

Different aspects of onboarding for IoT/Edge Devices

draft-nordmark-iotops-onboarding

Erik Nordmark
Zededa

Is onboarding about networking?

- Need to grant access to (local/external) network
- In general
 - Could involve laptop/phone registering its MAC address by logging in using enterprise creds
 - Might have very restricted access until posture has been assessed (NEA - RFC5209, RFC7632)
- IoT with no/limited UI?
 - Nimble out-of-band authentication for EAP (EAP-NOOB)
 - Bootstrapping Remote Secure Key Infrastructure (BRSKI - [RFC 8995](#))
 - Device Provisioning Protocol (DPP)
- In some deployments, physical access to plug in Ethernet might be sufficient to get network access
 - That is, to send and receive IP packets

Finer-grained access control?

- Restrict device/application on device to reach certain destinations?
 - To protect the rest of the world from the device?
 - To protect the device from the Internet threats?
- MUD [RFC8520] can do that
- Assumes the application is defined by the device

Higher levels

- Can the device get e.g., the local OSPF configuration?
 - Pertinent for devices which are routers and switches
 - Requires mutual trust of some sort including remote attestation
- Can the device be managed by some management system?
 - Discovery of management system?
 - Device trusting the management system
 - The management system trusting the device is legitimate
- Device vs application on device?
 - At constrained device edge[1] devices run on or a few pre-determined applications
 - At smart device edge[1] applications can be deployed post device deployment
 - Onboarding device - to hardware maintenance and management system
 - Onboard each application to their “controller”

Roots of trust?

- Hardware manufacturer certificates
 - Can check with manufacturer that device is valid, but doesn't indicate management/controller
- Tracking the transfers of ownership through supply chain
 - Enables late binding to management/controller in FIDO[2]
 - The signature chain from manufacturer to end user establishes trust in controller
- Imprinting/configuring for/by the owner?
 - Including initial measurement/attestation of firmware/software

Example

- Project EVE[3] has a minimalistic but secure imprinting approach
- When software is installed (factory or elsewhere)
 - Imprint device which controller to trust (a root certificate) and initial URL to contact
 - Generate a device cert using the TPM
 - Extract the device certificate and pass to final user (paper, bar code, etc)
 - Perform initial measured boot to get baseline measurements
- Then in any order
 - User registers device certificate in controller
 - Device is installed and powered on and connects to its controller
- Now controller can specify which applications to deploy/boot/halt on device

Summary

- Don't assume a device runs (a single) pre-determined application(s)
- Support different policies for network access authentication
 - Is it the leader and onboarding to management/controller follows, or the other way around?
- Roots of trust and role of manufacturer is critical for onboarding

References

- [1] https://www.lfedge.org/wp-content/uploads/2020/07/LFEdge_Whitepaper.pdf
- [2] <https://fidoalliance.org/specs/FIDO/FIDO-Device-Onboard-RD-v1.0-20201202.html>
- [3] <https://github.com/lf-edge/eve>
- [4] <https://datatracker.ietf.org/doc/draft-irtf-t2trg-secure-bootstrapping/>