zone KSK Rollovers", 2019, <<https://www.icann.org/en/system/files/files/proposal-future-rz-ksk-rollovers-01nov19-en.pdf>>.
- [CABFORUM] CA/Browser Forum, ., "CA/Browser Forum Baseline Requirements for the Issuance and Management of Publicly-Trusted Certificates, v.1.2.2", October 2014, <<https://cabforum.org/wp-content/uploads/BRv1.2.2.pdf>>.
- [I-D.richardson-rats-usecases]
Richardson, M., Wallace, C., and W. Pan, "Use cases for Remote Attestation common encodings", Work in Progress, Internet-Draft, draft-richardson-rats-usecases-07, 9 March 2020, <<http://www.ietf.org/internet-drafts/draft-richardson-rats-usecases-07.txt>>.

[I-D.ietf-suit-architecture]

Moran, B., Tschofenig, H., Brown, D., and M. Meriac, "A Firmware Update Architecture for Internet of Things", Work in Progress, Internet-Draft, draft-ietf-suit-architecture-11, 27 May 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-suit-architecture-11.txt>>.

[I-D.ietf-emu-eap-noob]

Aura, T. and M. Sethi, "Nimble out-of-band authentication for EAP (EAP-NOOB)", Work in Progress, Internet-Draft, draft-ietf-emu-eap-noob-02, 12 July 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-emu-eap-noob-02.txt>>.

[I-D.ietf-rats-architecture]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-05, 10 July 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-rats-architecture-05.txt>>.

[I-D.birkholz-suit-coswid-manifest]

Birkholz, H., "A SUIT Manifest Extension for Concise Software Identifiers", Work in Progress, Internet-Draft, draft-birkholz-suit-coswid-manifest-00, 17 July 2018, <<http://www.ietf.org/internet-drafts/draft-birkholz-suit-coswid-manifest-00.txt>>.

[I-D.birkholz-rats-mud]

Birkholz, H., "MUD-Based RATS Resources Discovery", Work in Progress, Internet-Draft, draft-birkholz-rats-mud-00, 9 March 2020, <<http://www.ietf.org/internet-drafts/draft-birkholz-rats-mud-00.txt>>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <<https://www.rfc-editor.org/info/rfc8520>>.

[I-D.ietf-sacm-coswid]

Birkholz, H., Fitzgerald-McKay, J., Schmidt, C., and D. Waltermire, "Concise Software Identification Tags", Work in Progress, Internet-Draft, draft-ietf-sacm-coswid-15, 1 May 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-sacm-coswid-15.txt>>.

[RFC7168] Nazar, I., "The Hyper Text Coffee Pot Control Protocol for Tea Efflux Appliances (HTCPCP-TEA)", RFC 7168, DOI 10.17487/RFC7168, April 2014, <<https://www.rfc-editor.org/info/rfc7168>>.

[I-D.bormann-lwig-7228bis]
Bormann, C., Ersue, M., Keranen, A., and C. Gomez,
"Terminology for Constrained-Node Networks", Work in Progress, Internet-Draft, draft-bormann-lwig-7228bis-06, 9 March 2020, <<http://www.ietf.org/internet-drafts/draft-bormann-lwig-7228bis-06.txt>>.

[I-D.ietf-teep-architecture]
Pei, M., Tschofenig, H., Thaler, D., and D. Wheeler,
"Trusted Execution Environment Provisioning (TEEP) Architecture", Work in Progress, Internet-Draft, draft-ietf-teep-architecture-12, 13 July 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-teep-architecture-12.txt>>.

Authors' Addresses

Michael Richardson
Sandelman Software Works

Email: mcr+ietf@sandelman.ca

Jie Yang
Huawei Technologies Co., Ltd.

Email: jay.yang@huawei.com