Securing Internet of Things (IoT) Services Involving Network Operators

May 2017

Foreword

As a leading technology and solutions development organization, the Alliance for Telecommunications Industry Solutions (ATIS) brings together the top global ICT companies to advance the industry's most pressing business priorities. ATIS' nearly 200 member companies are currently working to address 5G, the All-IP transition, network functions virtualization, big data analytics, cloud services, device solutions, emergency services, M2M, cyber security, network evolution, quality of service, billing support, operations, and much more. These priorities follow a fast-track development lifecycle — from design and innovation through standards, specifications, requirements, business use cases, software toolkits, open source solutions, and interoperability testing.

ATIS is accredited by the American National Standards Institute (ANSI). The organization is the North American Organizational Partner for the 3rd Generation Partnership Project (3GPP), a founding Partner of the oneM2M global initiative, a member of and major U.S. contributor to the International Telecommunication Union (ITU), as well as a member of the Inter-American Telecommunication Commission (CITEL). For more information, visit www.atis.org.

Table of Contents

Table of Figures

Table of Tables

1 Introduction

The adoption of Internet of Things (IoT) services is rapidly growing. IoT services can provide significant advantages to consumers, enterprises, and government institutions. It is important that as IoT services are designed and delivered, full account is taken of the security considerations both to protect the IoT service itself and to prevent IoT equipment becoming a source of attacks against other service users.

In some cases, the network operator's role in delivering IoT services is simply to provide connectivity and there is no direct technical or business partnering between the operator and the IoT service provider. In other cases, the network operator may take a more active role where the IoT service includes technical and business aspects under the control of the network operator. In this report, several different scenarios are introduced that characterize different relationships and levels of partnering that may exist between a network operator and an IoT service provider. In these scenarios, shared responsibility for securing the service may exist and consequences of security failures may be felt by both the network operator and the IoT service provider. The security implications of the various scenarios are discussed and practices that can be used to proactively address security in these scenarios are provided.

No part of this document should be taken as normative. Its purpose is to document practices that may be helpful to the development of good solution security. As each situation is different, it is necessary for the security approach to be chosen by the parties involved appropriately for their service, priorities, and circumstances.

2 Acronyms, & Abbreviations

For a list of common communications terms and definitions, please visit the *ATIS Telecom Glossary*, which is located at < http://www.atis.org/glossary >.

ACL	Access Control List
API	Application Programing Interface
ARA	Architectural Risk Analysis
ATIS	Alliance for Telecommunications Industry Solutions
CPE	Customer Premises Equipment
IoT	Internet of Things
IPSO	Internet Protocol Smart Objects
LWM2M	Lightweight Machine to Machine
M2M	Machine to Machine
NFV	Network Function Virtualization
NIST	National Institute of Standards and Technology
NoT	Network of Things
OCF	Open Connectivity Foundation
OIC	Open Internet Consortium
OMA	Open Mobile Alliance
RPK	Raw Public Key

PSK	Pre-Shared Key
SIM	Subscriber Identity Module
SM	Security Module
SMS	Short Message Service
SW	Software
VNF	Virtual Network Function
WWW	World Wide Web
XMPP	Extensible Messaging and Presence Protocol

3 Overview of Market Segments

3.1 Example Scenarios for IoT Service Offers by Operators

Many IoT services are provided by organizations that focus purely on IoT activities and treat connectivity (either over Wi-Fi or cellular) as a dumb pipe. However, more sophisticated IoT services – especially those involving many devices distributed over a wide area – benefit from closer attention being paid to connectivity aspects. The following scenarios consider cases where network operators are involved in the provision of IoT services either to leverage the value of their brand or to take advantage of their wide-area connectivity assets.

Typically, an IoT service will involve communications, device, and network (cloud)-based elements. Network operators are well equipped to provide connectivity but may choose to work with third parties to provide the other solution elements. There are a variety of different scenarios for the relationship between the network operator and third party IoT partners. Each of these scenarios has different security implications.

3.1.1 Operator develops, brands, and markets its own service

In this scenario, the operator takes complete responsibility for all aspects of the IoT service. Though the operator owns and manages the complete service, it may procure hardware and software elements from third parties and integrate them into the service. In particular, network operators will work with hardware vendors to build IoT devices.

For example, home broadband operators may provide home gateway devices or set top boxes as part of their package. With users' permission, these devices may monitor their electrical supply voltage and AC frequency. This data may be collected by the operator and supplied to power companies to allow them to assess the health of their distribution network in real-time.

In this type of scenario, the operator acts as the technical design authority for the complete service. As such, the operator should apply good security practices during the service design. This should include a full assessment of the security implications of the service design and the creation of a cybersecurity risk management plan.

Where the operator includes hardware or software from third parties in the solution, the operator's security assessment should define the security requirements placed on these elements. The operator should then audit these elements to ensure that they meet the identified security requirements.

3.1.2 Operator procures a partner's solution; brands and markets the service directly to customers

In this scenario, the operator sells an IoT solution developed and operated by a third party using the operator's brand. The operator makes no technical changes to the IoT solution.

This scenario is most likely to arise when the operator wishes to have additional services to offer as part of their package. For example, the operator may extend their home broadband service by offering additional services such as home security monitoring. Technically, the home security solution is provided by an existing third party.

In this scenario, the network operator is dependent on the third party for the cybersecurity of the service being marketed under their brand. In order to protect their interests, the network operator should ensure that the third parties solution has an effective approach to cybersecurity. This may be done, for example, by requiring the third party to share their cybersecurity risk management plan and by auditing the third party's compliance to that plan.

Some areas of cybersecurity planning may fall under joint responsibility in this model. In particular the handling of public relations during a cybersecurity incident requires coordinated activity by the network operator and the third party solution provider.

3.1.3 Operator provides access and other network services for an IoT/M2M partner to bundle and market to their customers – Co-branded services

In this scenario, the operator sells access and other network services to an IoT/M2M partner. The IoT/M2M partner will bundle with their solution and sell as a co-branded service or a joint enterprise. When effectively done, co-branding provides a way for companies to combine forces so that their marketing efforts work in synergy. The IoT/M2M partner benefits from the operators existing customer base and marketing.

In this scenario, the operator is basically providing the backbone and would be responsible for the normal day-to-day cybersecurity on their existing network. The IoT/M2M partner would be responsible for the cybersecurity planning for the IoT/M2M solution.

Overall cybersecurity should be reviewed and part of the contract/agreement by both parties in this model. Due to the co-branding, a cybersecurity incident could have negative impacts on the network operator and/or the third party solution provider. Again in particular the handling of public relations during a cybersecurity incident requires coordinated activity by the network operator and the third party solution provider.

3.1.4 Operator provides access and other network services for an IoT/M2M partner to bundle and market to their customers - Operator and IoT/M2M provider each have an independent relationship with customers

In this scenario, the operator provides the access and other network services that enable an IoT/M2M service. The network provider and third party solution provider will each independently own the relationships with their respective customers.

In this scenario, both the operator and the third party are responsible to provide cybersecurity planning as well as validation of the others plan. Overall cybersecurity should be reviewed and part of the contract/agreement by both parties in this model.

A cybersecurity incident on either the network provider or the third party provider would have a negative impact on both. Again in particular the handling of public relations during a cybersecurity incident requires coordinated activity by the network operator and the third party solution provider.

3.2 Protecting Brand Value when providing IoT Services

As discussed above, IoT services may include elements from several different sources. All these elements contribute to the overall security of the service. Where operators market IoT services under their band any security failures will inevitably be associated with the operator's brand wherever the security vulnerability originated. This is particularly true because most network operators have a high profile in their operating countries and spend heavily on brand development.

While the brands of third party providers of service elements are less at risk in the event of a security breach, it is still possible that they will receive negative publicity if they become associated with a security failure. Therefore, third parties that offer solution elements should still have an independent cybersecurity management plan.

In order to protect their brands, operators marketing IoT services should have their own cybersecurity management plan as described in the previous scenarios. This will manage the risk of cybersecurity incidents. This plan should include a section dealing with incident response that includes the management of public relations during and after a cybersecurity incident. The cybersecurity plan should include a plan to enable service restoration following a cybersecurity incident.

4 Survey of Existing Work on IoT/M2M Security

The need to improve security for IoT/M2M applications has been widely recognized in the industry. There are a number of relevant activities summarized in the following table. During the analysis it was found that the various working groups have generally produced one of two types of recommendation. Some working groups produce general guidance on IoT security that is not tied to a specific technology platform (e.g., the work by the Cloud Security Alliance). These are labelled with a "G" in the table below. Other working groups define specific technology platforms to support IoT services that include platform-specific security features (e.g., oneM2M). These are marked with a "P" in the table below.

Table 4.1: Working Group Activities Relevant to IoT Cybersecurity

Working Group (Active Since)	Charter	Comments	Type G=General P=Platform
IPSO Alliance (Sep 2008)	Establish *Internet Protocol (IP)* as the network to interconnect smart objects and allow existing infrastructure to be readily used without translation gateways or proxies.	Extensively reuses existing, industry standard, IP protocols. We did not find significant, original, security content.	G
IoT-A (2010-2013)	Developed an architectural *reference model* to allow seamless integration of heterogeneous IoT technologies into a coherent architecture to realize 'Internet of Things' rather than 'Intranet of Things'.	This project was conducted as part of the EU's research agenda and aimed to define a framework for later industry consideration rather than create normative specifications for commercial application. Its security recommendations may be a useful input to security modelling.	G
AllSeen Alliance (2013)	Collaborate for an open, universal IoT software framework across devices and industry applications, based on AllJoyn *open source project*, originally developed by Qualcomm but now released to community developers.	See write-up below.	P
Industrial Internet Consortium (Mar 2014)	Accelerate development and adoption of *intelligent industrial automation* for public use cases.	The Industrial Internet Consortium has excellent participation from a variety of industry vertical sectors. Their published security guidelines are useful for their application domain and compatible with the approach in this document.	G

Working Group (Active Since)	Charter	Comments	Type G=General P=Platform
HyperCat (May 2014)	Develop an open specification for IoT that will make data available in a way that others could make use of it, through a *thin interoperability layer*.	Focus is on a high level object model for interoperability.	P
Open Interconnect Consortium (Jul 2014)	Define interoperable *device communication standards* (for peer-to-peer, mesh & bridging, reporting & control etc.) across verticals, and provide an *open source* implementation	See write-up below.	P
IEEE P2413 (Jul 2014)	Create a *standard interoperability architecture* and define commonly understood data objects, for information sharing across IoT systems.	This work is progressing but not yet mature enough (at the time of writing) for the significance of the security content to be clear.	G
OMA LWM2M (2014)	Proposed a new *Light-weight M2M protocol standard*, based on client-server model for remote management of M2M devices and related service enablement	See write-up below.	P
oneM2M (July 2012)	There is a need for a common, efficient, easily and widely available M2M Service Layer, which can be readily embedded within various hardware and software.	See write-up below.	P
Cloud Security Alliance IoT Security Guidance for Early Adopters of the Internet of Things (IoT)	The Cloud Security Alliance IoT Working Group focuses on understanding the relevant use cases for IoT deployments and defining actionable guidance for security practitioners to secure their implementations	The CSA IoT guidance helpfully addresses many common security issues for IoT. It is generally relevant to implementers and is compatible with the approach in this document.	G
NIST SP 800-183	A composability model and vocabulary that defines principles common to most, if not all networks of things, is needed to address the question: “What is the science, if any, underlying IoT?”	See write-up below.	G

Working Group (Active Since)	Charter	Comments	Type G=General P=Platform
GSM Association "Internet of Things (IoT) Security Guidelines"	To provide high-level discussion of challenges, IoT models, risk assessments, and solution spaces. To give the implementer of an IoT technology or service a set of design guidelines for building a secure product.	Good overview of security issues, but contents is generally well-known by security professionals. Much of the contents is familiar to experts and applies equally to IoT and non-IoT systems. Useful as a checklist of things for the Network Operator to consider when focusing on IoT.	G

The working groups shown in bold in this table are considered to be particularly relevant to the security of IoT systems and are discussed below.

4.1 Delivering IoT Security Through the Use of Specific Technology Platforms

As identified in the table above, there are several groups offering standardized platforms for delivering IoT services. These may take the form of system protocol standards and/or Open Source IoT platform implementations. Generally, these platforms provide a middleware software layer that provides commonly needed IoT facilities that can be accessed to deliver IoT services.

From the security point of view these platforms often include specific security features as part of the IoT facilities that provide. These security features may include:

- Data encryption to prevent eavesdropping.
- Data integrity protection to prevent data tampering.
- Access control covering access to software facilities and data.
- Authentication and authorization control to prevent identity spoofing.
- Credential management to allow security credentials to be managed and securely stored.

Platforms differ in how they encompass specific applications and how much control individual applications have over the security features used. One common model is that individual applications can access platform services via an Application Programing Interface (API). Applications may choose, as part of their software, which of the security features available on the platform to employ.

By providing defined security features a standard platform can enable IoT applications to be designed with good technical security without application designers and integrators having to invent their own security solutions. However, application designers have to ensure that they select and use appropriate security features from the platform. Therefore, the security needs of each application must still be analyzed for all aspects of the system.

It must also be recognized that technical security features embedded in technology platforms are only part of the holistic security approach that should be applied to IoT services. This is for two reasons:

- The technology required to build a whole solution will, in most cases, include elements that are outside the scope of the IoT security platform. For example, system management tools and back-office data processing are typically not covered in the standard IoT platforms.
- The whole security solution involves not just technical aspects but also aspects of company organization and processes. This is particularly true if multiple companies are involved in the delivery of services (as discussed in section 5).

The involvement of a number of security experts in the development of standard IoT platforms should give some confidence that security features in these platforms are well considered and fit for purpose. However, this should not be taken for granted:

- It is possible for security problems to remain undetected even in systems designed according to the best known processes.
- Most security considerations involve a trade-off between security and other features. Within the platform design certain trade-offs may have been implicitly or explicitly made that may not be appropriate for a particular application, particularly if this application was not foreseen by the platform designers.
- Where the platform exists as a specification only, with no software, an implementation of the specification will have to be made in software. During this implementation process it is possible that new security problems could be introduced.
- Where Open Source software is provided this software may contain security problems created as part of the implementation. This may occur either in the platform-specific code or in dependencies such as security libraries. Though Open Source software has generally achieved a good level of security there have been several incidents where widely used security libraries were found to contain serious bugs.

4.1.1 Recommendations for the use of IoT Platforms

The following recommendations are made for the use of IoT platforms to improve the security of IoT systems:

- Whether or not IoT services are going to be built using a standard IoT platform, the service should still be subject to a comprehensive security threat analysis leading to a secure solution design and secure operating procedures.
- Use of standard IoT platforms that contain pre-defined security features can help deliver secure IoT services. Benefits include:
    - Easy access to pre-defined security features will encourage these features to be used.
    - Security features developed collaboratively using “many eyes” should be of higher quality that security features designed by individuals, particularly if those individuals are not security experts.
- Based on the service security design, the service implementation should ensure it uses appropriate security features of the IoT platform.
- Service designers should ensure they properly understand the security features provided by the IoT platform and use them in an appropriate fashion.
- Service designers should not assume that security features provided by the platform are infallible and should provide layered security using multiple approaches to protect the whole solution.

4.2 AllSeen Alliance

4.2.1 References and Links

AllSeen Alliance web site: < https://allseenalliance.org/ >

List of certified products: < https://allseenalliance.org/certification/certified-products-directory >

4.2.2 Background and Purpose

“The AllSeen Alliance is a cross-industry consortium dedicated to enabling the interoperability of billions of devices, services and apps that comprise the Internet of Things.”

The AllSeen Alliance is part of the Linux Foundation. It develops and maintains the “AllJoyn” Open Source IoT platform originally defined by Qualcomm. AllJoyn was originally developed with a focus on peer to peer device communications within a single home environment. It supports applications like smart home appliances and smart home entertainment systems. AllJoyn is now expanded to include support for cloud-based solutions and industrial and commercial IoT applications.

4.2.3 Industry Adoption and Current Activity

The AllSeen Alliance currently has about 200 members. As of August 2016 the web site lists 24 certified products including lighting, home entertainment, and HVAC-related products. Windows 10 contains built-in support for AllJoyn and developer APIs[1].

Currently the AllSeen Alliance is continuing development on the AllJoyn software stack.

4.2.4 IoT Security in AllSeen Alliance

The AllJoyn software stack provides technical features to protect the security of operations between devices using the platform. Capabilities include:

- Authentication
- Certificate/key management and storage
- Data confidentiality protection on interfaces
- Access Control List (ACL) based policy

As shown in the diagram below, the architecture separates specific AllJoyn applications (“Apps”) from the common AllJoyn core that provides services to the applications. There is an API between the apps ad the AllJoyn core. In the figure two AllJoyn endpoints (the Consumer and the Producer) communicate via the AllJoyn Core. Individual Applications are in control over which of the platform’s security features they invoke.

As well as security communication according to the needs of applications the AllJoyn Core is responsible for the enforcement of certain security policies such as ACLs. In this diagram the End-User Security Manager is shown configuring policies in the AllJoyn Core of the endpoints.

[1] https://developer.microsoft.com/en-us/windows/iot/docs/alljoyn

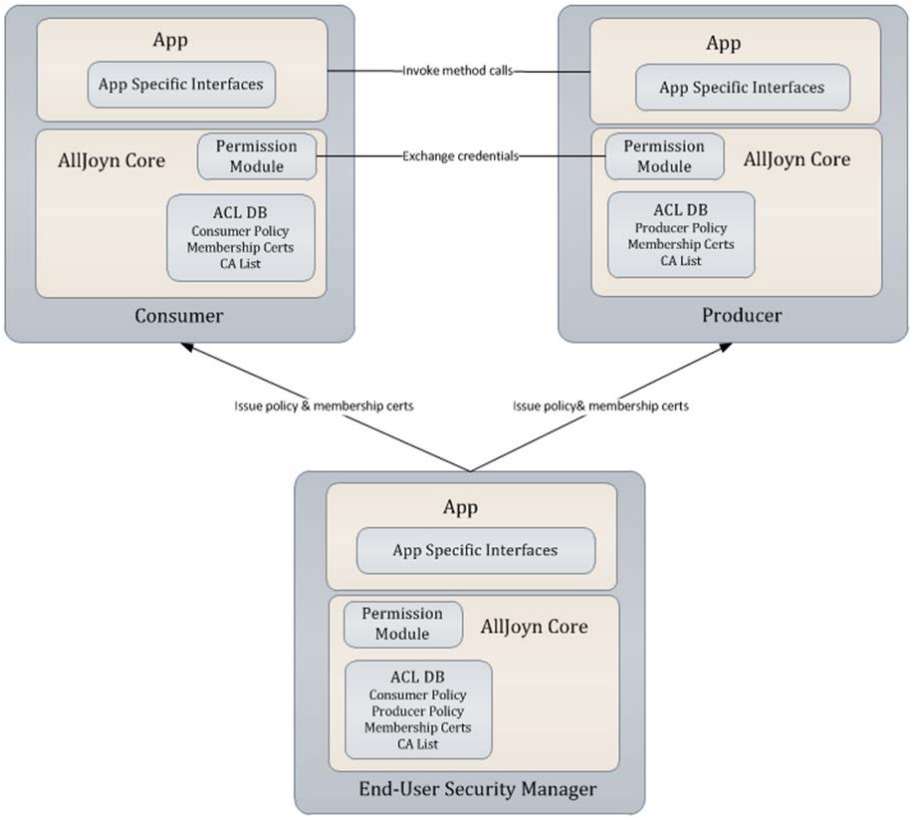

Figure 4.1: AllJoyn "Security 2.0" Architecture

4.2.5 AllSeen Alliance Conclusions

The AllSeen Alliance's AllJoyn platform is in active use by a number of IoT vendors, particularly for consumer applications within the home environment. For applications within an individual LAN, the AllJoyn Platform may not be very relevant to network operators. As AllJoyn's domain of applications expands to include cloud based services, it may become of more practical interest to network operators.

As discussed above, platforms like AllJoyn can address aspects of application technical security without application designers and integrators having to invent their own security solutions. However, application designers have to ensure that they select and use appropriate security features from the platform. Therefore, the security needs of each application must still be analyzed for all aspects of the system.

The AllSeen Alliance does not attempt to provide a comprehensive approach to the analysis of organizational or system security aspects.

4.3 One M2M

4.3.1 References and Links

oneM2M web site: < http://www.onem2m.org/ >

oneM2M Security Solutions Specification (Release 2):
< http://www.onem2m.org/images/files/deliverables/Release2/TS-0003_Security_Solutions-v2_4_1.pdf >

4.3.2 Background and Purpose

"The purpose and goal of oneM2M is to develop technical specifications which address the need for a common M2M Service Layer that can be readily embedded within various hardware and software, and relied upon to connect the myriad of devices in the field with M2M application servers worldwide. A critical objective of oneM2M is to attract

and actively involve organizations from M2M-related business domains such as: telematics and intelligent transportation, healthcare, utilities, industrial automation, smart homes, etc."

oneM2M is a partnership project with support from several international SDOs including ATIS and ETSI. They have specified a platform for M2M communication that acts as "network operating system" or "middle layer" providing common services needed by IoT systems. The overall architecture and service functions are illustrated below.

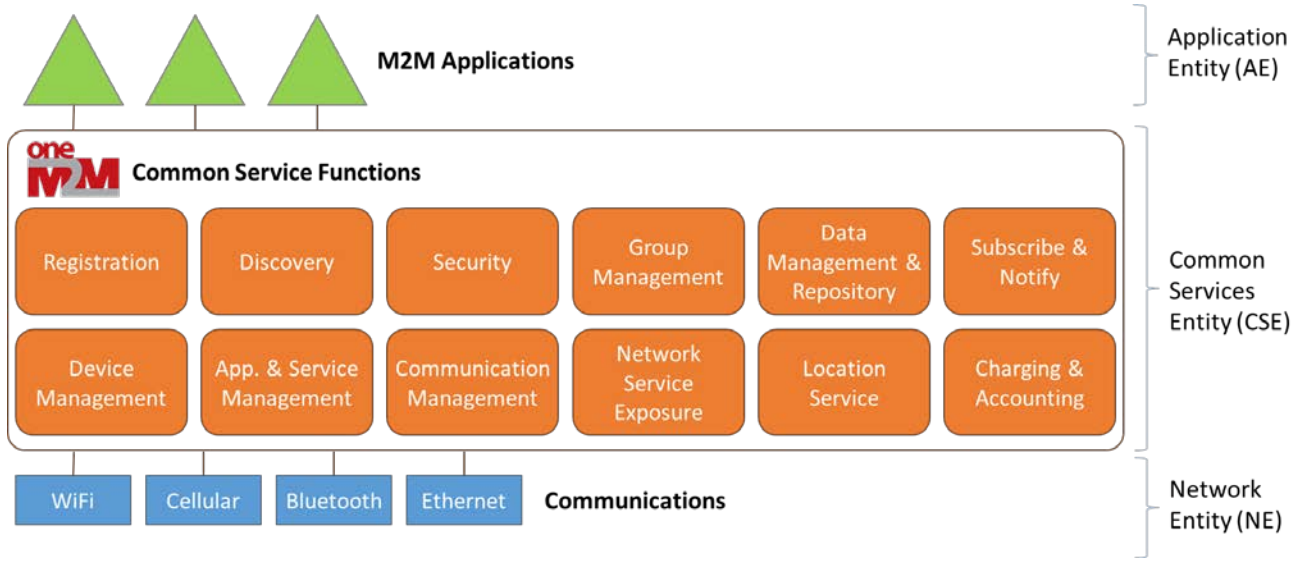

Figure 4.2: oneM2M Architecture

4.3.3 Industry Adoption and Current Activity

Several oneM2M-based applications are known to be available and industry interest in the standard appears to be growing. Three major open source implementations of oneM2M available:

- OpenM2M (Eclipse)
- IoT Data Management (IoTDM) (Open Daylight)
- Open allianCE for iot stANdard (OCEAN)

oneM2M published "Release 2" of their standard in summer 2016 and is now continuing with the development of Release 3.

4.3.4 IoT Security in oneM2M

oneM2M provides a variety of security features including:

- Identification and Authentication
  - Various flavors, e.g., pre-shared key, certificate-based, 3GPP identity-based.
- Security Association Establishment between communicating entities.
  - Provides confidentiality and integrity protection.
  - Uses (D)TLS in the signaling.
- Authorization
  - Policy based on a "resource" model concept.
- Remote Provisioning
- User Data Privacy
  - Layered privacy controls.

- Identity protection
    - Pseudonyms to create anonymity.
- Sensitive Data Handling
    - Storage of data in secure environments.
    - Specified in some detail for global platform smart-cards. Less detail for other types of security environment.
- Security Administration (related to device management)
    - Creates and administers dedicated Secure Environments and provisioning of master credentials.

4.3.5 oneM2M Conclusions

oneM2M is being actively adopted by industry and may become a commonly used technology in future IoT solutions. The standard provides a wide range of security related technical features as shown above.

As discussed above, platforms like oneM2M can enable IoT applications to address aspects of application technical security without application designers and integrators having to invent their own security solutions. However, application designers have to ensure that they select and use appropriate security features from the platform. Therefore, the security needs of each application must still be analyzed for all aspects of the system.

The oneM2M standard not attempt to provide a comprehensive approach to the analysis of organizational or system security aspects.

4.4 Open Mobile Alliance Lightweight M2M (OMA LwM2M)

4.4.1 References and Links

Open Mobile Alliance web site: < http://openmobilealliance.org/ >

OMA LwM2M Specifications web site: < http://technical.openmobilealliance.org/Technical/technical-information/release-program/current-releases/oma-lightweightm2m-v1-0 >.

4.4.2 Background and Purpose

The OMA is a standards development organization that produces open standards for the mobile phone industry. Founded in 2002, its membership includes wireless industry players such as equipment and mobile systems manufacturers and software vendors. To avoid overlapping its specifications with those of other standards organizations, the OMA coordinates with other standards bodies, including the Internet Engineering Task Force, (IETF), World Wide Web Consortium (W3C), and both the 3rd Generation Partnership Project (3GPP) and Project 2 (3GPP2).

The stated goals[2] of the OMA are to:

1. Deliver high quality, open technical specifications based upon market requirements that drive modularity, extensibility, and consistency amongst enablers to reduce industry implementation efforts.
2. Ensure OMA service enabler specifications provide interoperability across different devices, geographies, service providers, operators, and networks; facilitate interoperability of the resulting product implementations.
3. Be the catalyst for the consolidation of standards activity within the mobile data service industry; working in conjunction with other existing standards organizations and industry fora to improve interoperability and decrease operational costs for all involved.

[2] Source: OMA website at < http://openmobilealliance.org/about-oma/ >.

4. Provide value and benefits to members in OMA from all parts of the value chain including content and service providers, information technology providers, mobile operators, and wireless vendors such that they elect to actively participate in the organization.

4.4.3 Industry Adoption and Current Activity

The OMA produces specifications that address many topics important to the mobile industry, such as digital rights management, mobile broadcast services, mobile device management, and many others. Of particular interest is the OMA work on Mobile-to-Mobile (M2M) communications, for which they have developed the Lightweight M2M protocol standard (LwM2M). Their goal is to produce a light-weight M2M protocol standard based on client-server model for remote management of M2M devices (and related service enablement). The OMA vision for LwM2M is that of a fast deployable client-server specification to provide machine-to-machine service.

LwM2M provides the interface between M2M device and M2M Server It is primarily a device management protocol, though it is designed to be extensible to meet the requirements of applications, including the transfer of service and application data.

The OMA LwM2M Requirements and Architecture documents were published in 2013. The majority of the LwM2M specifications were released in April 2016. These include:

- Enabler Release Definition for Lightweight Machine to Machine: 2016.
- Lightweight Machine to Machine Technical Specification: 2016.
- LWM2M_Security: 2016.

Due to the short period of time between the release of those specifications and the writing of this section (September 2016), the details of industry adoption are unclear.

4.4.4 IoT Security in OMA LwM2M

Security services have been developed within an LwM2M Security Object, which features:

1. **Pre-Shared Key (PSK) Mode**: A binary shared secret key between the Client and Server. The default PSK Cipher Suites defined in this specification use a 128-bit AES key. The corresponding PSK Identity for this PSK is stored in the Public Key or Identity Resource. The PSK Identity is simply stored as a UTF-8 String. Clients and Servers MUST support a PSK Identity of at least 128 bytes in length.
2. **Raw-Public Key (RPK) Mode**: Requires a public key and a private key of the appropriate type and length for the Cipher Suite used. The default RPK Cipher Suites is a 256-bit ECC key.
3. **Certificate Mode:** Requires an X.509v3 Certificate along with a matching private key.
4. **SMS Payload Security**: Defines the format of Secret Key and Public Key and Identity resources of the LwM2M Server and LwM2M Bootstrap Objects when using SMS Payload security. These resources are used to configure keying material that a Client uses with a particular Server.
5. **Security Object Management**: Essentially, creation and deletion of a Security Object instance.

4.4.5 OMA LwM2M Conclusions

Market demand for a common set of standards for managing light-weight and low-capability devices on a variety of networks has provided the incentive for the OMA efforts in the M2M area. The industry expectation is that billions of devices will be connected to networks in homes and businesses, in public transportation, and in areas such as smart metering and connected medical devices. The OMA initiative recognizes that the capability to monitor, provision, and manage these billions of connections is absolutely essential; and that standards to support a tremendous assortment of M2M implementations will be essential to the expanding market.

By providing a security framework for some parts of the IoT service – particularly device management – the OMA LwM2M platform can help improve IoT security. However, as with any platform, it is important to apply LwM2M's capabilities in the context of an overall service security architecture.

4.5 Open Interconnect Consortium (OIC)/Open Connectivity Foundation (OCF)

4.5.1 References and Links

Open Interconnect Consortium web site: < https://openconnectivity.org/ >

OIC IoT Specifications web site: < https://openconnectivity.org/resources/specifications >

4.5.2 Background and Purpose

The Open Interconnect Consortium (OIC) is an industry group whose work is to advance standards and certification for IoT devices based on the Constrained Application Protocol (CoAP).The OIC was created in July 2014 by Intel, Broadcom, and Samsung Electronics. In 2016, it became the Open Connectivity Forum, which (as of this writing) has over 200 members spanning the globe.

4.5.3 Industry Adoption and Current Activity

To realize IoT (also termed the Network of Everything), the OIC has delivered a framework that enables the IoT requirements of easy discovery and trusted, reliable connectivity between things. It realizes these requirements via a specification, a reference implementation, and a certification program. In September 2015, the OIC released a candidate specification (version 1.0) that covers the core framework, smart home devices, resource types, security, and remote access capabilities. This specification is available to the public and is accessible for non-members without registration. It is divided into two sets of documents:

1. **Core Specifications**: These documents specify the OIC Framework; i.e., the OIC core architecture, interfaces, protocols, and services to enable OIC profiles implementation for IoT usages and ecosystems.
2. **Vertical Profiles Specifications**: These documents specify OIC profiles that enable IoT usages for different market segments; e.g., smart home, industrial, healthcare, and automotive. The Application Profiles Specification is built upon the interfaces and network security of the OIC core architecture defined in the Core Specifications.

These specifications (our items of interest) consists of five volumes.

1. **OIC Core Specification**: This specification describes the OIC architecture, which enables resource-based interactions among IoT artifacts; i.e., physical devices and applications. It takes advantage of existing industry standards and technologies to provide solutions for connecting and managing the flow of information among disparate devices (i.e., different form factors, OSs, service providers).
2. **OIC Remote Access Specification**: This brief describes the use of Extensible Messaging and Presence Protocol (XMPP) and Interactive Connectivity Establishment (ICE) using STUN[3] and TURN[4] to add Internet connectivity securely and “scalably” to 1) constrained device networks and 2) network topologies that obfuscate or otherwise inhibit general connectivity.
3. **OIC Resource Type Specification**: This document specifies the resources and properties that may be defined for OIC resources.
4. **OIC Security Specification**: This specification defines security objectives, philosophy, resources, and mechanisms that affect base layers of OIC Core specification. See Section 2.3.4 for further description of this specification.
5. **OIC Smart Home Device Specification**: This brief describes the constructs used for an OIC Smart Home Device and the resources mandated to be implemented for each OIC Smart Home Device.

[3] Session Traversal Utilities for NAT (Network Address Translation).

[4] Traversal Using Relays around NAT.

4.5.4 IoT Security in OIC

OIC Security Specification addresses nine specific topics:

1. **Security for the Discovery Process**: Security considerations, discoverability of security resources.
2. **Security Provisioning**: Device identity for devices with, device ownership transfer methods, provisioning.
3. **Security Credential Management**: Credential lifecycle, credential types, certificate based key management.
4. **Device Authentication**: Device authentication with symmetric keys, raw asymmetric keys, and certificates.
5. **Message Integrity and Confidentiality**: Session protection with Datagram Transport Layer Security (DTLS), cipher suites.
6. **Access Control**: Access Control List (ACL) generation and management.
7. **Security Resources**: Credential Resource (Key Formatting, Credential Refresh Method Details), Certificate Revocation List, Security Services Resource, ACL resources (e.g., Access Manager, Signed ACLs), Provisioning Status Resource.
8. **Core Interaction Patterns Security:** (under construction; no definition).
9. **Security Hardening Guidelines**: E.g., secure storage, secure boot, software updates.

This document is intensely detailed, as is needed for a specification of this type. It addresses security both broadly, as the preceding list shows; and deeply, sacrificing no detail. Written into it are points and pointers that suggest efforts to address security at an architectural level.

4.5.5 OIC Conclusions

The OIC/OCF is defining a wide-ranging communications framework, the goal of which is to enable emerging applications in all significant vertical markets. This framework is intended to facilitate many new methods of communication, such as Peer-to-Peer, Mesh and Bridging, Reporting and Control, and many others. It includes a self-consistent implementation of identity, authentication, and security, whether the implementation be standard User IDs, Enterprise and Industrial IDs, or other forms of Credentials. The framework supports a “building block” architecture and provides an Open Source implementation based on the standards and a reference implementation it defines.

The OIC framework provides a wide range of security features that may be advantageous for IoT services. As with any platform, it is important to apply OIC’s capabilities in the context of an overall service security architecture.

4.6 NIST SP800-183: Network of ‘Things’

4.6.1 References and Links

NIST web site: < https://www.nist.gov/ >

NIST SP800-183 URL: < http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-183.pdf >

4.6.2 Background and Purpose

The National Institute of Standards and Technology (NIST) is a non-regulatory agency of the U.S. Department of Commerce, charged with promoting innovation and industrial competitiveness. Founded in 1901 as the National Bureau of Standards (NBS), with the mandate to provide standard weights and measures and to serve as the national physical laboratory for the United States, NIST now provides industry, academia, government, and other users over 1,300 standard reference materials. Many of these references are recognized as having specific characteristics and content, and encompass topics such as calibration standards for measuring equipment and procedures, quality control benchmarks for industrial processes, and experimental control samples. Topics of NIST documents, reports, and publications include buildings and construction, semiconductors, electromagnetics, manufacturing, chemistry, physics, information technology, cybersecurity, bioscience, and (its original mission) metrology.

4.6.3 Industry Adoption and Current Activity

U.S. industry and government have adopted NIST standards, with the level of compliance to NIST standards being assessed. In some areas, NIST has taken a leading role in addressing industry. Examples of this proactivity can be found in NIST's 2015 creation of a new series of Special Publications – the 1800 series – devoted specifically to cybersecurity, and in the release of the July 2016 Special Publication SP800-183, *Networks of 'Things'* – the item of interest to us at this time for its treatment of IoT from an elemental perspective. The following section discusses this document.

4.6.4 IoT Security in NIST

SP800-183 is an attempt to abstract all IoT functionality into a small set of primitive functions that perform the basic functions of sensing their surroundings, performing computations from sensor data, communicating raw and aggregated data among the involved entities, and actuating other actions on the basis of the sensory data and the computations. Feedback loops provide data on the results of an actuated action, which is used by primitives within functional system to further hone the system's responses or to make other decisions.

SP800-183 first proposes that IoT is but one example of a more general Network of Things (NoT), and then goes on to present a common vocabulary to foster a better understanding of IoT and better communication among parties discussing IoT. This common lexicon has the advantage of offering a basis for understanding of IoT as an "organism" of sensing, computing, communication, and actuation. Among the terms SP800-183 attempts to define are the following:

Primitive: An item that consists of small pieces from which larger blocks or systems can be built. A primitive can be developed or derived from something else, but is itself the smallest building block within a NoT. (The primitive might be viewed as akin to an atom, which is the basic building block of matter, yet is itself composed of smaller units.)

Distributed system: "A software system in which components located on networked computers communicate and coordinate their actions by passing messages. The components interact with each other in order to achieve a common goal."

Thing: A deliberately undefined term that nonetheless is a basic construct of any NoT. The fact that it is undefined permits greater flexibility in applying it, and allows a deeper level of abstraction to be placed on a NoT. (A Thing can be favorably compared to the point, line, and plane, the most basic entities of mathematics; which, by virtue of being undefined, extend the scope of mathematics in untold ways by providing a level of abstraction and generalization without which the common features of many mathematical concepts could not reveal themselves.)

SP800-183 takes the position that "no simple, actionable, and universally-accepted definition for IoT exists" (which is probably true). Instead, it proposes a model and a vocabulary that reveal underlying foundations of the IoT and expose the ingredients that express how the IoT behaves without actually defining an IoT. In doing so, it considers an IoT to be one type of a NoT; and a NoT to be one kind of distributed system.

SP800-183 also presents five core system primitives that are seen as belonging to most distributed systems. System primitives provide a basis for discussions of IoT formal models, simulations, and reliability and security tradeoffs. These primitives are the basic building blocks for a NoT.

The five NoT primitives are the:

1. **Sensor**: An electronic utility that measures physical properties such as temperature, acceleration, weight, sound, location, etc. All sensors interface with a controlled process or open environment by mechanical, electrical, chemical, optical, or other means.
2. **Aggregator**: A software implementation based on mathematical functions that transform groups of raw data into intermediate, aggregated data. Raw data can come from any source. Aggregators help manage "big" data.
3. **Communications channel**: A medium by which data is transmitted (e.g., physical via Universal Serial Bus, wireless, wired, verbal).
4. **External utility (eUtility)**: A software or hardware product or service. This definition is deliberately broad to allow for unforeseen future services and products.

5. **Decision trigger**: A construct that creates the final results needed to satisfy the purpose, specification, and requirements of a specific NoT. As its name implies, it "triggers" an action based on input from other primitives within the system.

Most NoTs contain all of these primitives; all NoTs contain at least one.

SP800-183 introduces other concepts that facilitate the discussion and understanding of the NoT idea. The interested reader is encouraged to read the document, which is only 25 pages in length and is straightforward in its presentation of the ideas.

4.6.5 NIST SP800-183 Conclusions

Unlike most NIST Special Publications, SP800-183 is primarily concerned with introducing a modelling concept rather than normative standards. The material it presents is generic to all distributed systems that employ IoT technologies. It is unique in its ideas and manner of presentation.

This document acknowledges that the Internet is a network of networks, but takes the position that focusing on restricted NoTs in a bounded way provides a more effective way to address trustworthiness problems than an unbounded Internet does. To reinforce this perspective, it defines primitives (and additional considerations omitted from this overview), which help establish the parameters for reliability and security when defined and used as described. Primitives also allow for analytics and formal arguments of IoT use case scenarios.

The ideas presented in SP800-183 are worthy of consideration, though they must still be "proven" by mapping them to real IoT systems and constructs. Such mappings are needed to demonstrate the applicability of the abstract concepts SP800-183 presents.

5 Architecture and Trust Boundaries

5.1 General Architectural Model

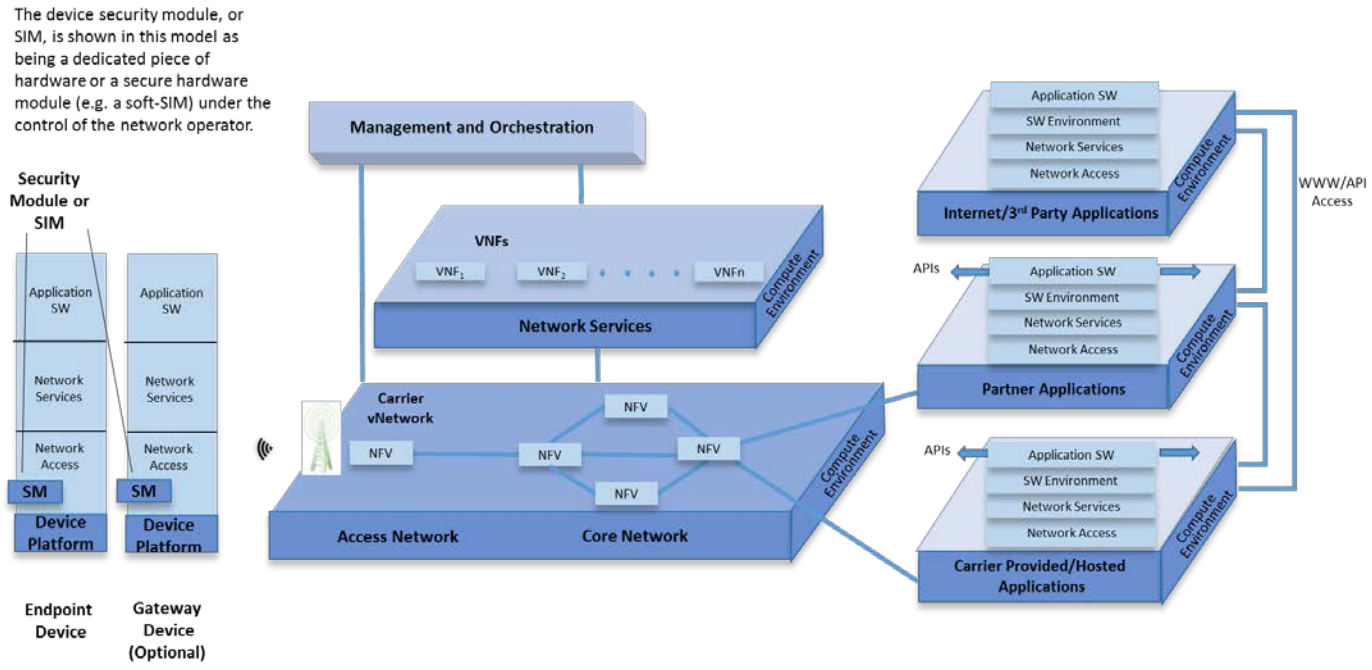

Figure 5.1: General IoT Architectural Model

The diagram above depicts a general network model for IoT applications – particularly those delivered using the partnership scenarios described in Section 3. This diagram extends the network model shown in [ARA process], which establishes a general model for analysis of network security.

The core blocks of the network (including the “Carrier vNetwork”, the “VNFs”, and the “Management and Orchestration” blocks) are essentially the same in this diagram as they are in the [ARA process]. Some modifications have been applied in other areas which are discussed below.

In the Customer Premises Equipment (CPE) or device part of the diagram the optional “Gateway Device” has been added. This reflects the practical deployment of many IoT solutions where the actual endpoint communicates to the network via a gateway rather than directly. For example, in a home automation system the gateway could be the home router and the endpoint could be a room thermostat. In another example related to a personal area network, the endpoint could be a fitness monitor and the gateway could be a mobile phone. Not all IoT solutions will include a gateway, but the inclusion of this entity in the diagrams highlights this as a possibility that needs to be considered.

Another feature shown in the CPE or device part of the diagram is the “Security Module or SIM”. Some (though not all) IoT devices and gateways will contain a security module that is under the direct control of the network operator. This contrasts with the main part of the device that is heavily controlled by the device vendor. The most common example of a such a security module is the SIM or Smart Card used to secure access to GSM, UMTS, and LTE networks. In the case of a SIM card the security module is physical module. In other cases a soft-SIM or similar secure enclave may be used. From a security point of view, the presence of these secure modules is important because it provides one part of the device which may be closely trusted by the network operator.

For IoT scenarios there are three major blocks in the applications domain. The “Carrier Provided/Hosted Applications” block represents applications that are under the direct control of the operator. This could include, for example, applications that the operator has developed itself and hosts on its infrastructure. The “Partner Applications” block represents applications that the operator is involved with due to a partnership with an IoT provider under one of the scenarios described in Section 4. The existence of a relationship between the network operator and the application partner implies some level of trust exists between these entities. Both the carrier provided/hosted applications and the partner applications may be able to access services from the carrier vNetwork as shown by the lines linking these blocks.

The final block of applications is the “Internet/3rd Party Applications” block. These are applications that do not have a strong trust relationship with the network operator. Instead of a negotiated access to operator facilities, these applications typically access data via an open public API. These public APIs typically provided a carefully controlled access to service features so that only a minimum level of trust is required between the API provider and the API consumer. The links on the right of the diagram labelled “WWW/API Access” represent this kind of public API. An example of an “Internet/3rd Party Application” could be an application that consolidates health data from a number of different sources to present the user with a consolidated dashboard of health metrics.

5.2 Overview of Trust Boundaries and Trust Domains

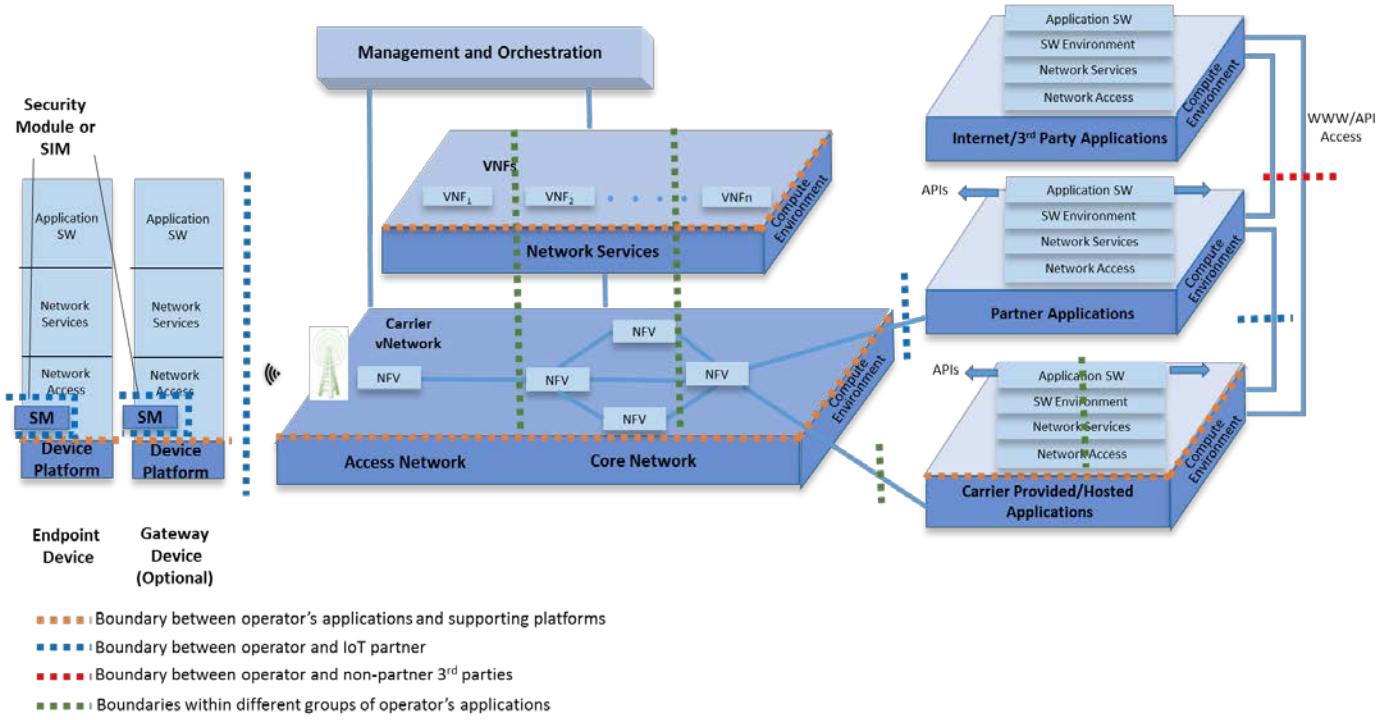

Figure 5.2: Overview of Trust Boundaries

In the IoT network diagram and the scenarios described in Section 3, several different entities are involved in delivering IoT services to the end user. Within this system the concept of a "trust boundary" may be used to indicate the places where entities at different levels of trust are connected together. Within a trust boundary all entities are treated with a similar level of trust. The term "trust domain" may be used to describe the entities within a trust boundary. At a trust boundary entities on either side of the boundary may be at different trust levels. To manage trust across boundaries and to provide layered security for the system, each trust domain should, within that domain, apply security controls to protect against possible attacks or other security threats originating from outside that trust domain. The trust boundaries are also places where authentication and authorization may be applied to ensure that only authorized entities gain access to certain resources and that the use of resources by entities outside the trust domain conforms to the permitted policy.

In the diagram different colors are used to show different types of trust boundary. The level of security protection and the range of permitted actions that may occur across a trust boundary will be influenced by the type of boundary. In this diagram we have focused on the perspective of the network operator looking outwards. Similar diagrams could be created to represent how other parties see the trust boundaries from their perspective.

Within a single network operator, the operator may choose to create several different trust domains – for example, high level user applications running in a cloud environment will likely be in a different trust domain to the SDN controllers for the carrier's virtual network. In this diagram trust boundaries within the operator's own network are shown symbolically by green dashed lines. How to structure these trust domains is beyond the scope of this document, but the overall purpose should be to ensure that critical or sensitive data and controls are well protected. The operator may define its own security policies within each internal trust domain.

Operator and third party services are built on top of hardware and software platforms that are typically provided by a platform vendor or integrator. The trust boundary between the network applications and this platform is shown by the brown dashed lines. In general, the operator has to place a high level of trust in the supporting platform because the platform underpins the operation of their network. Trust in the platform may be established by, for example, receiving from the vendor a security assessment of their platform or a certification that that platform conforms to a recognized security standard. Alternatively the operator may choose to independently audit their

platforms to establish their security. The management of platform dependencies is not specific to the IoT case and the details are beyond the scope of this document.

In Section 3, we describe an number of scenarios where a network operator delivers IoT services in a partnership with an IoT provider. Examples of where trust boundaries between the network operator and the IoT partner may fall are shown with blue dashed lines in the diagram. In this diagram we have assumed that the IoT devices (endpoint and gateway) are controlled by the partner though the security module or SIM in the device remains under the control of the network operator. The trust boundary between the network operator and the IoT partner is a critical security feature of these partnered IoT scenarios.

As discussed in Section 6, there are different security approaches which the operator may use with its IoT partners. In some cases there will be a high level of technical trust between the operator and the IoT partners (for example if a joint ARA process has been undertaken) and the trust boundary between the operator and the partner may permit the partner a high degree of latitude reflecting a high level of trust. In other cases the level of technical trust between the network operator and the IoT partner may be small and these partners may effectively only gain “arm’s length” access to network resources. In this case the trust boundary between the network operator and the partner should be tightly controlled within the network operator’s trust domain.

Finally, we have indicated that third party applications may gain access to data or services, for example using a public Internet API. These parties do not have a strategic relationship with the network operator and have merely met the conditions to be permitted API access. This may be as simple as registering an identity on a website. An example trust boundary for this access is shown by the red dashed line on this diagram. In this situation the level of trust is very weak to the point where it must be assumed that the external party could be a malicious or hostile actor. Security policies must be strictly applied to enforce control over what resources may be accessed and the transaction frequency (to prevent denial of service type attacks).

Table 5.1: Summary of Types of Trust Boundary

Trust Boundary Type	Level of Trust Across Boundary	Key Issues	Comments
Between trust domains within an operator (green)	Varies depending on the security design of the network.	How to partition network to provide best protection for critical and sensitive resources.	
Between the operator’s network and the supporting platforms (brown)	High by necessity.	How to ensure that the platform is trustworthy.	
Between the operator and an IoT partner (blue)	Varies depending on the level of trust established during the formation of the partnership.	What level of resource access does the partner need? How to authorize access and enforce security policies.	Key area for IoT
Between the operator and a non-partner 3rd party (red)	Low	How to protect the system from potential attacks while still providing using service access.	Key area for IoT

The table above summarizes the key points of the different types of trust boundary in terms of the possible level of trust and the key issues that relate to the positioning and the design of the trust boundary. As discussed, the two areas that are most strongly relevant for IoT and where IoT brings the most distinctive requirements are the boundaries between the operator and the IoT partners and between the operator and non-partner third parties that interact with IoT services.

We will now look at how this general model can be specifically applied to the scenarios described in Section 4.

5.3 Example Network Model for Scenario 1

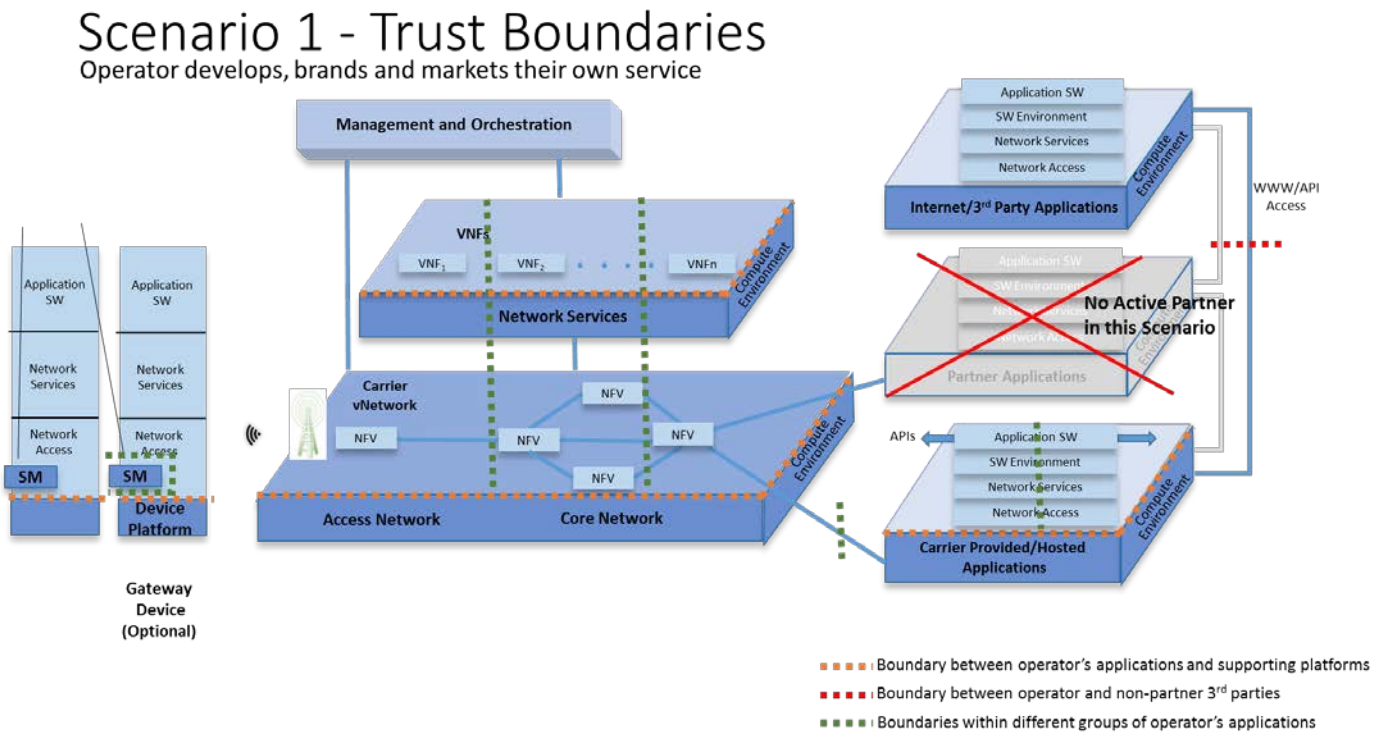

Figure 5.3: Scenario 1 Example Network Model

In this scenario there is no separate IoT partner involved. Instead the operator is responsible for the complete development, branding, and marketing of the IoT service. Hence, the IoT partner is removed from the diagram.

A key question in any IoT service is the security of the devices. Even though in this scenario the operator is responsible for the complete service delivery, it is likely that the operator will still obtain significant hardware and software components of the device from third parties or from open source components. In this diagram a trust boundary between the service's software elements and the device platform is shown. In reality the location of this boundary in the software stack will vary based on how much of the device's capabilities are inherited from the platform and how much has been designed as service-specific software.

Even if the devices are developed by the operator, they may for reasons of security or standards compliance choose to have a separate secure module in the device. Therefore these device security modules are included in this diagram even though both the device and the security module are under operator control. In this particular case the boundary between the device and the security module may be considered to be a domain boundary between different domains of the operator.

In this case the operator can undertake a system-wide ARA process to assess the security and plan attack mitigation strategies as defined in the [ARA process]. This analysis should encompass managing the security of platform dependencies, particularly on the devices. An output of this ARA process may be a decision to create different trust domains within the operator's own network for this service. These are shown symbolically by the green dashed boundaries in the diagram. As part of the ARA process the operator should generate a diagram that represents the actual service in their real network topology. This diagram may be used to show the trust boundaries for the specific service deployment.

Many services will provide public access to service data via an interactive web page or a public API. An example of a trust boundary for this type of interface is shown with a red dashed line in the diagram. These interfaces are inherently open to attack by malicious actors. Even for authorized users the trust level within the relationship is weak and attacks may originate both from authorized and unauthorized users. When performing the ARA analysis, the operator should pay particular attention to the risk of attacks arising from these interfaces.

In analyzing these network models the threat assets defined in the [ARA Process] were used. These are summarized in the figure below.

CON	Configuration Data	NFV	NFV Components
ICC	Inter-Component Communications	CSW	Customer Application SW
EUD	End User Data	COE	Compute Environment
API	APIs	EPD	Endpoint Device
STH	Service/Traffic Handling	PAS	Physical Assets, Power, Cooling
ENK	Encryption Keys	STF	Staff
IDB	Inter-Domain Admin Boundaries	BAR	Brand and Reputation

Figure 5.4: Summary of Threat Assets

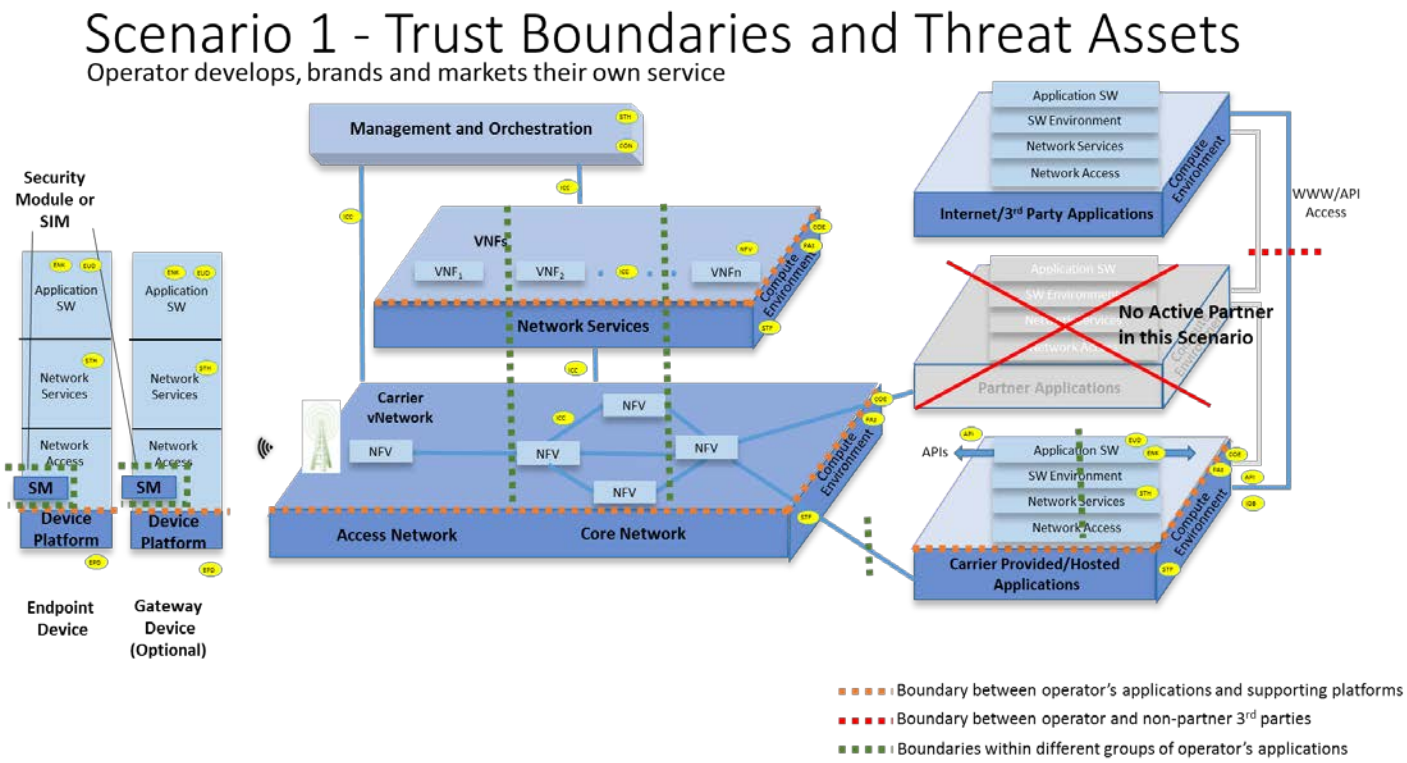

Figure 5.5: Example Threat Assets for Scenario 1

The diagram above shows how threat assets may be distributed in a scenario 1 example. This distribution is aligned with the case shown in the [ARA process]. As each service is different the operator's network design policy, the ARA process and the service-specific network diagram should be used to identify the precise location of threat assets for each service.

The explicit indication of a third-party API for in the IoT diagrams means that threat assets that are associated with this feature may be depicted.

One threat asset that cannot be localized to a particular place on the diagram is the threat to "brand and reputation". Damage that may occur to brand and reputation may arise from security failures at any point on the system. In this scenario the operator sells the service under its own brand and therefore the operator's brand is directly exposed to risks. When performing the ARA process the extent to which attacks may damage the operator's brand and reputation must be considered when scoring and prioritizing attacks.

5.4 Example Network Model for Scenario 2

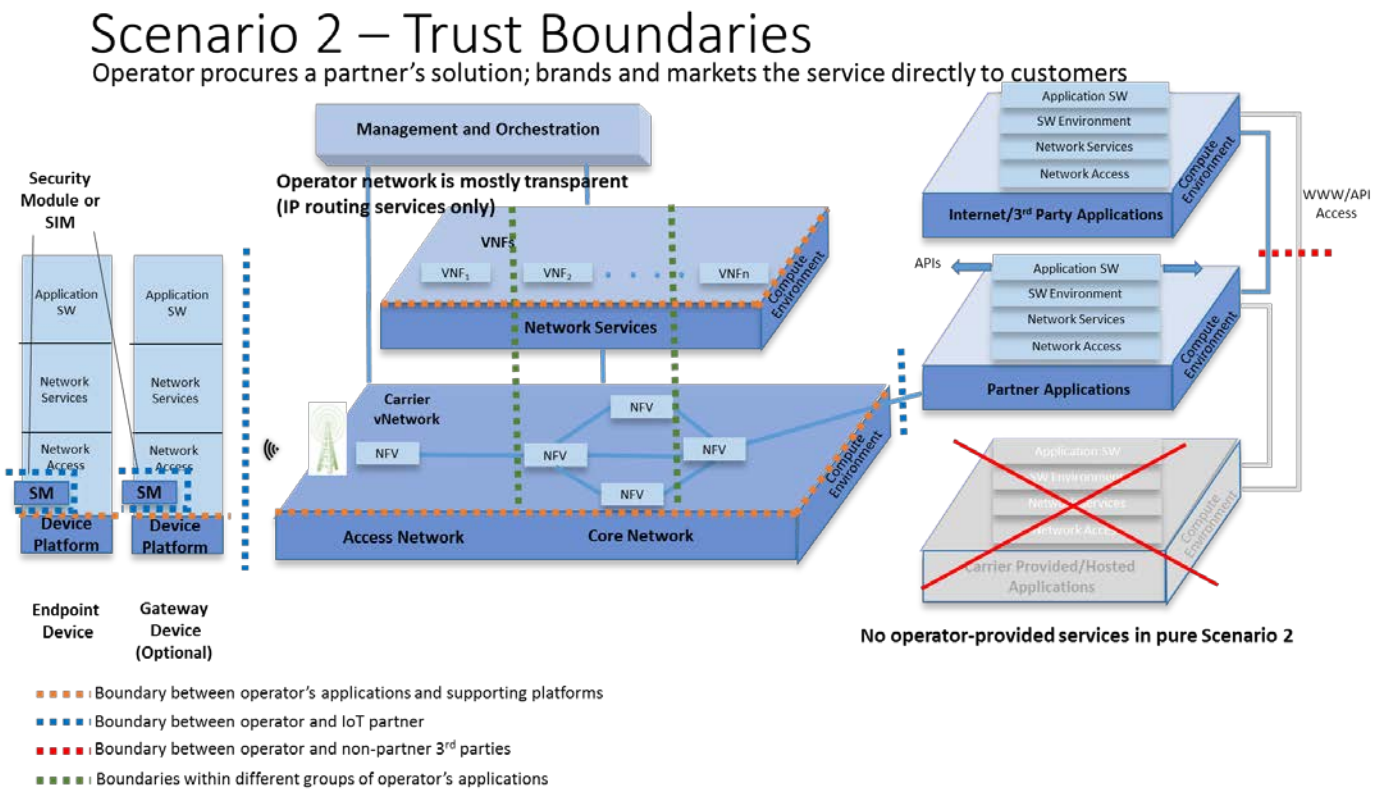

Figure 5.6: Scenario 2 Example Network Model

In this scenario, the operator procures an IoT solution from a partner but markets this solution under its own brand (e.g., using a "white label" offering from the partner). In practice in this scenario the partner's IoT solution will likely be concentrated in to two areas of the network:

- The IoT devices will be obtained from the partner.
- The IoT network application will be created by the partner and likely run in the partners chosen computing environment.

In this scenario, the technical role operator's network will predominantly be to provide connectivity between the IoT devices and the IoT network environment. The operator's network may have little or no customization to support the IoT service.

One point of control for the network operator in this scenario is via the SIM or other security module in the IoT device. The device gives the operator powerful control over the authentication of the device for mobile network access and the configuration of the device's communications stack.

Trust boundaries occur between the operator and the IoT partner at both edges of the network – as depicted by the dashed blue lines. Further the partner may offer open web access or APIs leading to the partner having a trust boundary with other third party applications (red dashed line). For the reasons explained in scenario 1, the security of the device and the security of any public API should be given special attention by the IoT partner when securing the service.

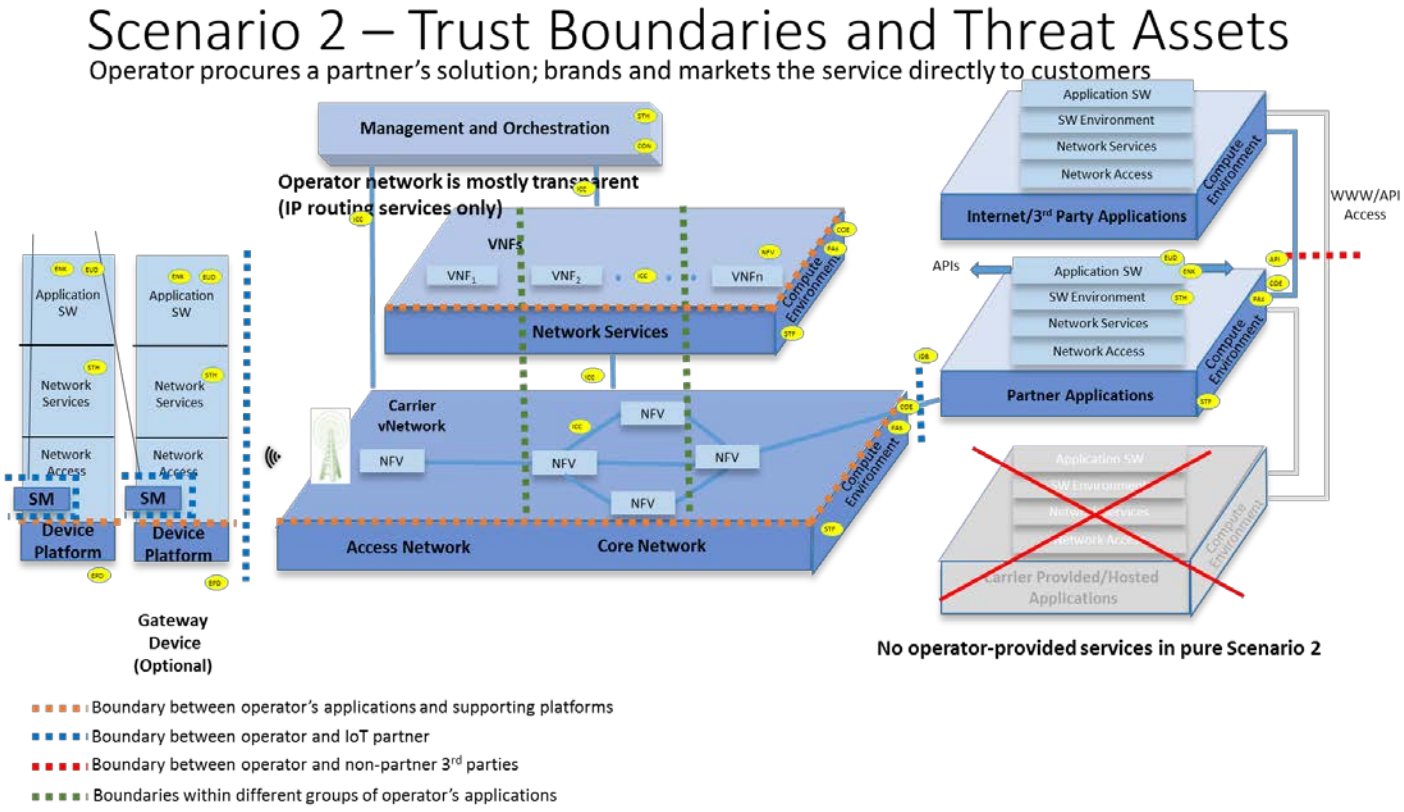

Figure 5.7: Example Threat Assets for Scenario 2

The diagram above show an example of the location of threat assets in this scenario. Notably many of the key service-specific threat assets are located in the trust domain of the IoT partner rather than the network operator. In this scenario, security failures that occur within the trust domain of the IoT partner can have direct impact on the operator's customers and the operator's brand reputation. To protect its assets, the operator must validate the security offered by its application partner. Methods to validate security when working with an IoT partner are discussed in Section 6.

5.5 Example Network Model for Scenarios 3 and 4

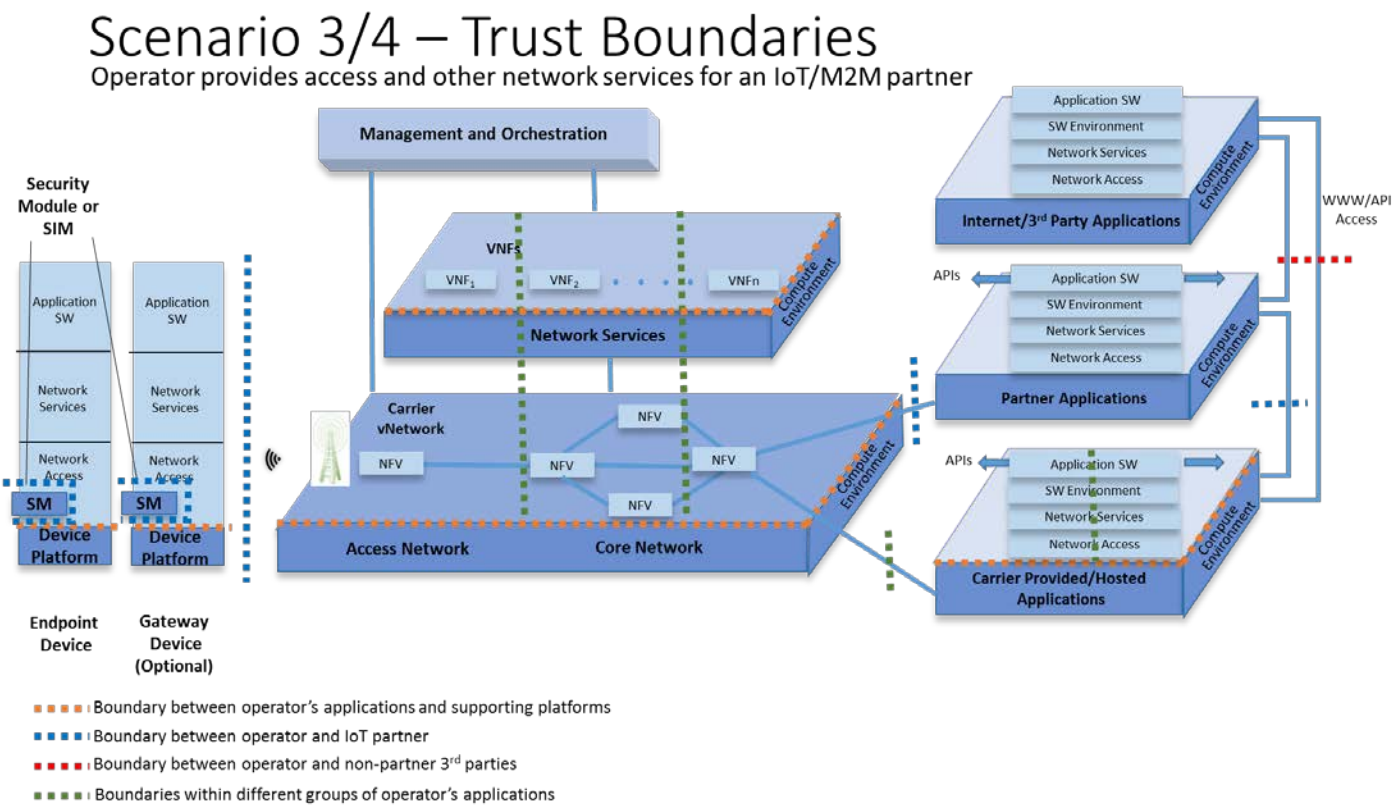

Figure 5.8: Example Threat Assets for Scenarios 3 and 4

These two scenarios are technically similar and are therefore depicted on using one common diagram. In both scenarios the technical provision of the service includes elements in both the operator's and the IoT partners' domains. In contrast to scenario 2, these scenarios include the possibility of a deeper technical role for the network operator – therefore the IoT-specific service elements are no longer solely contained in the partner's domain but may also be contained in the operator's domain. With all elements involved in delivering services, this diagram is essentially the same as the general diagram used to introduce this section. In these scenarios the trust boundary between the network operator and the IoT partner may occur in multiple places if functionality is distributed between both organizations.

In scenario 3 the services are co-branded, whereas in scenario 4 each entity separately sells the services under their own brand. In both scenarios customers of both entities and both brands are at risk in the event of a security incident. Damage to reputation and customers may occur in different ways, for example:

- Operators brand damaged by a security failure in a partner.
- Partners brand damaged by a security failure in an operator.
- Either brand may be damaged by a security failure of an underpinning platform.

The security analysis, for example using an ARA process, should include these risks under its consideration.

Both parties should establish mutual confidence in the security of the overall solution – for example by using an ARA process as described in Section 6.

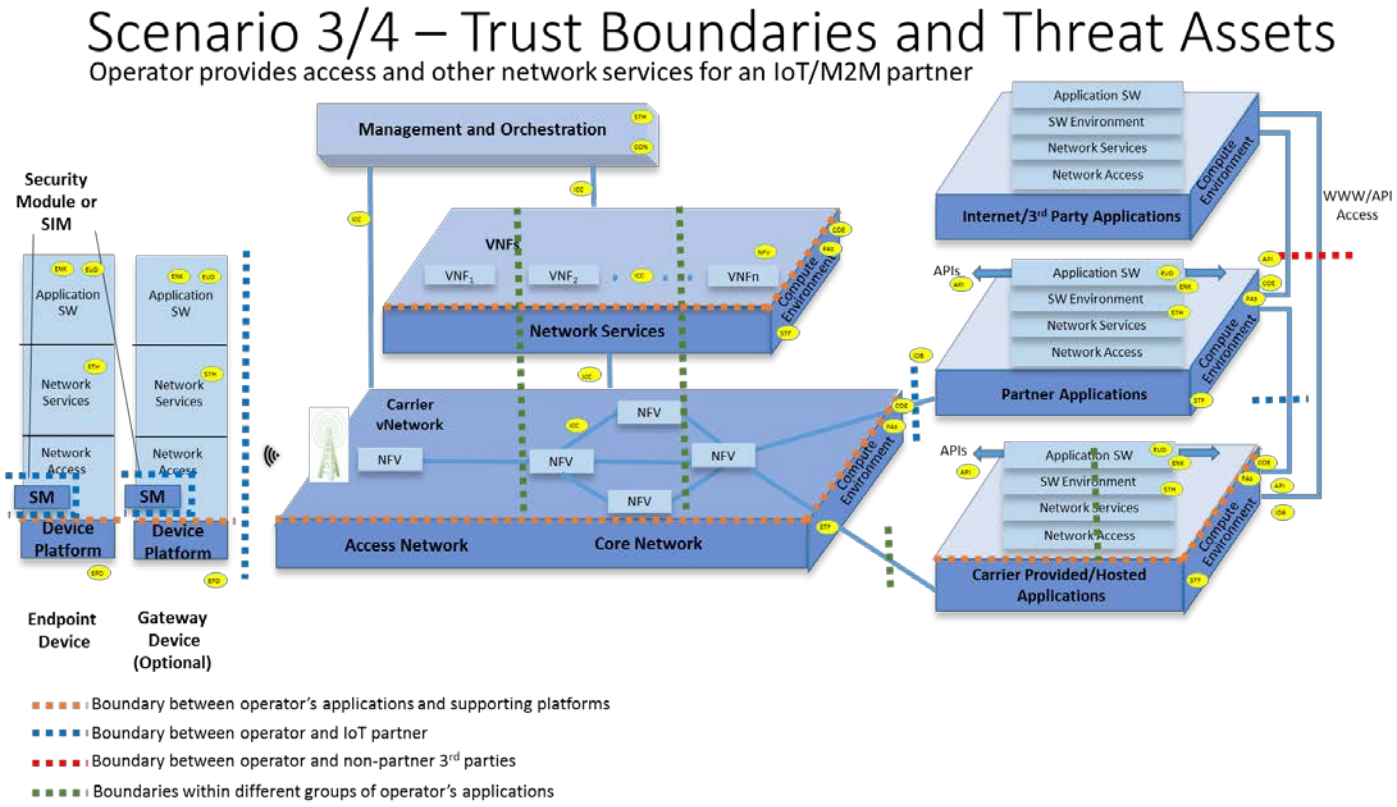

Figure 5.9: Example Threat Assets for Scenarios 3 and 4

In these scenarios the threat assets related to the service may be within the trust domain of the network operator or the IoT partner. The security analysis applied should ensure that risks are evaluated for the whole system and that assets are protected regardless of their location. Brand risks should be protected by:

- Secure design applied to whole system (including operator and partner elements).
- Continuing to validate the continued security of the design and applying secure operating procedures throughout the lifetime of the service.

6 Application of ARA Process to IoT Partner Scenarios

6.1 Introduction

The document [ARA process] defines a method of building security in to the design of ICT applications – particularly those delivered with the involvement of network operators. This section builds on the recommendations of the ARA process to provide additional information on how this process can be successfully applied to IoT applications – in particular IoT applications involving a partnership between a network operator and an IoT partner of the type described in Section 3.

6.2 Preparing for an ARA Process

Before starting an ARA process it is necessary to establish a context that supports the application of the process. This is particularly true in the case of a partnership scenario where there are multiple business entities involved in delivering the overall service. As discussed below, depending on the nature of the partnership there are different ways in which the ARA process can be applied.

6.2.1 Building security in to a partnership relationship

When establishing a partnership for an IoT service, security is one factor that should be considered at all stages of building the relationship. The partners should establish a mutual understanding of how they each view security and what security goals and practices they follow. Work to establish a security relationship should precede any contractual discussions between partners. At an early stage the discussions may establish an understanding of the general attitude and approach to security.

One question that both partners may wish to consider is whether they feel both parties in the partnership have similar priorities and outlooks on security issues. If one partner is significantly more security conscious than the other partner, then it is likely that this partner will want to consider what actions they can take independently or in conjunction with their partner to raise the solution security to a level they are comfortable with.

As the partnership matures, more specific and binding agreements on security-related matters should be formed. These may be formalized in to working agreements or contracts between partners. This step of formalization may take place after the ARA process described below has been applied.

The understanding of mutual security approaches established during early discussions in the partnership will help make decisions about how to apply the ARA process.

6.2.2 Define the service scenario under consideration

Section 4 of this report defines a number of scenarios that are representative of how network operators may work with IoT partners to deliver services. Before applying the ARA process to a partnered IoT service, it is necessary to consider how the partnership relates to the scenarios defined in Section 4. The nature of the scenario will inform the security objectives for the ARA process.

The product branding is a key aspect to understand in defining the scenario. Section 4 defines a number of models for how branding may be applied. Understanding the branding will define:

- The extent to which each brand's value could be impacted by security failures,
- Whose customers may be impacted by security failures and how, and
- How security failures in one entity may affect branding, reputation and customers belonging to a different entity.

As highlighted in Section 4.2, an important consideration in partner scenarios is that security failures in one entity can lead to consequences for brand and reputation that are primarily suffered by partner entities.

We have also shown how different scenarios lead to differences in the way that assets are distributed in to trust domains and which partner controls assets. This, in turn, implies that the scenario will also affect the technical parts of the ARA process and also how responsibilities for mitigation techniques may fall on the partners. In addition to the branding, the definition of the scenario should establish the areas of technical contribution of each party.

6.3 Methods to Apply the ARA Process to Partnerships

This section will describe two ways in which the ARA process could be applied in an IoT partnership scenario from an operator's point of view. The operator should choose the method that is most useful in their particular situation. The choice will depend on the nature of the security partnership established prior to the ARA process. This will include aspects such as:

- The degree of trust established between the partners.
- The willingness of partners to share technical details about their implementation.
- The willingness of partners to participate in a joint analysis of security.
- The degree or risk to which the operator's brand and reputation may be exposed.

The two methods of applying the ARA process are:

- The joint ARA process that fully involves both partners.

- The operator-only ARA process predominantly performed by the network operator.

6.3.1 Joint ARA process

In the joint ARA process, both the operator and the IoT partner participate in a single ARA process that spans the whole system. This requires a high level of cooperation and collaboration between the parties. This collaboration will involve sharing of detailed technical and operational information.

The advantage of the joint ARA process is that by considering the whole system it should deliver a better assessment of security risks and a better set of mitigation techniques than an operator-only ARA.

6.3.2 Operator-only ARA process

In the operator-only ARA process, the operator runs an ARA process with a primary focus on the aspects it controls. This allows the operator to apply the ARA process even if it does not have active collaboration from its IoT partner.

Even though the operator-only process is intended to be possible without collaboration of an IoT partner, it will generally be improved if some level of information sharing takes place. This could include the operator making a security audit of its IoT partner or receiving information from the IoT partner about its organizational security standards.

The operator-only ARA process offers the advantage of flexibility in that it does not require establishment of a detailed cooperation with the partner. However, the disadvantage is that the assessment and mitigation may be less comprehensive than in a joint ARA process.

6.4 Application of the ARA Process

The diagram below provides an overview of the ARA process. It consists of three “lanes” of related work topics. For each lane the key deliverables are identified on the diagram. The detailed explanation of the ARA process and its operation is provided in [ARA process]. It is expected that readers of this document will be familiar with at least the high level principles of the ARA process.

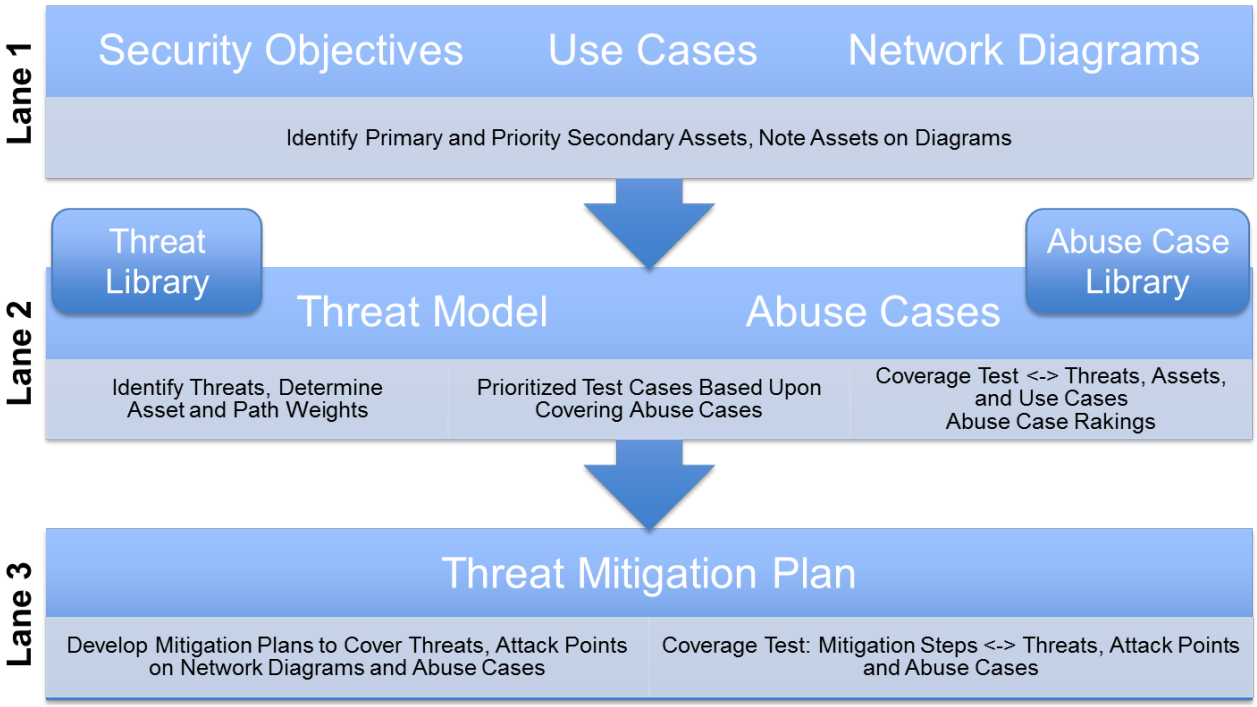

Figure 6.1: Overview of ARA Process

The following sections describe, for each lane, how to apply the ARA process to a partnered IoT scenario.

6.4.1 ARA Process Lane 1

The starting point of the ARA process is to define the security objectives for the process. A major input to this for a partnered IoT scenario is the preparatory work described in Section 6.2. When defining the security objectives, what interests are being addressed in this process should be considered:

- A joint ARA process should consider objectives that are relevant to both the network operator and the IoT partner.
- In contrast, an operator-only ARA process will primarily be focussed on the limited objective of protecting the operator's interests.

Typically, the network operator is expected to be the initiating partner in the ARA process. The operator should frame its security objectives by considering the following questions:

- What are risks to its brand, reputation, customers, and network?
- How much trust does the operator have in the IoT partner's level of security?
    - Will the IoT partner share the burden of security?
    - Should the operator minimize security dependencies on the partner and build the best security possible entirely within its domain?
- What knowledge and human resources does the operator require from the IoT partner to complete the process?

Once the security objectives are established, the use cases can be defined following the normal ARA process.

Lane 1 includes the creation of network diagrams for the solution. For a joint ARA process, the network diagrams should cover the whole solution spanning both the operator and the IoT partner domains. For an operator-only ARA, the IoT partner may have to be treated as a "black box" or covered in only limited detail. The ARA process acknowledges the possibility that is may be necessary to run the ARA process without having full information about all parts of the solution. The process document highlights that it is better to run an ARA process to the highest standard possible with incomplete information than to postpone the process while waiting for additional information that may never be forthcoming.

When developing network diagrams, these should explicitly include the trust domains and the trust boundaries as explained in Section 3 of the report. The diagrams should:

- Identify location of boundaries between trust domains
- Identify the type of each boundary (e.g. intra-operator, between an operator and an IoT partner)

6.4.2 ARA Process Lane 2

The normal ARA process can be applied to assess the threats sources and abuse cases. Two points of considerable importance to IoT scenarios should be considered:

- In IoT, the devices can be a particular source of vulnerabilities and threats due to this source should be especially considered.
- Similarly, exposure points to the public Internet, including interactive web pages to access IoT services and public APIs, are often vulnerable.

When considering the threat model the relationship to the IoT partner can be two pronged:

- The IoT partner could/should help in identifying and mitigating threats, though this may be limited in an operator-only ARA.

- Threats may originate within the IoT partner's trust domain if its security is compromised: Particularly for an operator-only ARA, the IoT partner's domain may be considered as potentially hostile in some cases (e.g., botnet IoT attacks based on IoT devices).

Similarly, from the IoT partner's point of view, an ARA process may consider the case where the operator's trust domain is compromised and becomes a threat to the IoT partner's trust domain.

An important part of lane 2 is the creating a mapping of assets on to the network diagram. When identifying assets, also identify which trust domain contains the assets. In some cases, the domain owner for an asset may not place the same security value on the asset as their partner does. In this case there is a potential conflict of interest between the partners in terms of the effort that should be applied to secure the asset.

The threat analysis should consider threats that originate within the same trust domain as the assets as well as threats that cross trust domain boundaries.

6.4.3 ARA Process Lane 3

The scope of the mitigation plan will depend on the type of ARA process being applied. In the case of a joint ARA process, the mitigation plan will cover both the operator and the IoT partner. In the case of an operator-only ARA process, the practically possible mitigations may be limited to the network operator only. However, in some cases, the network operator may still decide that threat mitigation is so important that even for an operator-only ARA they will negotiate additional agreements with the IoT partner to undertake certain security mitigation steps.

The use of multiple trust domains provides the possibility for layered mitigation techniques where different trust domains may contain independent mitigation processes against particular threats. This layering of mitigation techniques means that even if the security of one trust domain is compromised, the system still has possible protection in place in other trust domains.

The discussion on lane 2 of the ARA process highlights the possibility of attacks that originate within the trust domain of an IoT partner (e.g., if an IoT device has been compromised and becomes part of a DDoS botnet). Mitigation techniques should consider attacks that originate within partner organizations as well as outside the partnership.

Mitigation techniques should include organizational and operational techniques within the partnership to provide an on-going platform for supporting security objectives. These may include techniques for:

- Handling of new threats and service/technology changes.
- Joint responses to security incidents.
- Certification, training and authorization of personnel in both oganizations with access to security-related systems.
- Intrusion and threat detection.

7 Partnered IoT Security Analysis Template

The following sections provide a template that can be used as part of a security analysis for an IoT service involving a network operator partnering with an IoT provider. This template may be useful when applying the ARA process as described in Section 6. This template **does not attempt to cover all aspects of the ARA process that are described in [ARA process] but focuses on the aspects that are called out in this document as especially important to the IoT partner scenarios**.

By posing some of the key questions, this template encourages the use of a systematic approach to security analysis that takes account of the particular circumstances in each partnership. This template does not attempt to identify "right" answers, as every scenario must be assessed on its particular merits and according to the objectives and processes of the organizations involved. Therefore, this template should be used to develop particular security objectives for each partnership and to assess solutions and mitigation techniques against those objectives.

While some parts of this template may be shared between the partners, it may make sense for other parts to be managed confidentially within one partner to allow information to be captured candidly.

7.1 Preparing for an ARA Process

Question	Answer
Information about this analysis	
What is the product that this analysis relates to?	
Information about the approach to security in the partnership	
Which companies or organizations are partnering to deliver this product?	
What prior discussion of security issues has taken place between the partners and what conclusions have been reached?	
How would you characterize your partner's approach to security? How well aligned are you and your partner's security approach and objectives?	
Information about product branding and risks to brand value, service delivery partners, and customers	
How will this product be branded and sold (i.e., who does the user perceive the product is supplied by)?	
Which model from Section 4 best represents this partnership (if none of the models are applicable describe the partnership in your own terms)?	
To what extent could the brand value of each partner by impacted by security failures?	
Whose customers may be impacted by security failures and how?	
How may security failures by one entity affect branding, reputation, and customers belonging to a different entity?	
Method to apply the ARA process or other security analysis	
For this analysis are you performing the ARA, or other security analysis, yourself independently (operator-only ARA process) or jointly with your partner (joint ARA process)?	

7.2 ARA Process Lane 1

Question	Answer
Information about the objectives of the analysis from the operator's point of view	
What are risks to your brand, reputation, customers and network?	
How much trust do you have in the IoT partner's level of security?	
Will the IoT partner share the burden of security?	
Do you wish to minimise security dependencies on the partner?	
For a joint ARA process only: What are your partner's security objectives?	
What knowledge and human resources does the operator require from the IoT partner to complete the process? How can this information be obtained?	
Describing the solution and trust boundaries	
Show the network diagram for the solution. If full information is not available the IoT partner may have to be treated as a "black box" or covered in only limited detail.	
Where are the boundaries between trust domains on this diagram?	
For each trust domain boundary, identify the type of boundary (e.g., intra-operator, between an operator and an IoT partner).	

7.3 ARA Process Lane 2

Question	Answer
Device threats In IoT, the devices can be a particular source of vulnerabilities and threats due to this source should be especially considered.	
What possible security threats may originate from the device in this IoT product?	

Question	Answer
What security credentials (e.g., for user access, network access or service access) are stored on the device?	
How are security credentials for users and service access established and managed? Are weak, shared, or hard-coded credentials used in any context?	
Can security credentials be stolen and reused on other devices to steal services or data?	
What other sensitive data may be stored on the device?	
What services (e.g., a local web server, open IP ports, etc.) are exposed on the device?	
What LAN, WAN, and other electrical interfaces are available externally and internally (e.g., debug ports) on the device?	
What security protection is applied to prevent eavesdropping or tampering with signalling between the device and the infrastructure?	
Services exposed to the public Internet Exposure points to the public internet – including interactive web pages to access IoT services and public APIs – are often vulnerable.	
What services does this product expose to the public internet?	
How is access to these exposed services secured? How may this security be breached?	
How may these services be attacked (e.g., to prevent access to services, steal data, or create disruption)?	
Threats from the partner domain	
What threat may originate from the IoT partner's trust domain if its security is compromised?	

7.4 ARA Process Lane 3

Question	Answer
Mitigation of device threats	
How are security credentials protected from theft or brute-force attacks?	
How is the client software protected against common attacks (e.g., malformed packets/data, buffer overflows, SQL injection, XSS)?	
How is software on devices maintained and updated?	
How is the software update process protected against attacks?	
How can compromised devices be detected and repaired?	
Does the device use a recognized IoT framework with security capabilities (e.g., oneM2M)? If so, which security capabilities from the framework are used?	
Layered protection	
What security protection is applied between each communicating pair of trust domains (especially between the domain of the network operator and of the IoT partner)?	
If one security domain in the solution is compromised, how can the remaining domains continue to protect the security of assets?	
If the security of a partner is compromised, what protection is available from the security domains under your control?	
Organizational and operational aspects in the partnership	
How are new threats and service/technology changes handled?	
How will the partners jointly respond to security incidents?	
How are personnel certified, trained, and authorized to access security related systems?	

ATIS-I-0000056

Question	Answer
How are intrusions and threats detected and communicated between partners?	