NIST Internal Report
NIST IR 8350

Foundational Concepts in Trusted IoT Device Network-Layer Onboarding

Enhancing Internet Protocol-Based IoT Device and Network Security

Final

Susan Symington

Blaine Mulugeta

William Polk*

Murugiah Souppaya*

Jeffrey Marron

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8350

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST Internal Report
NIST IR 8350

Foundational Concepts in Trusted IoT Device Network-Layer Onboarding

Enhancing Internet Protocol-Based IoT Device and Network Security

Final

William Polk*
Jeffrey Marron
Applied Cybersecurity Division
Information Technology Laboratory

Murugiah Souppaya*
Computer Security Division
Information Technology Laboratory

Susan Symington
Blaine Mulugeta
The MITRE Corporation
McLean, VA

Retired NIST Author*
**Former NIST employee; all work for this publication was done while at NIST.*

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8350

November 2025

U.S. Department of Commerce
Howard Lutnick, Secretary

National Institute of Standards and Technology
Craig Burkhardt, Acting Under Secretary of Commerce for Standards and Technology and Acting NIST Director

NIST IR 8350
November 2025

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at https://csrc.nist.gov/publications.

Publication History
Approved by the NIST Editorial Review Board on 2025-09-30

How to Cite this NIST Technical Series Publication
Polk W, Marron J, Souppaya M, Symington S, Mulugeta B (2025) Foundational Concepts in Trusted IoT Device Network-Layer Onboarding. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Internal Report (IR) NIST IR 8350. https://doi.org/10.6028/NIST.IR.8350

Author ORCID iDs
Jeffrey Marron 0000-0002-7871-683X
Blaine Mulugeta 0009-0004-1299-7482
Murugiah Souppaya 0000-0002-8055-8527

Contact Information
iot-onboarding@nist.gov

National Institute of Standards and Technology
Attn: Applied Cybersecurity Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 2000) Gaithersburg, MD 20899-2000

Additional Information
Additional information about this publication is available at https://csrc.nist.gov/pubs/ir/8350/final, including related content, potential updates, and document history.

Abstract

Internet of Things (IoT) devices are typically connected to a network. The steps performed to provision a device with its network credentials are referred to as network-layer onboarding (or simply, onboarding, assuming the network-layer context is understood). This paper proposes a definition for trusted network-layer onboarding. This paper is intended to introduce the reader to trusted network-layer onboarding; describe its capabilities, characteristics, and benefits; and explain the important role that onboarding can play in the protection of IoT devices and networks throughout the device lifecycle. By providing a common language that describes and clarifies various onboarding capabilities, this paper assists with discussion, characterization, and development of trusted onboarding solutions. This paper also describes a generic trusted onboarding process, defines onboarding functional roles and responsibilities, discusses onboarding-related aspects of IoT device lifecycle management, and explains how onboarding can enhance security capabilities that protect the device throughout its lifecycle.

Keywords

application-layer onboarding; authentication; bootstrapping; credentials; device lifecycle management; identity; internet of things (IoT); network-layer onboarding; onboarding

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analyses to advance the development and productive use of information technology. ITL's responsibilities include the development of management, administrative, technical, and physical standards and guidelines for the cost-effective security and privacy of other than national security-related information in federal information systems. The Special Publication 800-series reports on ITL's research, guidelines, and outreach efforts in information system security, and its collaborative activities with industry, government, and academic organizations.

Audience

The audience of this paper is intended to include IoT device manufacturers, integrators, and vendors; managers of networks to which IoT devices connect; service providers (internet service providers/cable operators and application platform providers) who want to simplify the IoT device connection process for their customers; industry consortia; standards development organizations; and any other individuals or organizations that are stakeholders in the effort to define open, standard, trusted, and scalable solutions for efficiently and easily providing IoT devices with the network credentials that they need to become operational.

Trademark Information

All names are registered trademarks or trademarks of their respective companies.

Patent Disclosure Notice

NOTICE: ITL has requested that holders of patent claims whose use may be required for compliance with the guidance or requirements of this publication disclose such patent claims to ITL. However, holders of patents are not obligated to respond to ITL calls for patents and ITL has not undertaken a patent search in order to identify which, if any, patents may apply to this publication.

As of the date of publication and following call(s) for the identification of patent claims whose use may be required for compliance with the guidance or requirements of this publication, no such patent claims have been identified to ITL.

No representation is made or implied by ITL that licenses are not required to avoid patent infringement in the use of this publication.

Table of Contents

No table of figures entries found.**List of Figures**

Acknowledgments

The authors wish to thank the following individuals for the generous contribution of expertise and time that they demonstrated by attending a Cisco-hosted meeting and expressing their views on many of the onboarding-related topics covered in this paper. We have developed the content of this paper, in large part, based on the discussion that ensued at that onboarding meeting. The contributors[1] are Allaukik Abhishek, Anurag Gupta, and Reed Hinkel of Arm; Darshak Thakore and Mark Walker of CableLabs; Owen Friel, Russ Gyurek, Eliot Lear, Peter Romness, and Bob Sayle of Cisco; William Barker of Dakota Consulting; Katherine Gronberg of Forescout; Nils Gerhardt of Global Platform; Saurabh Dadu of Intel; Drew Cohen, Geoff Matrangola, and Kevin Yeich of MasterPeace Solutions; Janet Jones of Microsoft; Parisa Grayeli, Josh Klosterman, and Blaine Mulugeta of The MITRE Corporation; Doug Montgomery, Ranga Mudumbai, and Monika Singh of the National Institute of Standards and Technology (NIST); Matt Tooley of NCTA – The Internet and Television Association; and Michael Montemurro of the Wi-Fi Alliance/Blackberry.

The authors also wish to thank the following individuals and organizations for lending their expertise to thoughtfully review earlier drafts and provide comments that have improved the final paper: Elaine Barker, NIST; Jon Boyens, NIST; Kevin Brady, NIST; Julie Chua, NIST; Barbara Cuthill, NIST; Cherilyn Pascoe, NIST; Bo-Chao Cheng, National Chung-Cheng University, Taiwan; Steve Clark, WISeKey; FIDO Alliance; Julio Merette, Device Authority; OPC Foundation; and Dr. Ravishankar.C.V., Sambhram Institute of Technology.

[1] The contributors listed here were employed with the listed organizations at the time of the workshop.

1. Introduction

Just like any other device, an Internet of Things (IoT) device needs appropriate credentials in order to connect to a network securely. A typical commercially available, mass-produced IoT device cannot be pre-provisioned with credentials for its local network during the manufacturing process. Instead, these credentials have to be provisioned to the device at deployment time using a set of steps known as *network-layer onboarding* (or simply *onboarding,* assuming the network-layer context is understood).

Network-layer onboarding is a critical and vulnerable part of the IoT device lifecycle. If onboarding is not performed securely, both the device itself and the network to which it connects are put at risk. This paper introduces trusted network-layer onboarding; describes its capabilities, characteristics, and benefits; and explains the important role that onboarding can play in the protection of IoT devices and networks throughout the device lifecycle.

Since the publication of this report as a draft, the National Cybersecurity Center of Excellence (NCCoE) completed a project to demonstrate Trusted IoT Device Network-Layer Onboarding and Lifecycle Management, resulting in NIST Special Publication (SP) 1800-36. That project provided practical implementations of various network-layer onboarding protocols and valuable insight that validates and expands upon many concepts described in this document. We have updated this document to include key lessons learned and other relevant information from that project.

1.1. Purpose

The purpose of this report is to:

- provide the reader with a thorough understanding of what trusted network-layer onboarding is, how it works, and the mechanisms, functional roles, and responsibilities that are associated with it.
- convey the potential of trusted network-layer onboarding to integrate with and enhance additional security protections to help protect IoT devices not just during the process of network credential provisioning, but throughout their entire lifecycle.

1.2. Scope

This document considers network access methods that use the Internet Protocol (IP). It assumes that IoT devices that use non-IP access methods such as Bluetooth Low Energy, ZigBee, Z-Wave, and 802.15 radios will connect to the IP network through a gateway. Non-IP networks are deployed extensively and have unique advantages and disadvantages that could be discussed with respect to onboarding. However, we have chosen not to address them within the scope of this paper. In addition, only network-layer onboarding using Wi-Fi and wired

Ethernet access technologies are in scope at this time. Although IP-based 802.15 networks and IP over cellular networks are also in common use, the discussion and illustrations primarily focus on Wi-Fi use case scenarios.

Note that this report mentions specific protocols, specifications, and mechanisms in order to cite real-world examples of the concepts under discussion. These references are for informational purposes only. They are not intended to imply that these protocols, specifications, or mechanisms must be implemented, nor do they imply recommendation or endorsement of them.

This document describes onboarding concepts and terminology. Readers looking for more implementation guidance may reference NIST SP 1800-36, which demonstrates these approaches in practice.

1.3. Report Structure

The remainder of this publication is organized into the following sections and appendices:

- Section 2 provides background on the challenges with current IoT device onboarding mechanisms, examining the security challenges and operational inefficiencies that exist in current onboarding approaches for both consumer and enterprise environments.
- Section 3 provides an overview of onboarding and related terminology, establishing fundamental concepts and definitions needed to understand trusted network-layer onboarding.
- Section 4 describes onboarding lifecycle management, presenting a comprehensive view of IoT device lifecycles from an onboarding perspective. The IoT device lifecycle is separated into two main phases: pre-market and post-market activities.
- Section 5 details the onboarding process steps, providing a generic, solution-neutral description of the activities performed to make an IoT device operational.
- Section 6 defines functional roles, describing the personnel roles and responsibilities for accomplishing onboarding throughout the IoT device lifecycle.
- Section 7 explains how trusted network onboarding serves as a foundation for ongoing IoT device security, integrating with and enhancing additional security capabilities that provide continuous protection throughout the IoT device lifecycle.
- The References section contains references cited throughout the publication.
- Appendix A provides an acronym and abbreviation list.

2. Challenges with Current IoT Device Onboarding Mechanisms

The number of IoT devices is increasing rapidly. With many billions of devices currently connected and many more expected to be added in the near future, it is clearly not realistic to expect to manage these IoT devices manually. We need scalable, automated mechanisms to safely manage these devices throughout their lifecycles. This report focuses mainly on the network-layer onboarding portion of the lifecycle during which the device is provisioned with its network credentials. This is a particularly vulnerable point in the device lifecycle because if this provisioning is not performed in a secure manner, then both the device and the network are at risk: devices are at risk of being onboarded to networks that are not authorized to control them, and networks are at risk of having unauthorized devices connect to them. The network-layer onboarding process of the device lifecycle is also important in terms of being able to serve as a foundation for initiating and enhancing additional security capabilities that can protect IoT devices on an ongoing basis.

The wide variety of IoT devices differ regarding power, memory, computation, and other resource characteristics. Another key difference among these devices is in how they are onboarded. Ideally, the first operation of an IoT device is to onboard itself [1], and the onboarding process should be open, standardized, trusted, scalable, and flexible enough to meet the needs of various use cases. These use-cases may differ widely, including consumer settings (i.e., home or small business networks) that do not have the support of dedicated, IT-knowledgeable personnel; enterprise settings, which could be expected to be supported by full-time IT professionals; and specific industry sectors that may have unique requirements. Because IoT devices typically lack screens and keyboards, trying to provision their credentials can be cumbersome, and a network communication protocol may be needed to interact with the devices. For consumers, trusted onboarding should be easy; for enterprises, it should enable large numbers of devices to be quickly provisioned with unique credentials. The security attributes of the onboarding process ensure that the network is not put at risk as new IoT devices are added to it.

Mechanisms used to perform IoT device onboarding tend, in large part, to be inefficient or insecure. For example, typical devices that are onboarded to most consumer home Wi-Fi networks all use the same pre-shared key to connect to that network. If multiple networks are available, an IoT device selects the network to connect with and provides the network password (i.e., the pre-shared key). Without a screen or keyboard, the processes of selecting the correct network and providing the network password can be not only awkward, but difficult to do securely. To make these steps easier, some devices have been equipped with Wi-Fi Protected Setup (WPS). WPS enables a consumer to onboard IoT devices by simply pressing a button that causes the network router to provide the devices with the password they need. Unfortunately, WPS has been shown to suffer from several security vulnerabilities [2]. In addition, it also requires a physical button, which can be cumbersome. Due to a lack of a functional user interface, some devices use Wi-Fi to enable a user to interface with the device and insecurely provision credentials over an open network.

Given the threats faced in today's internet, there is a desire for more security than can currently be provided by the same shared password for all devices on a network. Under a shared-password model, if a device presents the correct password, the network will permit it to connect. The network does not take into account the device's identity or type. If the password falls into the wrong hands, unauthorized devices may use it to connect to the network. Furthermore, although networks can falsely identify themselves, the device is not typically provided with any way to verify that the network to which it is connecting is the intended network. To address these problems, the typical consumer network onboarding process needs to be improved [3].

In contrast to the home environment, onboarding in an enterprise environment is typically based on a more robust security model that requires each device to have its own distinct credential to connect to the network. In the past, this often meant that the onboarding process was complex and resource-intensive. Some onboarding processes could take more than 20 minutes per device, require coordination, and sometimes entail conflict and tension among installation technicians, information technology (IT) network/security operations, and operational technology teams [4]. When onboarding is performed manually, it is time-consuming and error-prone. If it requires individuals to have access to device credentials, it is vulnerable to the risk of those credentials being disclosed to unauthorized parties.

Some enterprises require the ability to perform bulk onboarding—i.e., to provide many IoT devices with their network credentials quickly—which necessitates that the onboarding process be automated and zero-touch (i.e., no human intervention required). However, some zero-touch solutions require that the device's network and other locally-significant credentials be built into the device at the point of manufacture [4]. This effectively requires a manufacturer to customize devices on a per-customer, build-to-order basis, which is complex, inefficient, and expensive. It requires the device manufacturers to collect each customer's unique requirements, a process that can take weeks to complete and requires the engagement of multiple parties [5]. Then, the manufacturer configures the devices to specific customer needs (e.g., credentials/keys specific to the device's target network are loaded by the manufacturer), which, once completed, requires multiple rounds of testing with various parties within the customer organization that may take as long as three weeks to complete [5]. Next, training is required, involving the preparation of unique instructions. When the customer receives the device, activation of the device on its target network often requires the customer to complete a long list of manual steps. The complexity of the process, combined with the fact that it is susceptible to human error, makes it vulnerable to security risks.

Customizing each device's network and other locally-significant credentials at the point of manufacture in this manner is inefficient, complex, and potentially insecure. To take full advantage of economies of scale, manufacturers seek to make devices as identical as possible for all customers. This desire implies an onboarding solution that defers the provisioning of a device's network and other locally significant credentials from the time of manufacture to the time of deployment. Scalable, trusted network-layer onboarding solutions address these challenges by equipping devices with unique authoritative device credentials (e.g., unique

identity, private key, possibly a certificate) by the manufacturer and then leveraging these unique authoritative credentials at the time of deployment to perform zero-touch trusted network-layer onboarding and thereby equip the devices with unique locally-significant network credentials.

The unique authoritative device credentials that a manufacturer installs on each device (also referred to as the device's *birth certificate* or *bootstrapping credentials*—see Section 3) should be targeted to enable and simplify the network-layer onboarding process for its intended users, while keeping the device manufacturing process efficient. Ensuring that such an onboarding solution is trusted requires the network credentials to be provisioned to the device over an encrypted channel without providing anyone with access to the credentials, thereby protecting them from disclosure to unauthorized parties. In addition, although having a secure mechanism for provisioning a device's network credentials is very important, it is not sufficient for ensuring device and network security. Devices must be managed throughout their entire lifecycle and continue to be patched, updated, and monitored on an ongoing basis to address potential vulnerabilities, constrain unauthorized traffic, and otherwise ensure that they remain secure. Ideally, trusted network-layer onboarding can provide a secure foundation upon which to establish these additional ongoing security capabilities, thereby helping to enhance the protections that they provide to the device throughout the remainder of its lifecycle.

3. Onboarding Overview and Terminology

An ***onboarding solution*** is a product/technology/service combination that an organization will deploy to onboard devices to its network. The solution usually requires the support of people and processes throughout its lifecycle in alignment with the organization's policies. But what does it mean to onboard a device to a network, and how is it accomplished in a trusted manner?

This section provides a high-level overview of onboarding and defines basic onboarding-related terminology. As shown in Figure 11, an IoT device may undergo two levels of onboarding: one at the network layer, which enables it to connect securely to the network, and one at the application layer, which enables it to become operational at the application layer and thereby perform its intended function. The subject of this report is network-layer onboarding, but it is helpful to distinguish the two, understand their relationship, and appreciate that network-layer onboarding can provide a secure foundation for performing application-layer onboarding.

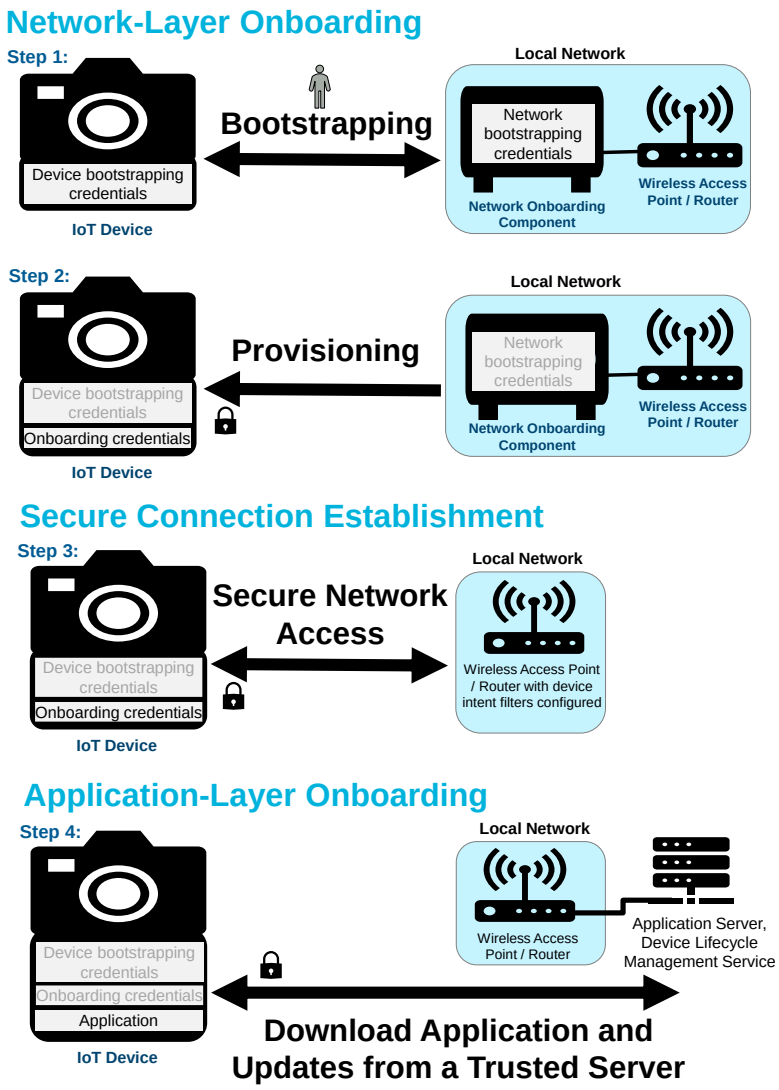

Figure 11 – General Overview of Onboarding Concepts

Network-layer onboarding consists of the actions required to provide an IoT device with the network credentials (and possibly other information) it needs to connect securely to a network.

As shown in Figure 11, network-layer onboarding consists of two subprocesses: bootstrapping and provisioning, which will be described in more detail later. After network-layer onboarding occurs, a device can use its newly provisioned credentials to establish secure access to the network. Upon establishing a secure connection, the device can automatically perform application-layer onboarding.

Application-layer onboarding consists of the steps required to provide an IoT device with the application-layer components (e.g., applications, updates, configurations) that it needs to execute its application and thereby perform as expected. Application-layer onboarding is analogous to network-layer onboarding; it occurs at the application layer rather than at the network layer.

Before network-layer onboarding can occur, certain elements are required to be in place: a networking onboarding component, device bootstrapping credentials, network bootstrapping credentials, and (if supported) a device information declaration. These are defined as follows.

A ***network onboarding component*** is a logical component that is authorized to act on behalf of the network to onboard devices that are authorized to connect to the network. The network onboarding component interacts with devices by using a ***network onboarding protocol***.

Device bootstrapping credentials are credentials that a device needs in order to establish communications with and be authenticated by the network onboarding component. Device bootstrapping credentials are typically installed on the device during manufacturing, and they must already be on the device before it initiates the network-layer onboarding process. The device bootstrapping credentials are authoritative credentials that are unique to the device and independent of context; they are sometimes referred to as the ***device birth certificate***. The device bootstrapping credentials always include some sort of secret (e.g., a private or secret pre-shared key), which the device will use to authenticate itself to the network onboarding component and to establish a secure communications channel with the network onboarding component.

In addition to a secret, device bootstrapping credentials may also contain information such as a device identifier and the Wi-Fi channel the device will use, if any. It may also include additional optional information associated with the device, such as device intent information (e.g., information the network can use to ensure that the device will be permitted to send and receive only the information that it requires to perform its intended function) and application-layer bootstrapping-related information (if any). The trustworthiness of the onboarding solution relies on the device bootstrapping credentials (especially the secret) being kept confidential. The bootstrapping credentials are a root of trust. It is assumed that the device bootstrapping credentials will not change over the lifetime of the device; the ability to onboard a device repeatedly depends upon its bootstrapping credentials remaining the same. If the credential is a certificate, it should not expire.

Network bootstrapping credentials are credentials that the network onboarding component requires so that the network can be authenticated by the device. They are authoritative credentials that are unique to the network. Network bootstrapping credentials have to be provided to the network onboarding component before the onboarding process is initiated, assuming the device will be authenticating the network as part of the network-layer onboarding process. Network bootstrapping credentials may include information such as the network identifier (e.g., X.509 certificate, service set identifier [SSID]) and a secret (e.g., private key). They will also identify the *network owner* (see Section 66).

A ***device information declaration*** is a trusted digital assertion with information about the IoT device, such as the identity or certificate of the device owner and possibly the identities or certificates of any other entities that the device owner has authorized to onboard the device on its behalf. When using an onboarding solution that supports proof-of-ownership verification, the IoT device's manufacturer will create a device information declaration for the IoT device that lists the device's owner and its authorized onboarders (if any). The manufacturer is usually considered the original device owner. As ownership of the device is transferred from one entity

to another, any ownership information in the device information declaration has to be kept up-to-date because the purpose of the device information declaration (or similar mechanism) is to prove who owns the device and who else (in addition to the device owner) is authorized to onboard the device. If the network that is trying to onboard the device (and thereby take control of it) is listed in the device information declaration as the device's owner or as an entity that is authorized to onboard the device on behalf of the owner, the device information declaration provides assurance to an IoT device that the network that is trying to onboard it is authorized to do so.

Figure 2 depicts a generic network-layer onboarding process to provide product-agnostic examples of the terminology introduced so far. Figure 2 shows a wireless network, but the network could also be a wired network. It also shows the bootstrapping credentials and the device information declaration that are present before network-layer onboarding begins. As shown in Figure 2, Network-layer onboarding consists of two subprocesses: bootstrapping and provisioning. Within industry, the term "bootstrapping" is used in many different contexts and has a variety of definitions [6]. In this report, we define bootstrapping as follows:

Bootstrapping is a subprocess of network-layer onboarding. It introduces just enough information to a device and a network onboarding component to enable them to trust one another and establish a secure channel. The trust established between the device and the network onboarding component through bootstrapping may be either mutual or one-way. The introduction within bootstrapping may be performed as an out-of-band (OOB) process requiring human interaction. For example, the ***device onboarder***—i.e., the person or a proxy that is performing the onboarding—may provide the network onboarding component with information regarding the device (e.g., the device's public key) or provide the device with information about the network onboarding component (e.g., the network's public key) or about the device's owner. However, not all bootstrapping mechanisms require human interaction. Bootstrapping mechanisms that do not require a device onboarder are referred to as ***zero-touch*** solutions.

In the context of network-layer onboarding, ***provisioning*** is the process of the network onboarding component providing onboarding credentials to an IoT device.

Onboarding credentials include the unique credentials that the IoT device requires to connect securely to the local network (e.g., network identifier, unique network password, X.509 certificate and associated private key, etc.). They may also include additional configuration information (e.g., URLs and public keys for reaching and authenticating controllers and servers) to enable the device to perform application-layer onboarding and thereby become operational at the application layer once it has securely connected to the network. In contrast to device bootstrapping credentials, which are independent of context, device onboarding credentials are locally significant.

As shown in Step 1, bootstrapping begins with the exchange of information required to enable the device and the network to authenticate each other. For example, the network onboarding component may be provided with the device's public key, and the device may be provided with the network's public key. This information exchange may be facilitated by a human or performed automatically without manual intervention. The device and the network use the

information that they receive in this exchange to authenticate each other, securely exchange cryptographic keys, and use those keys to establish a secure channel.

In Step 2, the network onboarding component uses the secure channel to provision the device with its onboarding credentials. Once these credentials are provisioned, network-layer onboarding is complete, and the network onboarding component's work is done. The IoT device thereafter interacts with the network directly rather than via the network onboarding component unless and until the device needs to be provisioned with new onboarding credentials.

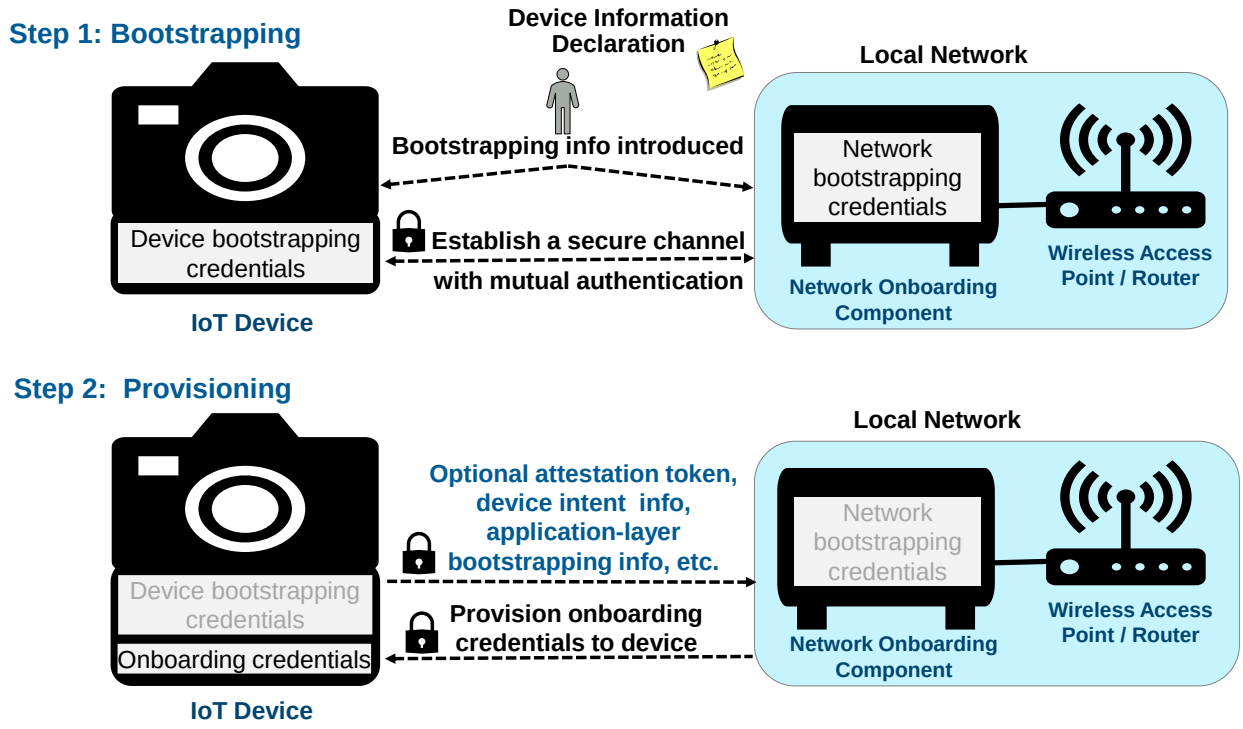

Figure 2 – Network-Layer Onboarding Steps

Step 2 may also optionally include the exchange of additional information (depicted in blue text) between the device and the network onboarding component that is designed to support and enhance security capabilities that can help protect the device on an ongoing basis after it has connected securely to the network. In Step 2, we list three examples of such information: a device attestation token, which the network can use to determine whether it trusts the device sufficiently to permit it to be onboarded; device intent information, which the network can use to ensure that the device will be permitted to send and receive only the information that it requires to perform its intended function; and application-layer bootstrapping information, which can be used as input to perform application-layer onboarding after the device connects to the network.

Similar to how network-layer onboarding expects the device and the network to be pre-provisioned with bootstrapping credentials, application-layer onboarding expects the device and the application servers or controllers involved in application-layer onboarding to be pre-

provisioned with application-layer bootstrapping credentials. It also requires an initial introduction of application-layer bootstrapping-related information to enable the device and the application server or controller to authenticate each other and establish a secure association. The secure channel that is established during network layer onboarding can serve as the mechanism for performing this initial introduction of application-layer bootstrapping-related information, as shown in Step 2. In this way, network-layer onboarding can support and enhance the security of application layer onboarding.

As shown in Figure 3, once the device has been provisioned with its network credentials, it uses those credentials to connect to the network securely. Note that if the optional device intent information had been sent to the network in Step 2, when the device connects to the network in Step 3, the network's routers have already been configured with access control filters to enforce the device's communications intent.

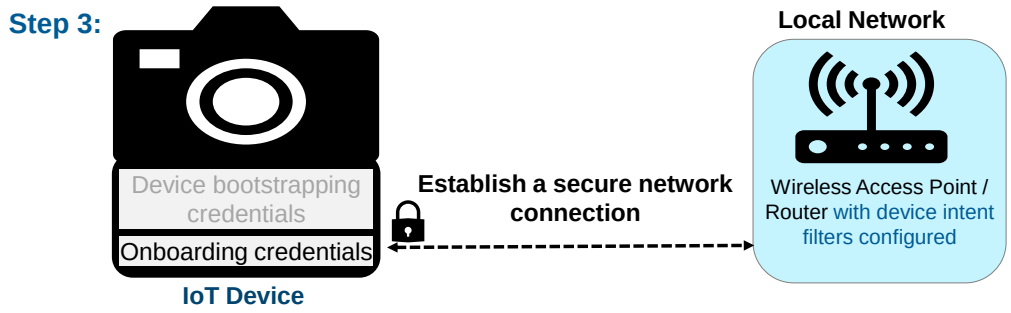

Figure 3 – Secure Connection Establishment

If the optional application-layer bootstrapping information had been exchanged in Step 2, then, as shown in Figure 4, application-layer onboarding can be automatically performed after the device connects to the network. The device and an application server can authenticate each other and establish a secure association that the application server uses to download the latest version of the application to the device. Similarly, the device and a lifecycle management service can authenticate each other and establish a secure association that enables the lifecycle management service to keep the device patched and updated on an ongoing basis.

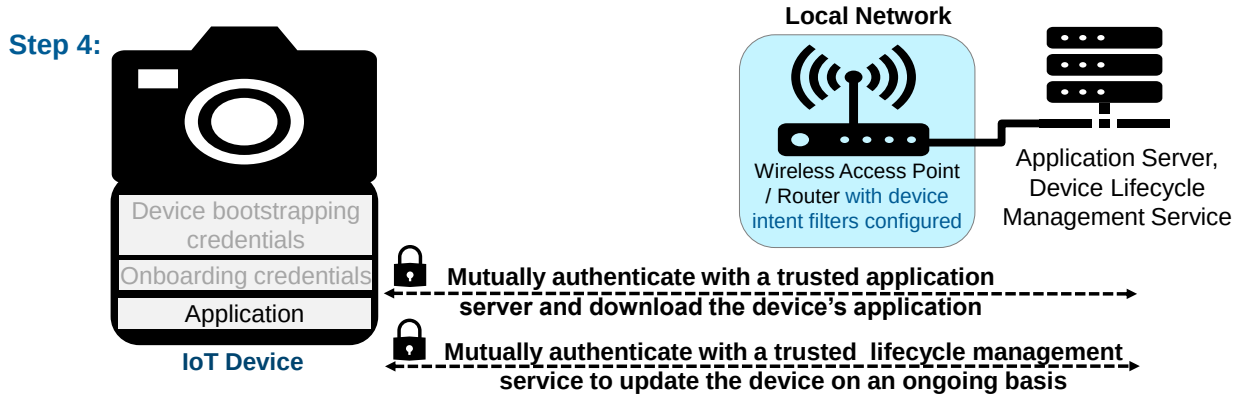

Figure 4 – Application-Layer Onboarding

The above depiction of the network-layer onboarding process is intentionally generic and product and protocol-agnostic. It may include steps that are not present in all onboarding solutions, and it may omit steps that are found in some onboarding solutions. In particular, it depicts an example of what this report defines as *trusted* network-layer onboarding.

Trusted network-layer onboarding is a network-layer onboarding process that:

- provides each device with unique network credentials,
- provides the device and the network an opportunity to mutually authenticate,
- is performed over an encrypted channel (to protect network credential confidentiality),
- does not provide anyone with access to the network credentials, and
- can be performed repeatedly throughout the device lifecycle (to enable network credentials to be replaced).

4. Onboarding Lifecycle Management

Lifecycle management refers to the operations that are performed to manufacture, configure, secure, use, update, and otherwise manage IoT devices and their credentials through all phases of the devices' existence. Ideally, all aspects of lifecycle management should be performed securely. This section depicts the lifecycle of a generic IoT device with a focus on the various aspects of the lifecycle that are relevant to the device's interaction with the network and, in particular, onboarding. The diagrams in this section and the definitions of the lifecycle phases they depict are informed by [7], [8], [9], and [10].

Although onboarding is only one (possibly recurring) phase in the device lifecycle, the onboarding mechanism may impact and be impacted by numerous other phases in the device lifecycle. It is important to understand how onboarding affects and is affected by the various phases of the device lifecycle to ensure that any onboarding solution being considered for use adequately integrates with and addresses all aspects of the device lifecycle.

Not all devices will experience all the phases and events depicted in this generic lifecycle, nor will they experience them in the same order. The specific phases and operations through which any device transitions depend on the purpose of the device, the context of its deployment use case, the onboarding solution it is equipped to support, and any specific circumstances that may arise.

At the highest level, the device lifecycle as we define it consists of two general portions: a pre-market portion and a post-market portion. While in its pre-market portion, the device is, among other things, manufactured and shipped. While in its post-market portion, the device is, among other things, installed, onboarded, and used; it cycles through periods of maintenance and operation and is ultimately decommissioned, at which point it may be either reinstalled on a different network for further use or considered to have reached the end-of-life phase.

4.1. Pre-Market

Figure 5 provides a detailed depiction of the first portion of the IoT device lifecycle: pre-market, with a focus on those aspects that are significant to the device's interaction with the network and, in particular, onboarding. In both Figure 5 and Figure 6, the phases or states through which a device may transition are depicted as rectangles, with some of the relevant activity that occurs within that phase or state listed as bullets. Subphases of larger phases are depicted as rectangles within larger rectangles. All phases and subphases are labelled with names. Actions or events that cause the transition from one phase or state to another are depicted as labelled arrows. At various points in our discussion of Figure 5 and Figure 6, we mention onboarding functional roles that are relevant to particular portions of the device lifecycle. These roles are described more fully in Section 6.

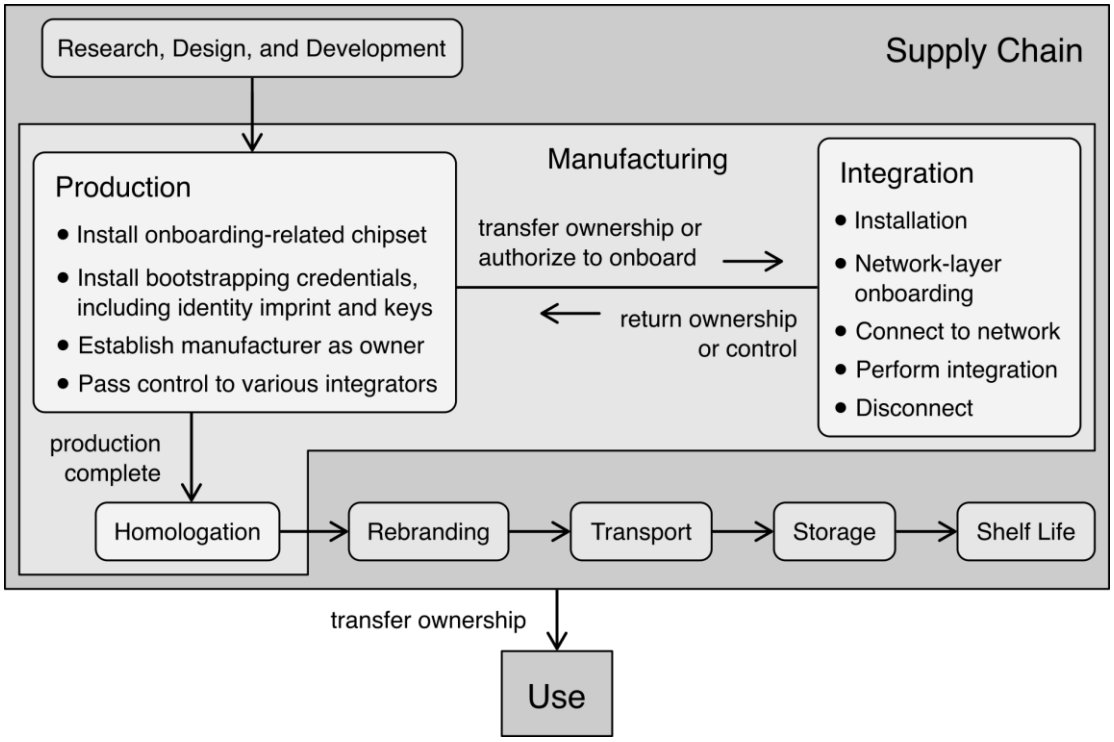

Figure 5 – Pre-Market Portion of the IoT Device Lifecycle from an Onboarding Perspective

As shown in Figure 5, the following phases and actions may be part of the device pre-market activities:

Research, Design, and Development Phase: This is the phase during which the device's onboarding (and other) requirements are defined and onboarding-related features are designed, tested, and refined. This phase may include trial production runs for review and improvement of potential onboarding-related capabilities. Device manufacturers can reference NIST IR 8259, *Foundational Cybersecurity Activities for IoT Device Manufacturers,* for more information about recommended pre-market activities.

Manufacturing Phase: This is the phase during which the device is produced, assembled, and approved. Although the device manufacturer is responsible for the overall manufacturing process, the manufacturer may rely on various integrators to create and install specific parts of the device. Hence, the manufacturing phase is comprised of a production subphase, an integration subphase, and a homologation subphase. The *device manufacturer and device system integrator* functional roles (See Section 6) are relevant to this phase of the device lifecycle:

Production Subphase: During production, the device manufacturer performs all parts of the manufacturing process that it will perform directly, which includes installing onboarding-related hardware, firmware, and software, including a security hardware module or hardware root of trust, random number generator, secure storage, and other components required to support trusted network layer onboarding. The device's bootstrapping credentials are also installed. Section 5.1.15.1.1 provides a more detailed list of steps performed during the production subphase.

The manufacturing phase may involve not just in-house device production but also integration with components supplied by various part manufacturers, including the installation of open-source or other software on the device. If the manufacturer relies on one or more other organizations to install and integrate parts or software on the device, the device may have to pass through a succession of device system integrators, and some or all of those device system integrators may need to connect the IoT device to their own networks for the short time necessary to install and integrate the desired component. If supporting an onboarding solution that provides proof-of-ownership verification, then during the production phase, the manufacturer will create a device information declaration that lists itself as the device owner and may also list some or all of the device's system integrators as authorized onboarders of the device, as needed. When the device is passed to an integrator, if the integrator is not listed as one of the device's authorized onboarders, then ownership of the device will have to be transferred to the integrator in order for the integrator to be able to onboard the device. When the device is passed to an integrator, it enters the integration subphase of the manufacturing phase.

> **Integration Subphase**: During the integration subphase, the device system integrator installs, onboards, and connects the device to its network. The device becomes operational only for undergoing the specific integration process required by that integrator. It is then disconnected from the integrator's network, and control (and possibly ownership) of the device is passed back to the manufacturer.
>
> The manufacturer may onboard the device to its network for further production or pass it to another device system integrator, which onboards the device to its network, and so on, until all system integration is complete and the device is ultimately transferred back to the manufacturer. Once the device has been returned to the manufacturer from its last integrator and production of the device is complete, the device enters the homologation subphase.
>
> Homologation Subphase: This is the subphase of manufacturing during which the device is approved and certified as being compliant with the set of relevant requirements that are specific to the purpose or industry for which the device is designed. Homologation is the last subphase through which the device passes before leaving the manufacturing phase.

Once manufacturing is complete, the device may enter rebranding, transport, storage, and shelf life phases.

Rebranding Phase: This phase may occur if a device is rebranded by a vendor other than the original manufacturer. If the device has a device information declaration or a device intent information file, the responsibility for maintaining these artifacts may be securely passed from the manufacturer to the vendor that has rebranded the device. In some cases, a value added reseller may install brand-specific applications onto the device as part of the rebranding process. Depending on policy, the authenticity of these applications may have to be attested to as a prerequisite to performing network-layer onboarding of the device. These applications may also need to be configured as part of the device's application-layer onboarding process, and any

additional communications that they require may have to be added to the device's device intent information profile.

Transport, Storage, and Shelf Life Phases: These phases occur after the device has been manufactured but before it is purchased and installed by its first post-market owner. Note that if the device bootstrapping credentials were to expire during any of these phases, trusted onboarding as we envision it would not be possible. This is why any certificates in the device bootstrapping credentials should not expire.

Transfer Ownership: The transition from pre-market to post-market is an event in the device lifecycle that occurs when the IoT device is sold or otherwise transferred to its first post-market owner, as shown in Figure 5. We will discuss the use portion of the device lifecycle in detail below. Before doing so, however, it is important to discuss the onboarding-related activities that may be required to accompany ownership transfers. As has already been mentioned, some onboarding solutions support a proof-of ownership mechanism that enables verification regarding which entity owns the device. They may also support an authorization to onboard mechanism that enables verification regarding a network's authorization to onboard a particular device. Both mechanisms are designed to enable a device to determine if the network that is trying to onboard it is authorized to do so. (In terms of the functional roles discussed in Section 6, if the network owner is the same is the device owner, then the network is implicitly authorized to onboard the device. If the network owner is a device authorized onboarder, the network is explicitly authorized to onboard the device.)

If an onboarding solution does not support either a proof-of-ownership or authorization-to-onboard mechanism, then no special onboarding-related activities are involved in the transfer of device control or ownership from one entity to another. On the other hand, if an onboarding solution does support a proof-of-ownership verification or authorization-to-onboard mechanism (or both), then the manufacturer of the device is responsible for performing (or having a proxy perform) certain activities so the device's owner and the list of device authorized onboarders can be tracked and verified as they change during the device lifecycle. For example, during the production phase, the device manufacturer must create a device declaration information that lists the manufacturer as the first owner of the device and/or specifies what other entities (if any) are authorized to onboard the device (i.e., what other entities are device authorized onboarders—see Section 66). When ownership of the device or authorization to onboard the device changes, the manufacturer is responsible for ensuring that the device declaration information is updated to reflect the new device owner and the new list of device authorized onboarders (if any). The device information declaration must be kept up to date because the device will only permit itself to be onboarded to a network that is either the owner of the device or an authorized onboarder of the device, as documented in the device information declaration.

To be useful, the device information declaration has to be trusted by the entities that are consulting it, (i.e., the IoT device and the onboarding component). Such trust could be established, for example, by having the current owner and/or the device manufacturer sign the device information declaration or by ensuring that the device information declaration is available from a widely trusted, well-known server. As ownership or other information within

the device information declaration changes, it needs to be updated and re-signed as appropriate.

In Figure 5 and Figure 6, these device information declaration update actions coincide with the events that are depicted by the arrows labelled “transfer ownership” or “authorize to onboard” (or both), and they occur at several transitions in the device lifecycle. In the pre-market portion, they occur on the arrows that transition back and forth between the production subphase to the integration subphase to denote that the manufacturer may either transfer ownership of the device back and forth between itself and a succession of system integrators or designate each system integrator to be an authorized onboarder of the device and merely pass control of the device back and forth between itself and the succession of system integrators.

The transfer ownership action also occurs on the transition from the pre-market portion to the post-market portion of the device lifecycle to denote that when the manufacturer of the device transfers ownership of the device to its first post-market owner, the device enters the post-market portion. In terms of the functional roles discussed in Section 66, the device purchaser will designate the device owner and any device authorized onboarders, if any, and the device manufacturer will update the device information declaration with this new information. Either the device owner or one of the device authorized onboarders must be the same as the network owner, i.e., the owner of the local network to which the device will be onboarded.

In Figure 6, the transfer ownership action occurs on the transition from decommissioning to installation to denote that when a device is re-sold for use by a new owner, it would need to be re-installed by that new owner on the new owner’s network. Prior to such an ownership transfer event, it would be advisable for the current owner to delete all information on the device except the device’s bootstrapping credentials. Again, the device purchaser will designate the device owner and determine any device authorized onboarders, if any, that will be inserted into the device information declaration, and one of these must be the same as the owner of the new local network to which the device will be onboarded.

4.2. Post-Market

Once ownership of the device is transitioned to its first post-market owner, the device leaves the pre-market portion and enters its post-market portion. Figure 6 shows a detailed depiction of the post-market portion of the device lifecycle, once again with a focus only on the device’s interactions with the network and those aspects that are significant to onboarding. The *device manager* functional role (See Section 66) is relevant at numerous points throughout this portion of the device lifecycle. Before the post-market portion can begin, the local *network administrator* is assumed to have performed certain functions in preparation for onboarding the IoT device, as discussed in Section 5.1.2).

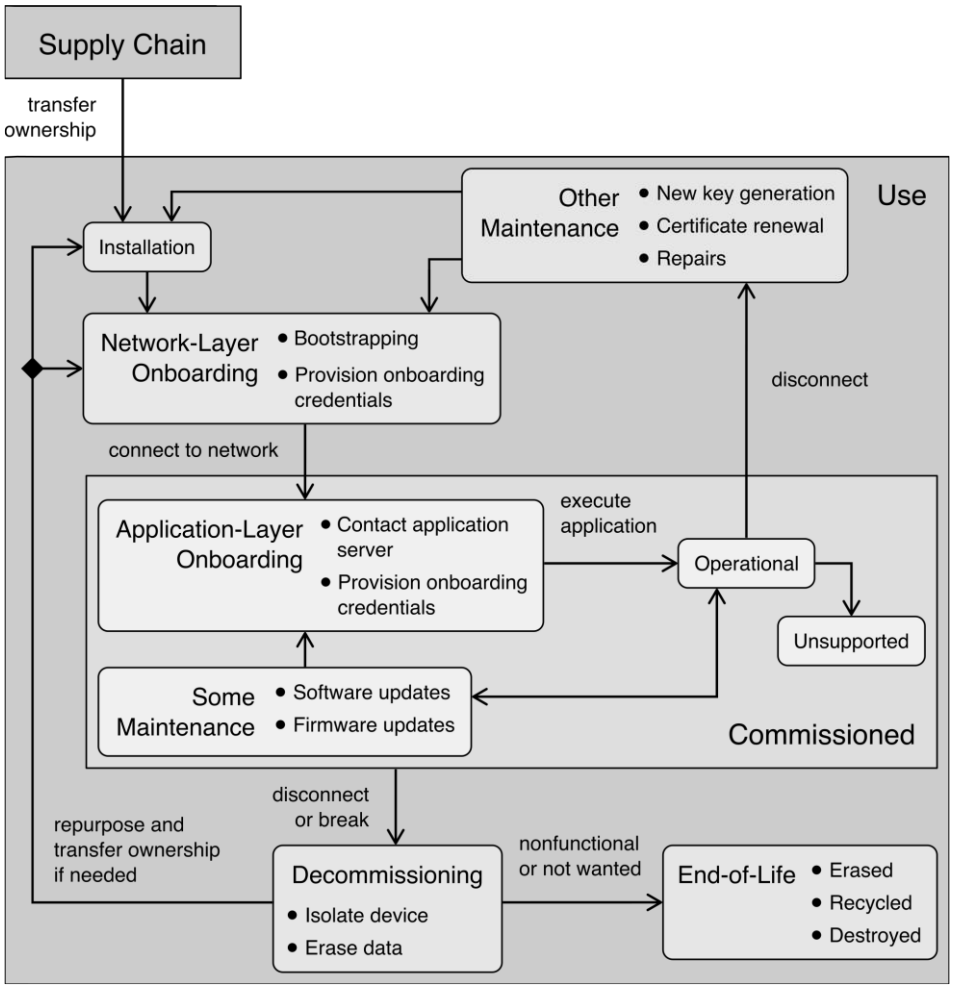

Figure 6 – Post-Market Phase of the Device Lifecycle from an Onboarding Perspective

Figure 6 depicts the phases that may comprise the device's post-market period. There is no one single path that every device follows through the post-market portion of its lifecycle, and many devices will cycle through periods of operation and maintenance as shown in the diagram. One path that a device may take as it transitions from installation to end of life might be as follows:

Installation Phase: This is when the device is physically placed into position, turned on, and, if it will have wired network access, physically connected to the network. If any buttons need to be pushed, antennae need adjustment, or the device needs to otherwise be prepared for onboarding, those operations are performed as part of the installation. (In some deployments, the installation phase may be performed after network-layer onboarding rather than prior to it. For example, in some deployments, an IoT device is required to be sealed underground or elsewhere and not accessed for many years. In these deployments, it would make sense to perform onboarding before installation, if possible, to ensure that onboarding was successful before sealing the device away.) The *device installer* functional role (See Section 66) is relevant to this phase of the device lifecycle.

Network-Layer Onboarding Phase: This is when network-layer onboarding, as previously described in Section 3, is performed. A more detailed list of steps performed during the

network-layer onboarding phase is provided in Section 5.2. In addition to being provisioned with unique network credentials, the device may be personalized with a local identifier by the network administrator (i.e., a device name that is meaningful to the network administrator and/or device manager). The *device onboarder* functional role (See Section 66) is relevant to this phase of the device lifecycle.

Connect to Network Action: The device uses its newly-provisioned credentials to connect securely to the network.

Commissioned Phase: Once the device has connected securely to the network, it enters the Commissioned phase. In this phase, the device is operational at the network layer, meaning that it can communicate securely with other devices on the network. The following subphases comprise the Commissioned phase:

- **Application-Layer Onboarding Sub-phase**: As previously described in Section 3, during this sub-phase, the device establishes a trusted association with one or more trusted application servers, controllers, or cloud services that will securely provision application-layer functionality to the device. Once this application-layer functionality is installed, the device is assumed to have everything it needs to function as intended and fulfill its purpose. In some cases, application-layer onboarding may have to be performed manually by an *application installer* (see Section 66) rather than being executed automatically upon network connection.
- **Execute Application Action**: The device begins executing its application, thereby transitioning into the operational phase.
- **Operational Sub-phase**: In this phase, the device's application is executing as intended; the device is performing its intended purpose, and it can be used by a *device user* and *an application user* (see Section 66). Ideally, the device stays in the operational state for most of its lifetime. Periodically, however, the device may need maintenance that requires it to leave the operational state. Alternatively, it could stop being supported, in which case it would enter the unsupported state.
- **Unsupported Sub-phase**: A device is considered to be in the unsupported phase if it is functional, but it is no longer supported by its manufacturer or by one or more of the manufacturer's integrators (either because the manufacturer or integrator has gone out of business or because the device has been deprecated). The device is still operating on the network, executing its application, and able to be used despite potentially having unpatched, known vulnerabilities and no longer being covered under the manufacturer support contract. In this phase the device may have reduced functionality. An unsupported device stays in the unsupported phase either until it is explicitly disconnected from the network, at which time it should be decommissioned, or until it

breaks, at which time it should be decommissioned and will reach the end-of-life phase by virtue of no longer being functional.

- **Some Maintenance Sub-phase**: A device is in this sub-phase when it is not operational and is undergoing maintenance that can be performed with minimal disruption. For example, many routine software or firmware patches or updates may be performed while the device continues to remain connected to the network in the commissioned state. Once this maintenance is complete, the device may be able to transition directly back to the operational state, or the device may only need to perform application-layer onboarding again to update its application before returning to the operational state. Either the *device manager* or the *application manager (*see Section 66), or both, may be responsible for performing or overseeing activities performed within this sub-phase.

Other Maintenance Phase: A device is in this phase when it is not operational and is undergoing maintenance that is so disruptive (e.g., replacement of security keys, certificate renewals, encryption library updates, some security patches or upgrades, and some physical repairs) that it requires the device to be disconnected from the network and possibly even uninstalled, thereby causing the device to exit the commissioned state. This maintenance may be needed as a result of a network compromise, regular credential replacement, or device malfunction. After this type of maintenance is complete, the device may need to be reinstalled, and it will have to go through network-layer onboarding again, be reconnected to the network, and go through application-layer onboarding before returning to the operational phase. The *device manager* (see Section 66) typically performs or oversees activities during this phase.

Disconnect Action: The device manager does this to remove the device from the network so that the device can be either maintained or decommissioned. Disconnection removes the device from the commissioned state.

Break Event: This results in the device no longer being functional. A device that becomes nonfunctional and is beyond repair needs to be decommissioned and will reach its end-of-life phase. Breaking removes the device from the commissioned state.

Decommissioning Phase: During this phase, the *device manager* and *application manager* perform the operations needed to ensure that the device permanently stops performing its intended function on the local network, i.e., they decommission the device. A *device manager* may decide to decommission a device if it stops functioning (e.g., breaks) and cannot be repaired, or when it is determined that the device should no longer be used to perform its intended function on the network (perhaps due to becoming out-of-date or obsolete or losing software support). A device that has been decommissioned may be replaced on the network by a newer-model device. The decommissioning phase includes disconnecting and isolating the device so that it can no longer affect the network. It also involves erasing all sensitive data from the device as required by organizational policy, including all onboarding information and all application-related data (e.g., logs and user data that have been collected), so that the only information that is left on the device is its original bootstrapping credentials. A factory reset

may be required to ensure the removal of the sensitive information. After a device has been decommissioned, it may either reach its end-of-life phase or be repurposed.

It should be noted that there is a distinction between a device being decommissioned and reaching the end-of-life phase in terms of its network connectivity and a device being decommissioned and reaching end-of-life in terms of its real-world functionality. Figure 6 depicts only the device's lifecycle in terms of its network connectivity. A device that is decommissioned from the network may continue to be used while disconnected. For example, a connected washing machine may reach the end of its software support, leading its owners to disconnect it from the network and decommission it (in terms of network connectivity) so that it will not be vulnerable to a network-based attack due to unpatched software. This decommissioned device may still function well as a washing machine and may continue to be used to wash clothes. In terms of the decommissioned device's interaction with the network, however, it has reached its end-of-life phase because it will not be used to connect to a network again.

End-of-Life Phase: This is the phase that a decommissioned device enters if:

- The device is nonfunctional and cannot be repaired.
- The device is functional but is no longer deemed useful for any purpose, not even on a secondary market.
- The device will not be connected to a network again; it no longer needs those components it uses to interact with the network.

Upon reaching the end-of-life phase, a device should have all its sensitive data removed (to the extent possible), and it (or at least the components it uses to interact with the network) should be destroyed. Some of its parts (precious metals, batteries) may be recycled for use elsewhere.

Repurpose Action: A *device manager* does this on a decommissioned device that is still usable in terms of interacting with a network. Repurposing means putting a device to a different use. It may be put to a different use by its current owner, or it may be sold on a secondary market and used by a new owner. When the device is repurposed, it essentially loops back to an earlier phase and begins proceeding through a new path in its lifecycle.

- If the device is remaining with its current owner but needs to be onboarded to a new network, it will loop back to the installation phase and then proceed through its lifecycle.
- If the device is remaining with its current owner and will be used on the same network but in a different role, the device will loop back directly to the network-layer onboarding phase and then proceed through its lifecycle.

If the device will be sold to a new owner, the current owner will execute an ownership transfer event before repurposing the device (e.g., by ensuring that the device information declaration is updated to reflect the change of ownership), assuming that the onboarding solution supports this feature. After ownership has been transferred to the new owner, the device will loop back

to the installation phase and then proceed through its lifecycle on a different network—one that belongs to or is authorized by its new owner. After being repurposed, the device will proceed through various lifecycle phases as it did before—perhaps looping through the operational and various maintenance phases for some time, perhaps being repurposed one or more times—before ultimately reaching the end-of-life phase.

5. Onboarding Process Steps

This section describes the steps that are performed to make an IoT device operational. It is intentionally general and describes each step in a generic, solution-neutral manner. These steps can be viewed in four general stages: pre-onboarding, followed by the three stages that were depicted in Figure 11: network-layer onboarding, secure connection establishment, and application-layer onboarding.

The steps described in this section assume an onboarding solution that supports proof-of-ownership and integrates with the optional mechanisms (attestation, device intent, and application-layer onboarding) depicted in blue in the provisioning step of Figure 2. Not all onboarding solutions will support these mechanisms, and any given onboarding mechanism may support any combination of them. We include them here because they may be relevant to some onboarding solutions.

5.1. Pre-Onboarding

The pre-onboarding stage of the IoT device occurs before the device is associated with any given network. The goal of the pre-onboarding stage is to equip both the device and the network with bootstrapping credentials and other capabilities they may need to perform network-layer onboarding. Some pre-onboarding activities are performed on the device during the manufacturing process. These may be performed by the device manufacturer or a trusted party, such as a device system integrator. Other pre-onboarding activities are performed at the network to which the device will be onboarded. These are performed by the network administrator, as authorized by the network owner.

5.1.1. Pre-Onboarding Activities Performed during Manufacturing

The activities of the pre-onboarding stage that occur as part of the manufacturing process consist of six general steps, as depicted in Figure 77:

- Step 1 is the installation of all onboarding-related chipsets, hardware, and software, and installation of the device's bootstrapping credentials. This is the only pre-onboarding step that is mandatory. To prevent the secret in the bootstrapping credentials from being disclosed, it should be safeguarded in a hardware-backed secure storage element that prevents it from being easily extracted or modified without detection [11].
- Step 2 involves the manufacturer generating and possibly signing a device information declaration. This step is only performed if the onboarding solution supports a proof-of-ownership verification capability.
- In step 3, the manufacturer creates the information that is needed to verify the device's attestation claims and makes it available. For example, the manufacturer could make this information available on a verification service that is run by either the manufacturer or a third party, as shown in Figure 77.

- In step 4, the manufacturer creates the device's device intent information and makes it available as required by the device intent mechanism being used. For example, the manufacturer may post the device intent information to a device intent file server that is available on the internet. If the device intent information ever changes, the manufacturer will be responsible for updating this server so that it makes the latest version of the device intent information available. The manufacturer also adds a device intent information URL that identifies the location of this file to the device's device bootstrapping credentials. Step 4 is only performed if the device is device-intent-capable and the onboarding solution supports conveyance of device intent information. If the device is intended to be device-intent-capable even though the onboarding solution does not support the conveyance of device intent information, the manufacturer will still create and post the device's device intent file. However, a different mechanism (e.g., the Dynamic Host Configuration Protocol [DHCP header]) will be used to enable the device to convey its device intent information to the network. If the device is not intended to be device-intent-capable, no part of step 4 will be performed.
- In step 5, the manufacturer posts the device's application to an application server so that it will be available for the device to download after the device is connected to the internet. As the manufacturer continues to update the application, the manufacturer is responsible for updating this server so that it makes the latest version of the application available. The manufacturer also installs any credentials and information (e.g., the device's public key) that the application server might need to enable it to establish a secure association with the device. The manufacturer also adds an application server URL that identifies the location of this application to the device's bootstrapping credentials, along with any information that the device might need to be able to trust the device (e.g., the application server's public key). Step 5 is only performed if the device has an application server and the onboarding solution supports application-layer onboarding. If the device will have an application server, but the onboarding solution does not support application-layer onboarding, the manufacturer will still post the device's application. However, a different, possibly manual, mechanism may be used to download the latest versions of the application to the device. If the device does not have an application server, no part of step 5 will be performed.
- Finally, in step 6, the device information declaration is updated, possibly signed, and transmitted to the next owner once the next owner's identity is known. The manufacturer may also make this device information declaration available so that it can later be accessed by the device and/or network in real-time, on-demand, during the bootstrapping phase of network-layer onboarding. As discussed with respect to the

Transfer Ownership event in Section **Error! Reference source not found.**, the manufacturer is responsible for ensuring that the device information declaration is kept up to date on an ongoing basis. As ownership of the device or authorization to onboard the device changes, the manufacturer must ensure that the device declaration information is modified to reflect the new device owner and the new list of *device authorized onboarders* (if any). Step 6 is only performed if the onboarding solution supports the proof-of-ownership and/or onboard-only-to-authorized-networks mechanisms.

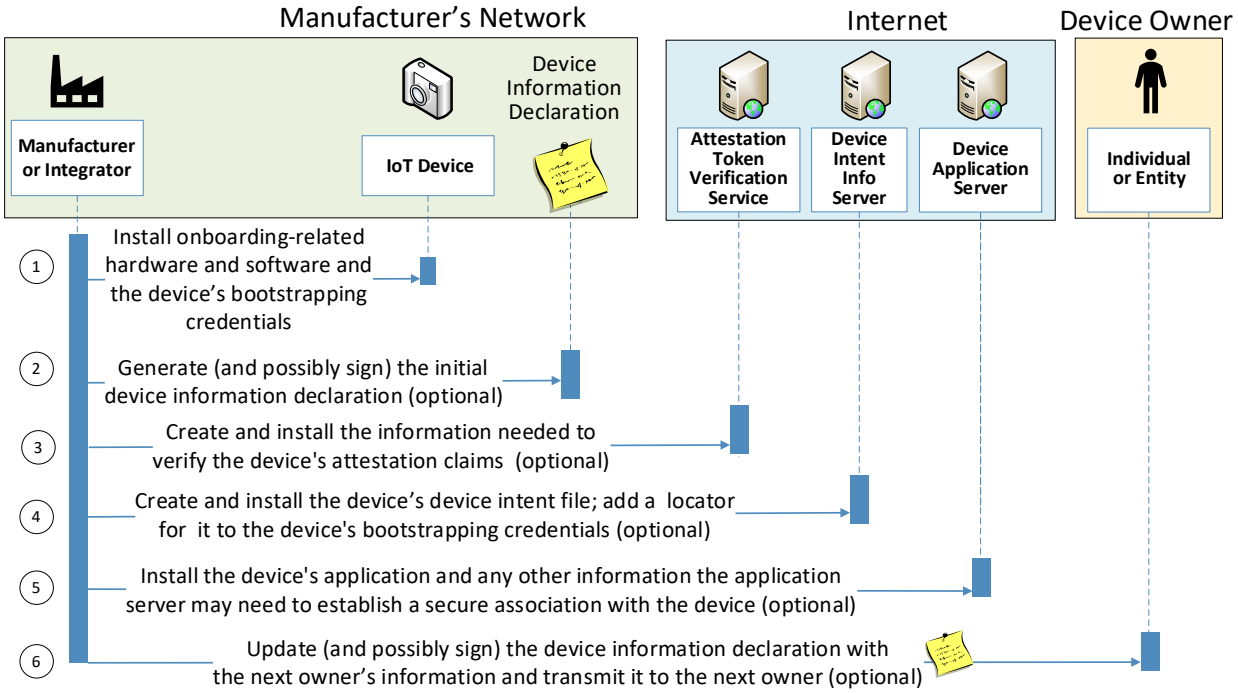

Figure 77 – Pre-Onboarding Activities Performed during Manufacturing

5.1.2. Pre-Onboarding Activities Performed at the Onboarding Network

In addition to the pre-onboarding activities that are performed during manufacturing, some pre-onboarding activities may also need to be performed at the network to which the device will be onboarded. Figure 8 illustrates three of these activities, which would be performed by the network administrator (see Section 6):

- The first activity, installing the network bootstrapping credentials on the network onboarding component, is required if the onboarding solution requires the network to authenticate to the device, or if the onboarding process supports a proof-of-ownership or onboard-only-to-authorized-networks mechanism. If the onboarding solution supports a proof-of-ownership or onboard-only-to-authorized-networks mechanism, the *network owner* (see Section 66) must be identified in the network bootstrapping

credentials. Those aspects of the network bootstrapping credentials that need to be kept secret should be safeguarded in a secure storage component to which the network onboarding component has access [11].

- The second activity, installing the device information declaration on the network onboarding component, is performed only if the onboarding solution supports a proof-of-ownership mechanism and there is not assumed to be a mechanism for retrieving the device information declaration in real-time during the bootstrapping process. (During bootstrapping, the device requires access to the device information declaration in order to be able to determine whether the network is authorized to onboard it. The device information declaration could have been signed and sent from the manufacturer to the network owner upon purchase of the IoT device (as depicted in Figure 77) and installed on the network onboarding component by the network administrator before onboarding begins (as depicted in Figure 8), or it could be retrieved in real time during the bootstrapping process. Real-time retrieval requires internet access during the onboarding process but ensures receipt of the most up-to-date device ownership information.
- The third activity, configuring the device authorization policy at the authorization service, is performed only if the network has an authorization service. The authorization policy should be configured to enforce local network policy regarding the device (e.g., what access privileges it will have, what constraints must it meet). With respect to onboarding in particular, the policy may, for example, require the device to include specific claims in its attestation token before the device is permitted to onboard to the network, or it may stipulate that the device be assigned to a specific network segment based on its type, location, or other characteristics.

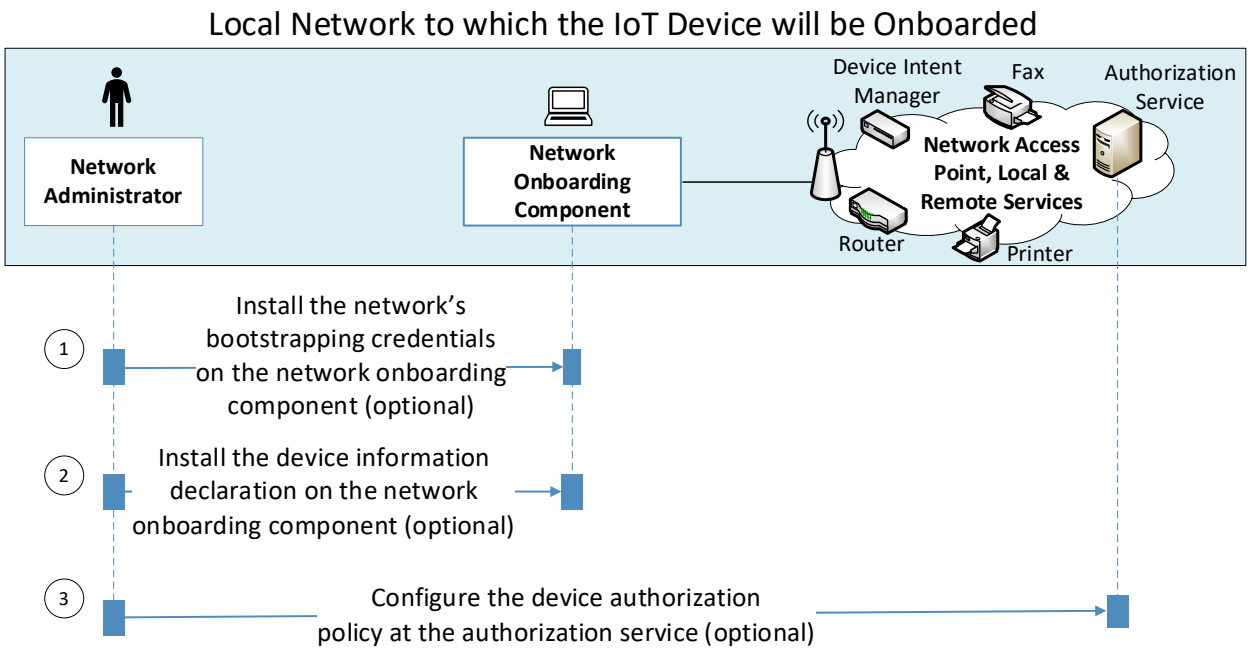

Figure 8 – Pre-Onboarding Activities Performed at the Onboarding Network

5.2. Network-Layer Onboarding

Once the IoT device has been purchased and the pre-onboarding activities have been completed at the onboarding network, the device owner provides the device to an installer for installation. Once installed, a *device onboarder* (see Section 6) initiates the network-layer onboarding process.

Figure 999 depicts a more detailed, yet still generic, version of the network-layer onboarding process than is depicted in Figure 2, and it includes steps related to additional security mechanisms, such as attestation, device intent, and application-layer onboarding, that are assumed to be integrated with the network-layer onboarding process. Not all steps depicted in the figure necessarily occur each time a device is being onboarded. The steps that are included depend on the capabilities of the particular onboarding solution being used, the capabilities of the device being onboarded, and the local security policy. For example, some policies may require only one-way rather than mutual authentication between the device and the onboarding component; some IoT devices may not support device attestation, and so would not send a device attestation token to the network; some devices may not support device intent, in which case the onboarding solution would not convey device intent information to the network; and some devices may not support proof-of-ownership, and so would not have a device information declaration. In these cases, the corresponding steps would be omitted. Also, the steps may not necessarily occur in the exact order depicted.

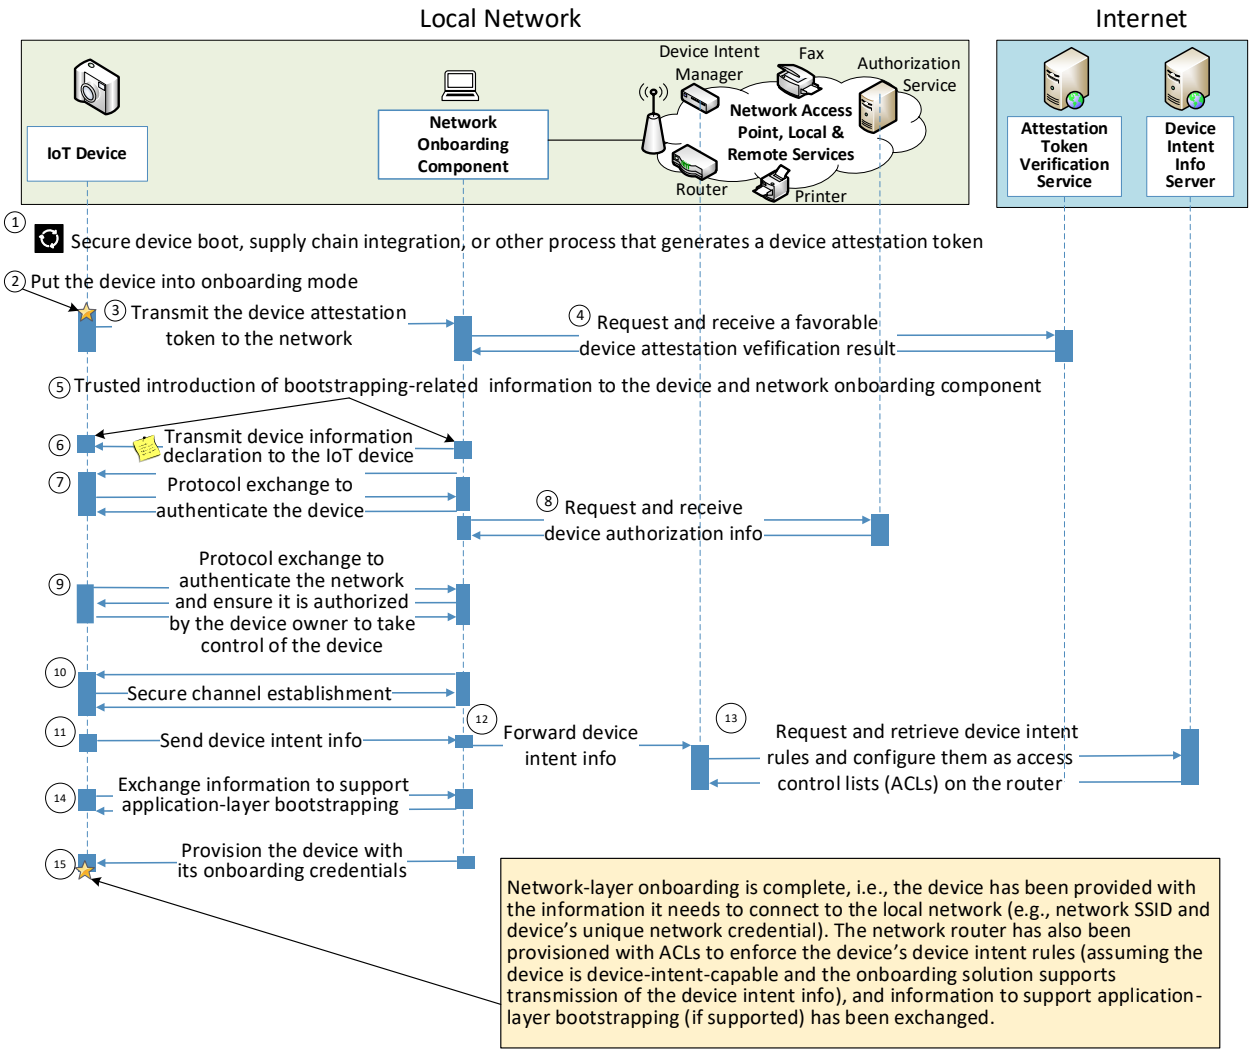

Figure 99 – Network-Layer Onboarding Process Steps

- Step 1 is for the *device onboarder* to boot the device securely, so that it generates an attestation token that makes claims about its authenticity and integrity, and perhaps other characteristics; alternatively, if the device is not capable of a secure boot, step 1 may involve generation of a device attestation token based on supply chain tools. (Steps 1, 3, and 4 are performed only if the device supports attestation.)
- Step 2 is for the *device onboarder* to put the device in onboarding mode, which means that the device begins transmitting and/or listening for onboarding protocol messages. Because the device is not yet securely connected to the network, it is permitted limited network communications capabilities to perform the network onboarding process, and it interacts only with the network onboarding component. For example, a device may communicate with the network onboarding component at Layer 2, i.e., the data link layer, or at Layer 3, i.e., the network layer, of the Open Systems Interconnection model

during onboarding, but it will not receive a routable IP address until onboarding is completed. The device interacts with the network onboarding component using the onboarding protocol, which should be well-defined. The details of the onboarding protocol exchanges are specific to the particular onboarding protocol used; in this section, we describe those exchanges in a generic manner.

- Step 3 is for the device to transmit the attestation token that was generated in step 1 to the network onboarding component.
- In step 4, the network onboarding component forwards the attestation token to an attestation token verification service that evaluates the token and returns a response indicating whether or not the claims that the device has made have been determined to be valid. If the device is not determined to be trustworthy, the network onboarding component terminates the onboarding process. If the device is determined to be trustworthy, the onboarding process continues.
- Step 5 involves the trusted introduction of information to the device and the network onboarding component. For example, the network onboarding component is provided with the information it needs to authenticate the device (e.g., the device's public key), and the device is provided with the information it needs to authenticate the network (e.g., the network's public key). Note that the information that is introduced does not include either the device's or the network's secret. This trusted introduction can consist of a variety of different bootstrapping mechanisms, some of which may be considered more trustworthy than others. For example, some solutions may depend upon a trusted human (i.e., a *trusted device onboarder*) to perform the introduction of bootstrapping-related information to the device and/or the network; other solutions may be able to perform this trusted introduction in a zero-touch manner, based on device and network hardware roots of trust that include an X.509 certificate and access to a public key infrastructure (PKI). The particular trusted introduction mechanism that is used to bootstrap any given onboarding solution will play a significant role in determining the overall level of security assurance that the onboarding solution provides.
- In step 6, the device information declaration is sent to the device. This device information declaration could be sent to the device from the network onboarding component, assuming it had already been obtained from the manufacturer (as depicted in Figure 77) and installed on the network onboarding component by the network administrator (as depicted in Figure 8). Alternatively, the device could retrieve it from the manufacturer in real-time.
- Step 7 is authentication of the device. The device presents its identity to the network onboarding component, which authenticates the device, most likely by using a challenge

response protocol. The network onboarding component uses the device's public key to verify that the device has the identity that it claims to have by virtue of being in possession of the device's corresponding private key (bootstrapping secret).

- In step 8, the network onboarding component retrieves the device's authorization policy from the authorization service so that it can ensure that the device is constrained to doing only what it is authorized to do, if applicable.
- Step 9 is authentication of the network. The network onboarding component presents the network's identity to the IoT device, which authenticates the network, verifying that the network has the identity that it claims to have by virtue of being in possession of the network's private key (bootstrapping secret). The IoT device also uses the *device owner* or *authorized device onboarder* information in the device information declaration (which it received in step 6) to ensure that the network is authorized to onboard the device by verifying that the *network owner* (as identified in the network onboarding component's bootstrapping credentials) is the same as either the *device owner* or one of the *device authorized onboarders*. The details of the protocol exchanges that need to occur to perform the authentication processes in steps 7 and 9 are specific to the onboarding protocol used.
- Step 10 is secure channel establishment. The device and the network onboarding component establish a shared secret key to encrypt their subsequent exchanges. This encrypted connection, which, preferably, has a unique one-to-one binding between the device and the onboarding component, ensures the confidentiality of the device's onboarding credentials while they are in transit between the network onboarding component and the device in step 15. At this point, the bootstrapping process is complete.
- In step 11, the device sends its device intent information locator, which is specified in its bootstrapping credentials, to the network onboarding component over the secure channel. This enables the network to learn the device's intent so the network can enforce appropriate device communications.
- In step 12, the network onboarding component forwards the device intent locator to the device intent manager.
- In step 13, the device intent manager receives the device intent information locator from the network onboarding component, retrieves the device intent information file, and configures the network router to enforce the device intent access control list (ACL) rules for the device.

- In step 14, the device and the network onboarding component exchange application-layer bootstrapping-related information that will be used later to support application-layer onboarding. This application-layer bootstrapping-related information will help enable the device and one or more application-layer components to authenticate each other and establish a secure association.
- Finally, in step 15, the network onboarding component provisions the device with its onboarding credentials, i.e., the information it needs to connect to the local network, as well as any other information that it may need to support application-layer onboarding.

5.3. Secure Connection Establishment

After the device has been provisioned with its onboarding credentials, the network-layer onboarding process is complete, and the device is no longer in onboarding mode. It is no longer listening for or generating onboarding protocol messages, and the network onboarding component and the network-layer onboarding protocol are no longer active or used. The device is ready to connect directly to the network (rather than to the network onboarding component) by presenting its newly provisioned network-specific credentials to establish a secure network association. In some cases, the device may automatically connect to the network immediately after being provisioned with its onboarding credentials.

In some onboarding solutions, immediately after the device successfully connects to the network, it may be desirable for the device to report this fact back to the network onboarding component as a diagnostic feature so the network onboarding component can be aware of the status of the device. In this case, there would be a brief period during which the device would communicate with the network onboarding component after the device has connected to the network.

If the device needs to be provisioned with different onboarding credentials due to events that affect its current credentials (e.g., credential expiration, security updates, security breach, cryptographic library bug, or certificate renewals) or to the device being repurposed or resold, then the device's current onboarding credentials would be deleted, and the network-layer onboarding process would be repeated, starting with step 1, to provision the device with new credentials.

5.4. Application-Layer Onboarding

If the device has been provided with application-layer onboarding information as part of its onboarding credentials, then after connecting to the network, this information could direct the device to a particular controller, application server, or cloud service that, when contacted by the device, will securely provision application-layer functionality to the device. This application-layer functionality can include authentication/authorization with a cloud service, application

provisioning, subscription to firmware updates, device ownership assignment, and device lifecycle management.

6. Functional Roles

This section describes proposed personnel roles for accomplishing onboarding. These roles may be filled by the same or different people or entities, depending on the use case. (In a home setting, for example, many of these roles would fall to the device owner.) Also, the persons or entities filling these roles may change as a device moves through its lifecycle (e.g., its owner and its authorized onboarders may change). The roles and responsibilities are described below and are grouped by the general stages defined in Section 4.

Throughout the device lifecycle, trust needs to be established and maintained between the device and the entities playing these various roles. For example, a medical device might need to trust a network owned by one entity but also connect to and trust cloud servers owned by another entity. Also, as the device moves through its lifecycle, some of the above human roles move in and out of relevance to the device.

Pre-Onboarding by the Manufacturer

- The **device manufacturer** creates the device, installs the device's bootstrapping credentials, and is the first owner of the device. The manufacturer knows the intent of the device but is not able to imprint anything on the device that is unique to the device's specific network deployment. The manufacturer is responsible for creating and possibly signing the device information declaration, and also for creating and signing the device intent information and making it available. The manufacturer (or a trusted third party) also has related ongoing responsibilities that extend well beyond the manufacturing process. The manufacturer is responsible for ensuring that the device information declaration remains up-to-date as device ownership and device authorized onboarders change, that the device intent information remains up-to-date and available as the device's communications intent changes, and that device firmware and software libraries are updated as needed.
- The **device system integrator** is responsible for integrating a subcomponent (including software) of the device onto the device during the manufacturing process.

Pre-Onboarding at the Local Network

- The **network owner** is the individual or entity that owns the network on which the IoT device is deployed. The network owner must be identified as part of the network onboarding component's bootstrapping credentials. In the consumer use case, the network owner may be the same as the device user (i.e., the consumer), but in the enterprise use case, the network owner is typically a company. In some deployments, the device owner may be different from the network owner. For example, in a connected grid deployment, the connected grid of IoT sensors and other devices may be owned by one company, but the actual network on which the connected grid is running

may be owned by a different organization. In onboarding solutions that support proof-of-ownership verification and mechanisms to grant authorization to onboard, where the device owner is not the same as the network owner, the network owner needs to be a *device authorized onboarder* (see below) of the device.

- The **network administrator** is the individual or entity that manages the network and configures, updates, maintains, and monitors networking-specific components, including the network onboarding component (but not necessarily those of the network's IoT devices). The network administrator expresses its wishes through policy and enforces them via mechanisms such as the authorization service. The network administrator is determined by the network owner.

Network-Layer Onboarding

- The **device owner** is the only individual or entity authorized to:
    - onboard and use the IoT device,
    - grant another individual or entity the authority to onboard and use the IoT device, and
    - transfer ownership of the IoT device to another individual or entity

  An IoT device may have only one owner at any given time. If an onboarding solution supports a proof-of-ownership mechanism, the device owner will be recorded in the device information declaration. The device's owner may change at various stages in the device's life. The device owner is also the entity that has the authority to determine the device's installer, onboarder, manager, and users as well as the application's owner, installer, manager, and users.
- The **device authorized onboarder** is an individual or entity that is also a *network owner* and that is authorized to onboard a given device to its network. This authorization comes from the device's owner, by virtue of the fact that the owner has listed the device authorized onboarder in the device information declaration; and this authorization may also be revoked by the device's owner by removing the device authorized onboarder from the device information declaration. The device authorized onboarder does not have the authority to designate any other entity as a device authorized onboarder; only the device owner is able to do this. A device may have multiple device authorized onboarders at any given time. If a device authorized onboarder needs to delegate onboarding ability to another party, it should request that the device owner add that party to the device's list of authorized onboarders.
- The **device purchaser** is the individual or entity that pays for or, in some cases, leases the IoT device. The device purchaser designates what individual or entity will be granted

ownership of the device by the manufacturer when the device is acquired. The device purchaser is not necessarily the same as the device owner, onboarder, manager, or user.

- The **device installer** is the individual or entity (e.g., the IT team) that places the device at its deployment location, connects any network or other cables to it, and may turn it on.
- The **device onboarder** is the individual who performs device onboarding. It is important to understand the distinction between the *device authorized onboarder* (described above) and the *device onboarder*. A *device authorized onboarder* is used to determine whether a given network is authorized to onboard a device, whereas a *device onboarder* is the person who performs the onboarding (and who may or may not be a network owner). A *device authorized onboarder* is listed in the device information declaration, and if it matches the *network owner (e.g., Doe Family)*, as listed in the network bootstrapping credentials, then the network is authorized to onboard the device. This authorization holds true regardless of who (e.g., John or Jane Doe) the device onboarder is. Depending on the capabilities of the device being onboarded and the capabilities supported by the onboarding solution, the device onboarder (e.g., John or Jane Doe) may or may not need to be trusted to some extent. If the onboarding solution does not rely on the device onboarder to provide any input into the onboarding process, there is no requirement that the device onboarder be trusted. Such solutions are most appropriate for enterprise settings. In other onboarding solutions, it may be desirable to make a tradeoff between security and convenience by permitting a device onboarder to play a limited role, such as determining to which network the IoT device should be onboarded. In this case, the device onboarder must be trusted to select the correct network, but beyond that selection, the onboarding solution does not require any additional trust to be put in the device onboarder. In onboarding solutions in which the device onboarder must be trusted, the device onboarder's role should ideally be as limited as possible so as to minimize the amount of trust that must be placed in this individual. For example, the device's bootstrapping and onboarding credentials should always be kept confidential and should never be disclosed to a device onboarder, even a trusted one.
- The **device manager** is the individual or entity responsible for managing the device. The device manager performs or oversees device software and firmware updates and configuration, as well as oversees all other device repairs and maintenance. When it is time for the device to be decommissioned, the device manager disconnects and isolates the device and erases all sensitive data, possibly performing a factory reset. When the device reaches its end-of-life phase, the device manager removes all data from and destroys the device, possibly selecting certain parts for recycling. When the device is to

be repurposed, the device manager transfers control of the device to its new authorized onboarder or new owner (as directed by the device's current owner).

- The **device user** is the individual or entity that uses the IoT device. From the viewpoint of the device and the network, the user is represented by their credentials.
- The **service provider** is the entity that operates the network that traffic transits to be sent to and from the internet from the IoT device's local network.

Application-Layer Onboarding

- The **application owner** is the individual or entity authorized to install, manage, and use a specific application on the IoT device. The application owner can grant others the authority to install, manage, and use the application. The application owner may be different from the network owner and from the device owner. For example, a consumer might have a solar panel set up on their home's roof. The solar panel is an IoT device that may be owned by either the consumer or the solar energy company. The solar panel is running a solar-energy-related application. The solar energy company owns the application, but the consumer owns the local network over which the solar energy application will send data back and forth to the cloud.
- The **application installer** is the individual or entity (e.g., the operational technology team) that onboards and installs the application to the IoT device. In some IoT devices, application installation may occur automatically during the application-layer onboarding process, based on the application-layer bootstrapping credentials that were included as part of the device's onboarding credentials.
- The **application manager** is the individual or entity responsible for managing the application. The application manager oversees application installation, initiates execution of the device's application, and helps manage the application by overseeing periodic application software updates. In addition, when the device is decommissioned, the application manager ensures that all application-specific sensitive data such as passwords, keys, logs, and user data that has been collected is erased.
- The **application user** is the individual or entity that uses the application on the IoT device to cause the device to perform its intended function. From the viewpoint of the application, the user is represented by their credentials.

7. Onboarding as a Foundation for Ongoing Device Security

Trusted network-layer onboarding can provide security benefits that extend well beyond the process of securely provisioning devices with their network credentials. It has the potential to integrate with and enhance additional security capabilities that provide ongoing protection of IoT devices for as long as they remain connected to the network.

Potential capabilities include, for example:

- The traffic filters that were specified by the device intent information are enforced to ensure that communications to and from the device are restricted to only those that are required. Local network policy can also be applied in addition to the device intent-specified policy, if desired.
- The device's firmware, software, and configuration are updated and patched by the trusted lifecycle management service as needed to address vulnerabilities.
- The device and its trusted lifecycle management service periodically perform mutual attestation to renew their confidence in each other's trustworthiness.
- The asset management system periodically cross-checks its discovered devices with the onboarded IoT devices to ensure there are no discrepancies. The asset management system also monitors and remediates the devices' software and configurations to identify known vulnerabilities and ensure compliance with policy expectations. This system can also verify whether a device is still supported by its manufacturer.
- Device profiling and behavior analysis are performed to help discover unexpected devices or anomalous behavior that may indicate the presence of unauthorized or compromised devices.
- The device can be reassigned to a particular network segment to restrict its ability to communicate with other local devices, for example, based on information received as a result of ongoing mutual attestation, asset management, threat intelligence, or device profiling and behavior analysis. The device can be dynamically reassigned to another segment, (e.g., for the purpose of quarantining the device if its trustworthiness comes into question).

By integrating with these or similar protections that build upon each other to comprehensively protect IoT devices, beginning with initial device boot-up and extending through the entire period during which the device remains connected, network-layer onboarding can serve as a lynchpin of both IoT device and network security.

References

[1] Kaiser Associates, Inc. (2017) IoT Onboarding: A Device Manufacturer's Perspective. Kaiser Associates IoT Onboarding for Device Manufacturers White Paper

[2] National Institute of Standards and Technology (2025) Trusted Internet of Things (IoT) Device Network-Layer Onboarding and Lifecycle Management. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP) 1800-36. https://doi.org/10.6028/NIST.SP.1800-36

[3] U.S. Department of Homeland Security, Cybersecurity and Infrastructure Security Agency (January 23, 2013) Wi-Fi Protected Setup (WPS) Vulnerable to Brute-Force Attack, Alert (TA12-006A). Available at https://www.us-cert.gov/ncas/alerts/TA12-006A

[4] Thakore D, (April 9, 2019) Micronets Deep Dive, National Cybersecurity Center of Excellence, Mitigating IoT-Based DDoS meeting, Rockville, MD, unpublished.

[5] Intel Corporation Product Brief (2019) Intel Secure Device Onboard. (Intel). h https://www.intel.com/content/dam/www/public/us/en/documents/IoT/sdo-product-brief.pdf

[6] Pandey AK (2019) AutoAdd—Automatic Bootstrapping of IoT Devices. (Internet Engineering Task Force [IETF]). Automatic Bootstrapping of IoT Devices

[7] Thakore D (November 4, 2019) IoT Device Onboarding & Lifecycle Management presentation, slide 9. (IoT Device Onboarding & Lifecycle Management meeting, Washington, DC, unpublished

[8] Vermillard J (2015) Bootstrapping device security with Lightweight M2M. Bootstrapping device security with | by Julien Vermillard

[9] Garcia-Morchon O, Kumar S, Sethi M (2019) Internet of Things (IoT) Security: State of the Art and Challenges. (Internet Research Task Force [IRTF]), IRTF RFC 8576. IoT Security: State of the Art and Challenges

[10] Atkinson S (December 7, 2017) IoT and connected device lifecycle management. (CIO). IoT and connected device lifecycle management | CIO

[11] Abhishek A (November 4, 2019) IoT Device Onboarding & Lifecycle Management. Cisco Systems IoT Device Onboarding & Lifecycle Management meeting, Washington, DC, unpublished

Appendix A. List of Symbols, Abbreviations, and Acronyms

Selected acronyms and abbreviations used in this paper are defined below.

5G
5th Generation

AAA
Authentication, Authorization, and Accounting

ACL
Access Control List

CDI
Compound Device Identifier

CRL
Certificate Revocation List

DAA
Direct Anonymous Attestation

DICE
Device Identifier Composition Engine

DPP
Device Provisioning Protocol

EAT
Entity Attestation Token

EPID
Enhanced Privacy ID

eSIM
embedded subscriber identity module

ETSI
European Telecommunications Standards Institute

FOIA
Freedom of Information Act

GPS
Global Positioning System

GSMA
Global System for Mobile Communications Association

ID
Identifier

IDE
Integrated Development Environment

IEC
International Electrotechnical Commission

IETF
Internet Engineering Task Force

IoT
Internet of Things

IP
Internet Protocol

IRTF
Internet Research Task Force

ISO
International Organization for Standardization

ISP
Internet Service Provider

IT
Information Technology

ITL
Information Technology Laboratory

JSON
JavaScript Object Notation

MAC
Media Access Control

MUD
Manufacturer Usage Description

NCCoE
National Cybersecurity Center of Excellence

NIST
National Institute of Standards and Technology

NIST IR
NIST Interagency or Internal Report

OOB
Out of Band

PSK
Pre-Shared Key

R&D
Research and Development

RFC
Request for Comments

SP
Special Publication

SSID

Service Set Identifier

TPM
Trusted Platform Module

URL
Uniform Resource Locator

Wi-Fi
Wireless Fidelity

WPS
Wi-Fi Protected Setup