Device Identity rooted at TPM

As of now device certificate(`/config/device.cert.pem`) is created from ECDSA key pair generated in software. For enhanced security, private keys should be securely generated inside a Trusted Platform Module (TPM). When device certificate is generated out of these secure key pair, device identity becomes secure against spoofing.

The following sections describe various changes required for this feature.

Notion of "TPM enabled" device:

Software should be able to take advantage of TPM presence to provide enhanced security. At the same time, on devices where TPM is not available, the software should be able to provide existing functionality without any hindrance. To help this, there will be a provision on software to query the current device mode (say "tpm-enabled"). Based on current mode, software modules can choose whether to enable security features or fall back to legacy mode. e.g. TLS library will, based on this mode, decide whether to talk to TPM for signing the challenge.

When TPM keys are used, a file will be created under /persist/. E.g. /persist/config/tpm_in_use. Presence of this file can be used in other software modules to determine if the device is in "Tpm-enabled" mode.

E.g. Zedagent may use this file to report TPM operational information in the device Info message to the Controller. The file will be rewritten after every boot, to prevent accidental overwrite or deletion of /persist/config directory.

Taking Ownership of TPM and its Keys

There are 2 scenarios here:

One is when the device is used for the first time. E.g. device has been procured from a reseller, and one wants to initialize TPM for the first time. This is the factory reset mode, and USB installation can be treated as one of the factory reset methods. In this scenario, TPM keys will be erased, and new TPM keypair will be generated. This is to create a new identity for the device.

The other one is when device was previously onboarded and just the device certificate is deleted from /config directory. In this scenario, TPM key pair will not be erased, instead the existing keys will be re-used(or generated afresh if none exists) to generate the device certificate again. This is to make sure that upgrade/deletion of device certificate does not take away the device identity.

Image compatibility management

Devices with TPM capability running old software(without TPM support) will need seamless transition to new software that makes use of TPM. Similarly downgrade from new software to old software would need a seamless transition. To address this, the following is proposed:

"tpm-enabled" mode in software will be set only when device certificate is being generated; I.e. device has TPM chip, and we are generating device certificate . In other scenarios, software will use device certificates generated in software.

State Transition Diagram

