# Q&A Session for “Trusted IoT Onboarding: An Introduction to Draft NIST SP 1800-36” Webinar

**Date:** Wednesday, November 1, 2023
**Time:** 10:00 a.m. to 11:30 a.m. (Eastern)

The NIST National Cybersecurity Center of Excellence (NCCoE) IoT Onboarding team would like to extend our sincerest thanks to each of you who attended our webinar on November 1, 2023. Your active participation, insightful questions, and shared experiences truly enriched the event and made it a great success!

We have prepared this Q&A response document to ensure all questions were addressed. This document encompasses all the questions submitted during the webinar and our comprehensive responses to each, along with the corresponding transcript and timestamps for those that were addressed live (see the post-event video). It is our hope that this further facilitates understanding and stimulates continued discussion on the subject. We truly appreciate your engagement and look forward to many more insightful conversations in the future.

---

## Webinar Q&A

**[1:19:50]** Paul Watrobski, Principal Investigator, NIST NCCoE

**Q: The first question is about the network onboarding service in the architecture diagram. Does this reside within the firmware of the device or is this external? Where does that live?**

Michael Richardson, Chief Scientist, Sandelman Software Works

A: Yes, you can’t do network onboarding without some firmware in the IoT device. You also can’t do network-layer onboarding without some device in the network to talk to that. What’s not always clear is that some of these things like WPA run over special packets, they’re called, that are part of 802.11 spec DPP runs completely over that, for instance, so we have to have a protocol that runs over that, and there’s a variety of different ones.

Brecht Wyseur, Senior Product Manager and Product Strategy, Kudelski IoT

A: There is, indeed, some part that needs to be in device—namely, some of the format that will make connection to some service—but when we talk about the part that’s on the right bottom of the notional architecture that we have, that should be outside the device. One of the reasons is because the network onboarding service is responsible to basically attribute ownership to a device and then say it can be onboarded on that network. And that information is something that is not existing in the device upfront. So during manufacturing, they do not know where the device should connect to, both network and application layer. So that's exactly why you need this external onboarding service. So service itself is outside of the device.

**[1:21:40]** Paul Watrobski, Principal Investigator, NIST NCCoE

**Q: The next question we got was regarding the credentials being unique for each device. For example, are these private keys unique for each of the devices that are connected to the network? This [question] came in during Build 1, so I'll turn it over to Dan and Danny.**

Dan Harkins, Fellow, HPE Aruba

A: Bootstrapping keys are definitely unique. So, these are generated—again, if you refer back to the factory use case, the idea is that when the device is manufactured, the manufacturing process is going to generate a unique key pair: either the device will generate it itself or as Michael described, the factory can generate it. But that key is just used for network-layer onboarding. Once the device has been onboarded, it's going to get another special credential that we'll use to access the network. So, part of the DPP onboarding process, it'll either get a, say, a passphrase for WPA3, it can get a certificate for WP3 enterprise, or it can get a connector if it wants to do the DPP AKM. A connector, again, is a unique credential; it's a signed public key. The public key that's being signed for the connector is different than the bootstrapping key that the thing used to do DPP, though. But again, there will be a unique credential for onboarding, and there will be a unique credential for bootstrapping, and then a unique credential for the device to get network-layer onboarded.

Michael Richardson, Chief Scientist, Sandelman Software Works

A: Absolutely the same story, [for BRSKI] you need to have a unique key, you need a certificate attached to that key (public key) and then it needs to be connected to the device in some way already in the factory.

Brecht Wyseur, Senior Product Manager and Product Strategy, Kudelski IoT

A: That is correct, [the same is the case for build 4] you must have different identities for each and every device, just because of the fact that two different devices might be owned by two different people, so you must be able to have that segmentation. The baseline of security is to have unique identity in the devices.

**[1:24:05]** Paul Watrobski, Principal Investigator, NIST NCCoE

**Q: On the note of ownership—is the assumption that all the network owners own the devices on the network?**

Brecht Wyseur, Senior Product Manager and Product Strategy, Kudelski IoT

A: First of all, ownership can be different between network-layer and application layer. Secondly, ownership of a device can be independent from network or application-server ownership. Thirdly, ownership can even change – can be transferred over time.

The real challenge will be the establishment and this management of ownership over time. The way how that can be done may also be very dependent on the use-cases. For example, you may have an out-of-band establishment of transfer of ownership from silicon vendors to device manufacturers and then operators through unique chipset identifiers; where ownership claim is then done contractually between those entities – or could be done on a first-come-first-served – or based on proof of possession, etc. The way how that will work may also be very different between a consumer IoT context, and an industrial context where you have more B2B or professional installation use-cases.

What is important is that at the moment of onboarding, the device and the onboarding service have the information to be able to associate device ownership to the network or application-service where the device needs to be onboarded to.

Dan Harkins, Fellow, HPE Aruba

A: For DDP-enterprise case, it I assumed that the network is the legitimate owner of all the devices. That's really the basis of the bootstrapping trust. So, the legitimate owner of these devices is going to possess the DPP URI or all the things that it wants to onboard, and it's possession of those things makes it a legitimate configurator, basically, and it's again, foundational to the trust of DPP.

Brecht Wyseur, Senior Product Manager and Product Strategy, Kudelski IoT

A: So maybe to challenge the previous comment, the device owner, does not necessarily need to be the network owner. For example, industrial systems or think about, for example, metering systems. You may have a certified meter, where you have metering data that goes towards a utility, but that network connectivity can go over another network from the local payments owner, for example. So you could use a different local network for enabling communication, but then what concerns the application layer application, that data or that ownership of that data can be from another party. So you can have different kinds of relations of ownership at different levels.

Darshak Thakore, Principal Architect, CableLabs

A: I think that's a good point to highlight. Is that one of the thing that we always talk about in this architecture is when we talk about ownership, right, there's actually multiple levels of ownerships here. When we talk about network ownership, I think what probably Dan said is somewhat true, right, is that you or whoever owns the network needs to be able to at least have ownership of the network side or to be able to provide network connectivity to those device, but when it comes to the application layer, it can be a completely different ownership, right? Somebody else might own it from the data and everything else that's going into the application layer. That is a distinction that we kind of make in Build 2, the network owner and the application layer owner might be different entities.

Brecht Wyseur, Senior Product Manager and Product Strategy, Kudelski IoT

A: You might even have transfer ownership during a lifetime of the device, over one million.

Nick Allott, CEO, NquiringMinds

A: And just to muddy the waters even more, not only can it be hierarchical, it can be an end-to end relationship. It's not always mutually exclusive, so it can get quite a naughty problem.

**[1:28:13]** Paul Watrobski, Principal Investigator, NIST NCCoE

**I'll pair these two questions together in the interest of time as they are related. So it was mentioned that these devices might need to be re-credentialed or kind of change ownership over their lifetime. We have one question first of: What is the process for offboarding a device and the timeframe for monitoring, as well as talking with a third-party market of refurbish devices being sold, how does that IoT onboarding process apply to those devices?**

Michael Richardson, Chief Scientist, Sandelman Software Works

A: I think that the millions of third-party devices are unfortunately landfill. It seems very unlikely that they're going to get a proper onboarding update firmware update from their manufacturer anytime soon, or at least that device isn't going to get it. The question written says, you know, it says that some of the device wiping doesn't even work. So, in that case, if you can't even remove the old credentials, and the odds of you being able to install new ones is very, very, very low. And that also argues, unfortunately, very much that they are going to be landfill. I think this is very unfortunate and I think there are some interesting opportunities to refurbish devices rather than just recycle them.

Paul Watrobski, Principal Investigator, NIST NCCoE

A: These solutions are kind of forward looking as opposed to trying to patch and somehow update these existing devices. They're kind of falling sort of into the legacy market, where you don't have the necessary means to be able to kind of implement what you guys are demonstrating here.

Brecht Wyseur, Senior Product Manager and Product Strategy, Kudelski IoT

A: Well, it all depends on what information already exists in the device. So, it might be that, for example, the firmware is not yet complete, but at least if you've got some hardware root of trust, for example, that allows us to do attestation, you could bootstrap a few activities already in there. For example, all the businesses, we've got a background, and then also the digital TV systems with setup boxes, we've got setup boxes about 30 years old in the field. And so, you want to figure out, what can you still do with them? You don't want to have to have them in landfill. There's regulation that says that you can't throw away those kinds of setup boxes if they are, let's say, too young. So if you want to have longevity and this is where you would need a minimum of, for example, hardware root of trust with some trusted credential? Once you have that you can try to figure out what functionalities are available. If I do a firmware update, can I authenticate that update and then bring in new functionality? It's going to be a challenge because you have to look at each system individually, but sometimes it's not necessarily landfill.

Darshak Thakore, Principal Architect, CableLabs

A: The best you can do is to do with the old devices is to isolate them out on a separate subnet or a separate network, right. And there are multiple ways of doing it, because as Michael and Craig also mentioned, some of these devices don't even have any of these mechanisms of having an individual key or having a secure onboarding mechanism, like DPP or anything else. So then you're left with nothing but providing some sort of a shared WiFi password, right, which what we do in Build 2 a little bit is we have to use full device, WPA2-PSK to give to those devices. So then that way, at least you can keep those devices separate. Having said that, there's a big caveat that WPA2 itself is sort of out the door. So you kind of have to put those devices on a separate subnet or a separate network and verify, isolate them out, and hope that eventually they'll be replaced.

Dan Harkins, Fellow, HPE Aruba

A: So, I kind of understood the question to be a little bit different. I agree with the comments about legacy devices not really being able to be onboarding using these technologies, but devices that have been onboarded, I think it is possible to do a factory reset. It basically clears all the credentials off that the device had: both it's network access credentials an any sort of credentials it had gotten in application layer onboarding. And when that happens, the device basically for DPP it reverts back to its chirping state and it's going to wait for someone to onboard it. And again, that's going to require a transfer of ownership of the DPP URI to the new owner. And one that happens, this device will be discovered by the new owner and onboarded. So I think the process will work for these devices that are being manufactured of these technologies, but I do agree that for older legacy devices, they are pretty much going to be pretty much landfill.

Nick Allott, CEO, NquiringMinds

A: I think there’s a lot of detail to work through here. I mean, fundamentally, you can’t trust an untrustworthy device to respond to a command. Now, whether it’s untrustworthy because it’s a legacy device? Was untrustworthy because the compromise has been found because it fails one of the checks? I think the two scenarios sort of collapse into each other. I think absolutely the concept of offboarding is in scope of the project. In terms of what it actually means and what you deem as an acceptable method of doing it, I think there’s a little bit more detail to work through.

**[1:34:11]** Paul Watrobski, Principal Investigator, NIST NCCoE

**The last one I’ll throw over is regarding interoperability…So which software layer or who in the supply chain is responsible to ensure interoperability? Is it possible for interoperability between these different protocols to have, say, deployment with DPP and BRSKI and/or Thread that are on that IoT device?**

Dan Harkins, Fellow, HPE Aruba

A: I don’t think so. There is interoperability of the individual protocols. HPE did interoperate with the configurator with Build 2 and CableLabs did interoperate between the HPE configurator in Build 1, so we have demonstrated interoperability of the protocol itself, but there is a fundamental different between, say, BRSKI and DPP that make them incompatible.

Darshak Thakore, Principal Architect, CableLabs

A: I think BRSKI and DPP came out of different use cases and DPP does have some enterprise support, but it was mainly geared towards home IoT environment, whereas BRSKI started off more at the enterprise side of things.

Michael Richardson, Chief Scientist, Sandelman Software Works

A: Absolutely right, and they have different benefit-risk tradeoffs. Having said that, there’s nothing that prevents a device manufacturer that could support both and the factory process makes that so that there’s a lot of commonalities between the two phases. At a higher level, there’s nothing that prevents a light switch that was onboarded with DPP from speaking to a lightbulb that was onboarded with BRSKI, provided that the application layer protocol is actually compatible. So, I expect to see a lot of stuff there, particularly in the small office kind of space, where probably they’re going to have a mix of enterprise devices and a mix of more consumer devices. So, we can expect to see many places where there’s mixes of different things.

Steve Clark, Security Technologist, SEALSQ, a subsidiary of WISeKey

A: Let me jump in here and comment. The objective here is to network-layer onboard, and once it’s onboarded onto your network, you can use the device, as Michael just mentioned, for any application-layer use case, including Matter or whatever else that you’re using. The certificates for the application layer are independent. So basically, onboarding gives you the general purpose connectivity so that you can talk to the Internet.

Dan Harkins, Fellow, HPE Aruba

A: I think one thing to keep in mind is that all of these builds are still mapping back to the notional architecture that NIST came up with, and that's probably a good way to think about what Michael was talking about these devices interoperating in that sense, where you know, if you have one device run through the notional architecture of DPP, another one can run through it through it with BRSKI. The end result is going to be the same, and we're going to have devices on the network that can talk to each other. So, in that sense, I think there is an interoperability that can be had by following the guidelines that NIST has been laying out.

## Other Q&A Questions (Answered offline)

Q: Is the NL onboarding service (bottom right box) residing within the firmware of the IoT device?

A: No, it is not residing within the firmware, it's a separate service.

Q: Are the credentials unique for each device?

A: Yes, they are all unique.

Q: i.e. private key unique for each device?

A: Same as above.

Q: There are millions of refurbished IoT devices on the market being sold by 3rd party suppliers and manufacturers. How does onboarding apply to those devices since wiping the device doesn't always work as stated.

A: For legacy devices that don't support onboarding protocols there's nothing that can be done. However, for devices that are capable of leveraging onboarding protocols, wiping the devices should effectively allow for re-onboarding but if the devices can't be wiped then they can't be re-onboarded.

Q: Do you assume that the network owner owns all devices on the network?

A: No that is not an assumption nor is it a requirement to successfully onboard devices. The network owner must have the bootstrapping information for the devices.

Q: Can we have a link to this demo? I have 4 Raspberry Pis I can work with.

A: NIST SP 1800-36 Volume C has a step-by-step walkthrough for implementations. Link: https://www.nccoe.nist.gov/projects/trusted-iot-device-network-layer-onboarding-and-lifecycle-management

Q: What is the process for off boarding a device and the timeframe for monitoring

A: The network operator would invalidate the credentials of the device being offboarded and the device would thus be removed from the network.