IoThink Thought Leader Series

Security solutions and services for the IoT

Security is a major industry concern that could significantly slow IoT market growth. IOT security is a multi-layered problem with the added complexity of practical implementation challenges arising from supplier diversity and legacy systems. Some of the challenges which create the need to ensure security and trust across communications between multiple service providers include: a huge diversity of use-cases requiring different levels of security protection; constrained-application devices that lack the resources to support strong security solutions; and, interoperability scenarios.

The global oneM2M standard provides a horizontal architecture for common, IoT application enablement services, such as security. A key benefit of oneM2M's horizontal framework is its applicability across different vertical markets. The potential to design once and re-use many times is an enticing prospect for developer communities that seek to maximize their addressable market opportunity.

Another feature of the oneM2M architecture is to concentrate common, IoT-enablement services in a oneM2M Service Layer. This makes application development independent of the underlying communications infrastructure by abstracting the different network technologies that support any given IoT application. Because of this, it is the service layer, and not the application layer, that handles network management issues; developers can focus on the business logic for IoT-applications.

While end-to-end authentication and secure tunnelling are the ultimate services expected by IoT applications, a fundamental requirement is to manage the diversity of communication links that underpin different IoT applications. The oneM2M Service Layer employs a 'hop-by-hop' strategy to bridge the communications between all device-types in an IoT application. This ensures a consistent path across an extended communications chain at the networking layer of the oneM2M architecture.

By embracing a standards based approach within oneM2M's architecture, each IoT application can set its own security policies while still being able to interact with other IoT applications. This ability to re-use IoT enabling services combined with the sharing of platform costs among multiple users underpins long-term economies of scope and scale. Reusability also opens up a new world of opportunities by supporting interactions between applications from different providers. It also means that IoT service providers and enterprise users no longer have to create multiple security solutions for each and every application they deploy.

LEAD AUTHORS:

François Ennesser (Gemalto) **Yogendra Shah (InterDigital, Inc.)**

February 2016

The **oneM2M IoThink Series** features business insights about the IoT market and is produced as a collaborative initiative involving companies that are commercializing products and services based on the oneM2M standard.

Introduction

The deployment of effective security strategies will have a profound impact on how the Internet of Things market grows. The current market sentiment towards IoT security is one of grave concern. This stems from: the lack of widespread knowledge about IoT security threats; indiscipline in addressing potential security threats at the design stage; and the absence of a common framework to implement IoT security solutions.

Compared to traditional Internet security, it is important to note that security breaches in IoT systems can ultimately impact human safety. This may occur, for example, if system designers and service providers fail to secure actuators and their related control systems.

Traditional IT systems encompass relatively homogenous computing devices (mainly Personal Computers and server systems) organized in relatively stable architecture models. IoT systems, by contrast, encompass a huge diversity of devices and a multiplicity of architectural models both of which contribute to increased system complexity. Thus, while traditional IT system design is an important source for security best-practices, IoT security raises additional challenges. These arise from the multiple layers of complexity and interactions between silo systems within open eco-system operational systems.

In light of these technical and commercial security challenges, this paper firstly sets out to foster a common understanding of IoT security issues. Secondly, the paper describes oneM2M's IoT-enabling horizontal platform and common services framework, which addresses the individual security requirements of different IoT applications.

oneM2M is a global standard that was set up precisely with the aim of addressing the service-support requirements for IoT applications. Security is one of several key enabling services within the oneM2M design philosophy. By making security an intrinsic element of the oneM2M architecture, IoT application developers and service providers are able to collaborate within a common security framework to enable multiple IoT applications using shared service resources. This has a direct economic benefit in relieving application developers from having to implement multiple security solutions for each and every application they create.

The key topics covered in this paper, with a focus on security, are as follows:

- Characterizing security in the IoT context
- An illustration of the IoT security problem
- oneM2M's standardized architecture to enable IoT applications
- oneM2M's 'hop-by-hop' security strategy
- oneM2M security features and long-term road-map
- Benefits and issues for adopters

Characterizing security in the IoT context

The vast number of potential IoT applications represents a huge diversity challenge. While this is evident from the multitude of possible use-cases, there are practical diversity challenges created by the types and permutations of different sensors, devices, gateways, software applications and their respective interactions.

Many of these interactions occur autonomously, oblivious to human users at the edge of the IoT. They may exert direct control over the behaviour of powerful machines and have critical impact on our environment, in ways that potentially affect human safety. While Internet security is already an everyday challenge today, attacks rarely lead to such drastic outcomes since human users can spot anomalous and malicious behaviour and can intervene to thwart such attacks.

Another important characteristic of IoT applications is that many sensors and devices have limited resources in terms of computational processing capability, power, bandwidth, and memory. These 'constrained application' devices have a limited capability to support security-related functions. There is therefore an issue in harmonizing a system-wide security solution in the operational deployment of such devices. Then, there is a compounding effect in situations where applications have to operate with sensors and devices from different suppliers and product generations.

A further complication arises in the case of interoperability between traditional IoT application silos. These may occur across multiple service provider networks or involve interactions between two or more silo applications. The example of a smart city application involving utilities, transportation and health service providers typifies many IoT scenarios and the need to harmonize security roles across different service providers.

To illustrate this situation, consider security for a traveller transiting an airport. One service provider, the airline, will conduct its own electronic, passenger checks. A different service provider, airport agents, will conduct physical inspections of passengers and their luggage. And, yet a third service provider, the ticketing agents, will perform identity checks at the point of boarding. While these security checks share a common purpose, organizational needs, local conditions and resources are what determine specific actions at each stage. This scenario is analogous to the security solutions needed for IoT applications where a shared security purpose can be satisfied by collating and addressing the common security requirements in a harmonized manner.

It becomes significantly onerous to exchange information between different applications if each individual application addresses its own security needs (e.g. using encryption), independently of the others. For this reason, IoT scenarios that involve multiple applications, such as smart cities, require a federated security architecture that sits between the network and application layers.

An illustration of the IoT security problem

Beyond the coordination issues between multiple IoT stakeholders, security designers also need to consider application-related security threats such as: the potential hijacking of connected devices; theft or mis-use of sensitive data; and the threat to humans (life and limb) and property from compromised IoT applications.

To illustrate the different sources of risk, consider the following example of a connected-health application, in Figure 1, comprising a set of personal wellness monitors in a user's home. These monitors communicate via an in-home gateway, through a service proxy server to store data at a database-hosting service provider. The data is now accessible to an eHealth monitoring application in a physician's workplace. This type of ecosystem, involving multiple roles and players, is typical of IoT applications and requires the deployment of technologies that all stakeholders can trust.

Figure 1 Security threats in the IoT: remote health- and wellness-monitoring example

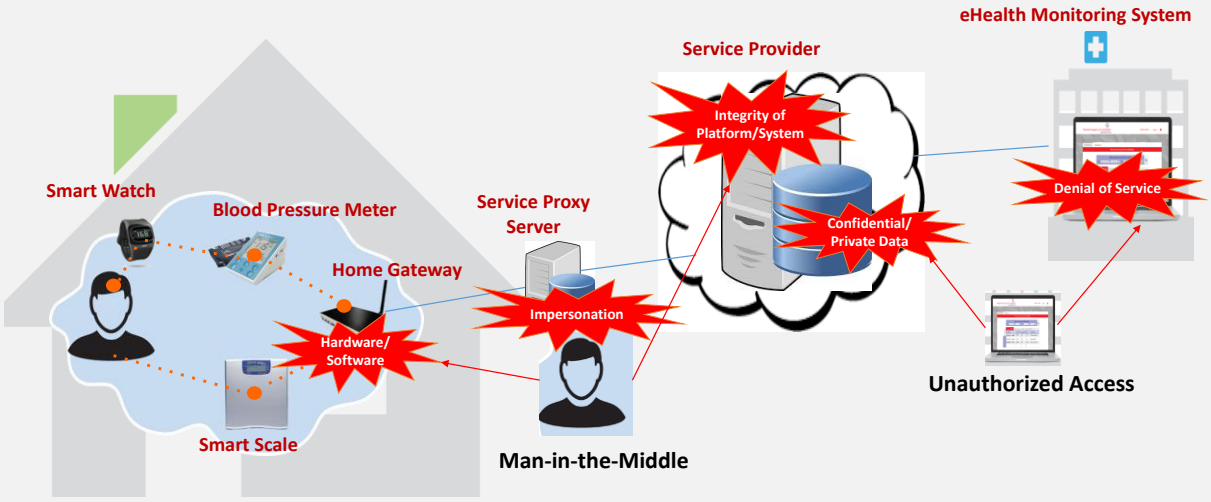

This arrangement is susceptible to multiple security threats including:

- Physical tampering with hardware and software attacks on the personal wellness monitoring devices or home gateway.
- Impersonation in the communications channel via a man-in-the-middle style of attack to intercept communications or alter the information stream.
- Insufficient integrity in the communications and data hosting platform/system.
- Unauthorized access to private or confidential data at various points, including intermediate entities and application endpoints, within the overall eHealth solution.
- Denial-of-service attacks to overload components of the end-to-end solution and deny or degrade services to legitimate users.

The specific architecture of each IoT application varies from one use-case to another so this is not an exhaustive list of security threats. oneM2M has studied[1] several use-cases to identify different threats and characterize the requirements for its standards development activities relating to IoT platform services.

oneM2M's standardized architecture to enable IoT applications

oneM2M is a global standards initiative to define a comprehensive IoT service layer solution to enable scalable and economic IoT applications. It aims to accelerate development and re-use of M2M/IoT data and components across a diverse set of vertical applications, networks, and devices. In early 2015, the oneM2M standardization body made a first release of its specifications for an IoT service layer platform.

oneM2M's platform architecture consolidates the essential components of any IoT application into a three layer model to ensure a consistent and modular framework for IoT application developers and users. To understand its benefits, consider the illustration in Figure 2, of a conventional, silo arrangement comprising two standalone applications. Devices and sensors related to Application#1 interact with their back-end business application via a dedicated middleware platform (MW1, Local and Remote) and underlying

[1] oneM2M Technical Report, Analysis of Security Issues, oneM2M TR-0008 (April 2014)

communications networks (local- and wide-area). An identical arrangement applies to Application#2 which relies on its own middleware (MW2, Local and Remote) capabilities.

Many of the middleware capabilities, such as device management and security, could be common. Indeed, they would benefit from being common to enable interaction via a shared infrastructure, from one IoT application to another. Recognizing this characteristic creates the opportunity for standardization of a horizontal platform architecture comprising applications, services and network layers.

Figure 2 oneM2M's horizontal architecture provides a unifying framework for silo applications

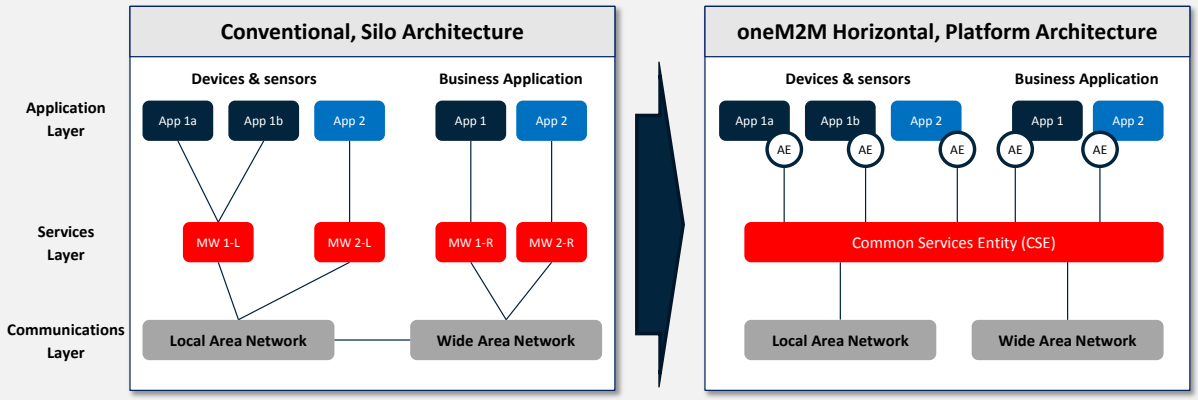

The oneM2M horizontal platform architecture illustrates a consolidated middleware layer designed to support resource sharing and interoperability. At the application layer, there is an Application entity (AE) within each IoT device or sensor, gateway and cloud server. These AEs have to deal with a wide variety of devices, each of varying capabilities. In some cases, they also have to interact with server platforms (called "nodes") that operate as hubs for multiple end-point devices or sensors. Some of these end-points may possess little or no security capabilities while others may provide more sophisticated security features.

In addition to new devices with advanced security features, IoT applications often need to accommodate legacy devices which are unlikely to include modern-day, IoT design and operational features.

The middleware service layer entity within the oneM2M architecture is the Common Service Entity (CSE). It resides within each server or device platform and provides a standardized interface. This allows applications to access commonly required M2M/IoT services such as store and forward, discovery of data across applications, notifications, and interactions with other applications. The CSE exposes an easy to use Application Programming Interface (API) to the Application Entities, thereby facilitating the development of oneM2M compliant applications.

At the bottom of the three layer model, the oneM2M standard abstracts connectivity technologies to enable the development of solutions that are agnostic to the communication networks, while still benefitting from specific capabilities exposed by the underlying network.

The standardized service layer, between applications (where data processing occurs) and networks (where communications capabilities reside), reduces the interoperability burden on the application layer. As such, this arrangement accommodates plug-in capabilities including security-management modules, for example. This offers a means of providing a security service to application developers, relieving them from the intricacies of implementing security in their applications. An additional benefit of oneM2M's service layer concept is that it facilitates interoperability with existing technology modules such as device management, for example.

By virtue of an open standard approach, oneM2M also ensures that developers can count on a competitive eco-system of security solution providers. As a result, they have greater latitude to avoid locking in to a proprietary single-vendor solution.

oneM2M's 'hop-by-hop' security strategy

Given the diversity of communication links that may be involved, the ability for IoT applications to perform a complete authentication and setup of a secure end-to-end channel or tunnel is not always as practical as it is for conventional communication systems. To alleviate this issue, the oneM2M architecture relies on propagating messages to a destination node through one or more intermediate nodes, over an established service layer session. The term that best describes this is 'hop-by-hop' communications.

As illustrated below, a hop is a client-server form of communication which occurs between two adjacent service layer entities, as in the case of the communication from AE1 to CSE1 in Figure 3.

Figure 3 oneM2M's hop-by-hop architecture

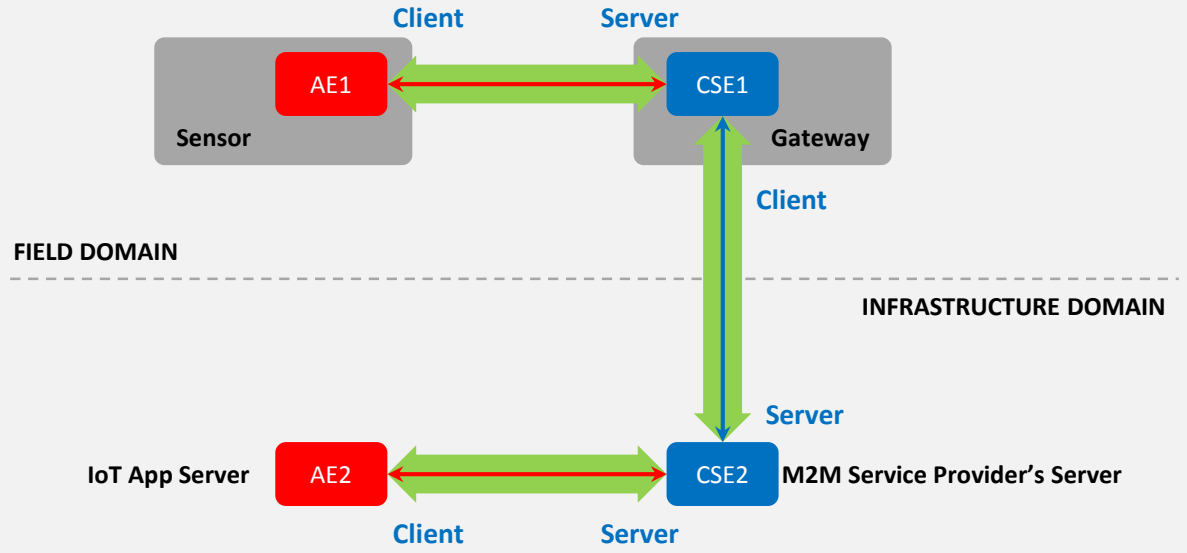

Additional intermediate nodes add to the number of hops required to reach a final destination node such as the communications from CSE1 to CSE2 and CSE2 to AE2. These communications sessions, at the application-layer, coexist with other service-layer sessions. The latter handle the overhead functions involved in establishing and maintaining the communications sessions, removing this responsibility from the applications themselves.

The oneM2M design philosophy allows a service layer session to operate on top of one or more underlying access network communication sessions/connections. This allows a service layer session to persist and continue to operate independently of the setup and teardown of lower layer underlying network sessions/connections.

The service layer functionality in the intermediate nodes handles all data communications including setup of secure communications. It achieves this by establishing a secure connection between adjacent nodes in a communications session. The service layer functionality on the co-operating intermediate nodes relies on a trusted entity to deliver an effective end-to-end security arrangement. Datagram Transport Layer Security (D)TLS provides the secure communications channel between the client and server for each 'hop' in a communications path. (D)TLS is a widely deployed protocol, which employs a handshake mechanism to exchange various parameters needed to establish a secure connection between the client and the server. After the handshake has occurred, a secure connection is established which prevents eavesdropping or tampering with any part of the encrypted communication between client and server.

A DTLS session provides for secure communications over UDP[2]. Alternatively, a TLS session applies in the case of secure communications over TCP[3]. Together, these methods provide flexibility and interoperability over a wide variety of underlying access networks.

oneM2M security features and long-term road-map

To accommodate the wide range of devices and deployment scenarios faced in the IoT in the most cost effective manner, oneM2M provides considerable flexibility for the implementation of service layer security features. This ranges from the initial provisioning of security credentials to providing authentication infrastructures as well as confidentiality, integrity and authorization services.

The service layer concept turns these features into shareable capabilities across multiple stakeholders and applications. It also includes the flexibility to implement security in a manner that reflects the capabilities of two adjacent communication-hop nodes but in a consistent manner.

The oneM2M Security Solution specification, TS-0003 (available at www.oneM2M.org) contains full details of the security features and protocols used for provisioning, authentication, authorization and establishing secure communications. The specification provides a variety of interoperable solutions (from device provisioning to centralized remote administration assisted by a oneM2M Enrolment Function) to deploy and manage credentials securely while accommodating both symmetric and asymmetric cryptographic solutions.

In addition, oneM2M intends to define service layer requirements for secure environments in IoT instances. These may be implemented as Hardware Security Modules (such as UICC), embedded Secure Element or a Trusted Execution Environment in a device processor. This will enable each stakeholder in an IoT ecosystem (network operators, service providers, and application providers) to manage and control their private credentials and security policies independently of each other by means of isolated storage and execution resources on the nodes.

oneM2M security services also enable each IoT application that creates or manipulates application information to apply fine-grain policies to control access to their data (i.e. for each piece of data, it is possible to specify WHO is allowed to do WHAT operation under WHICH context). A practical example from the healthcare scenario above is that the IoT application used by a physician at a specific place may be able to create or update prescriptions while the patient or other physicians would be limited only to reading the prescription information.

oneM2M standardization activities are part of an ongoing process punctuated by specific releases of the standard. The public launch of Release 1 occurred in January 2015. As far as security features are concerned, the table below summarizes the road map for different releases of the standard.

[2] User Datagram Protocol (UDP) is a common Internet protocol, typically used for streaming audio and video content. It offers speed because there is no form of flow or error control

[3] Transmission Control Protocol (TCP) is the most commonly used Internet protocol. It offers error correction, in part by determining when there is a need to re-send data to guarantee delivery.

Table 1 oneM2M security standardization road-map

Release 1.0	Release 2.0	Release 3.0
Q1-2015	**Planned for Q2-2016**	**Not yet confirmed**
• Addresses authentication and secure connection establishment at the service layer by means of (D)TLS using a pre-shared Key (PSK), raw public key or public key certificate. • Security credentials are provisioned, through trusted third-party entities, and relying on security associations from existing trust relationships. This may use local or remote mechanisms. • Authorization is performed using Access Control Policies (ACP) based on Attribute-based access control rules	• Will provide mechanisms for end-to-end message authentication and confidentiality. • Additionally, data confidentiality and integrity for both data at-rest as well as data in-transit. • Scalable authorization using well established dynamic authorization techniques that have been adapted for IoT. • Work is underway to standardize a secure abstraction layer that provides access to a Secure Environment (SE) via an API.	• oneM2M plans to address further security features that enable additional layers of security that more closely align to device capabilities.

oneM2M benefits and issues for adopters

oneM2M's standards release cycle demonstrates a long-term commitment to create a standard to enable IoT applications that proprietary and *de facto* approaches cannot rival.

The oneM2M service layer provides a federated architecture and a trust infrastructure (oneM2M Enrolment Function) to deploy security credentials as well as authentication (oneM2M Authentication Function) and authorization servers for M2M/IoT applications.

The common services provided by the oneM2M service layer relieves applications from having to worry about services such as setting up and administering communications links and security implementation. This is an important benefit for the developer community because it frees developers to concentrate their efforts on developing IoT applications. It also allows developers to focus on value added services and interact with other applications through the features provided by the service layer.

The standard enables all stakeholders in the M2M service chain (network providers, service providers and application providers) to leverage the security assets exposed by lower layer participants to address their specific security needs in the most cost effective manner. This has a direct commercial impact by making it possible to share the costs of deploying and preserving security, including personalization and deployment of trusted security anchors, with all the actors in the value chain.

About oneM2M

oneM2M is the global standards initiative that covers requirements, architecture, API specifications, security solutions and interoperability for Machine-to-Machine and IoT technologies. oneM2M was formed in 2012 and consists of eight of the world's preeminent standards development organizations.

oneM2M Standards Development Organization (SDO) Partners	
ARIB (Japan)	**ATIS (N. America)**
CCSA (China)	**ETSI (Europe)**
TIA (N. America)	**TSDSI (India)**
TTA (Korea)	**TTC (Japan)**

These SDO Partners collaborate with six industry fora or consortia (Broadband Forum, Continua Alliance, GlobalPlatform, HGI, Next Generation M2M Consortium, OMA) and over 200 member organizations to produce and maintain globally applicable, access independent technical specifications for a common M2M/IoT Service Layer. oneM2M specifications provide a framework to support applications and services such as the smart grid, connected car, home automation, public safety, and health.

oneM2M actively encourages industry associations and forums with specific application requirements to participate in oneM2M, in order to ensure that the solutions developed support their specific needs. For more information, including how to join and participate in oneM2M, see: www.onem2m.org.

oneM2M Eco-system

Gemalto and **InterDigital** are part of a growing eco-system of solution providers. The purpose of this collaborative article is to promote the oneM2M standard and to foster a competitive marketplace for oneM2M platforms, solutions and services.

Francois Ennesser (Gemalto) and Yogendra Shah (InterDigital) have co-authored this article on IoT security based on their respective areas of expertise in this domain.

François Ennesser is acting as chairman of the Security Working Group within the oneM2M Partnership since its inception in 2013. He has been involved in M2M / IoT security standardization and associated privacy related regulations since 2008. He is an employee of Gemalto, a global company developing products and services that enable its clients to offer trusted and convenient digital services to billions of individuals.

Yogendra Shah leads InterDigital's team focusing on developing security solutions and technologies to address the challenges and problems introduced by the IoT. Yogendra has been developing security technologies at InterDigital since 2006 and was driving InterDigital's contributions on security at ETSI TC M2M. InterDigital has been participating in oneM2M since the outset and Yogendra advises InterDigital's standards team on providing contributions to security standardization at various standardization bodies including oneM2M.