NOKIA

Using light device management to secure endpoints in the Internet of Things

Strategic White Paper

Connected devices are different within the Internet of Things (IoT). They are vast in number, often unmanned, and require repeated software and hardware upgrades over typically long life spans. Many of these devices will use different components and have different owners over time. What's more, cellular operators are creating separate IoT networks that present unique characteristics and tariff plans.

These differences have created a host of new security threats that, when combined with the expected scale of the IoT, can make devices costly and difficult to manage. Threats must be properly addressed as part of each device rollout. The use of a distinctly independent Lightweight Device Management client secures the software and hardware on the device in a cost-effective way that can be scaled to large populations of devices. It allows devices to be monitored, controlled and, in case of a breach, effectively remedied in accordance with the policies of the relevant security authority.

This paper offers insights that will help security authorities understand the growing security concerns within the IoT space and the benefits that Lightweight Machine-to-Machine (LWM2M) device management provides relative to fixed and mobile devices.

NOKIA

Contents

NOKIA

Introduction

The Internet of Things (IoT) offers unique concerns relative to security. In some instances, the challenges involved in managing mobile phones and personal computers are similar to those presented by IoT devices. For example, each of these technologies can experience compromises of firmware.

The difference with IoT is one of scale. Gartner forecasts indicate that 4.9 billion devices will connect to the IoT by the end of 2015, and that the IoT will have 25 billion connections by 2020 [1].

Given this exponential growth, a breach in security of an IoT device can have consequences that stretch far beyond technologies and geographic borders. It is critically important for all network operators, businesses, governments, and non-commercial enterprises with liability or another stake in IoT to ensure that networks, devices, and applications are secured effectively.

The IoT security threat is compounded by the fact that many deployments begin on a very small scale. Security is not always a high priority in this initial phase. As small-scale start-up deployments grow, security issues arise relative to manageability, user experience, and cost. If not tackled effectively, these issues can easily escalate and seriously harm users and device owners, putting their profitability, privacy, or physical wellbeing at risk. A simple example is a road sign, which must work effectively so that all of its users – commuters, in this case – can safely navigate the surrounding area.

The separation between owner and users is another factor that can obstruct IoT device security. As in the road sign example, users are very often not the owners of the device. They may not be aware that the device is experiencing an issue. On the other hand, there are devices, such as connected cars, where the user is the owner. A car's owner is most likely focused on ensuring that the car is equipped to travel safely from point A to point B. The owner is less likely to be concerned about managing the connected software embedded within the vehicle. This detachment in responsibility – either as user or owner – can create troubles when issues go undetected or where there is no point of contact for fixing an issue.

Location is also factor in securing IoT devices. This is particularly true for mobile devices, but even fixed devices can be located in environments that create security challenges. For example, devices stationed outdoors present security risks, ranging from theft and other challenges from malicious users to weather-induced corrosion, which can impact the physical device.

IoT devices now have long life spans. In part, this is due to the fact that they are integrated within equipment with long life spans. Cars and shipping containers are typically expected to last 10 years. Utility metering equipment

is not replaced for 20 or more years. These extended timelines increase device vulnerability. Plus, 10 to 20 years is a long time in technology terms. Devices will have to be upgraded with hardware, firmware, or software to keep up with technology advances that arrive during their life span. Device security can be compromised during upgrades.

Another key difference of IoT devices is how they are built. IoT devices, particularly as they emerge, are often built with different components or modules. Although many of these modules are off-the-shelf components, the combination of different elements may not be subject to the same security analysis as a manufactured device such as a mobile phone. This is common in the utility space, where the connectivity modem is "bolted" on to an existing meter. This assembly process can expose the device to threats, which can intensify as the device becomes more complex and incorporates more connection points.

Growth of IoT networks

Cellular operators are creating dedicated IoT networks to address the challenges that machine-to-machine (M2M) technologies will bring to mobile networks over the next few years. The number of devices connecting to the network is growing exponentially. But M2M will account for only 19 percent of mobile network connections and 4% of traffic, according to a May 2015 report by Machina Research [2]. The challenge comes from the fact that device behavior creates unpredictable traffic patterns.

The characteristics of IoT networks differ significantly from those of consumer cellular networks. IoT networks may be designed to align with M2M behaviors, which can combine long periods of little bandwidth use with bursts of activity. As a result, an IoT network will host a greater number of devices because the bandwidth required to support these devices will be much lower [2]. Network operators have devised tariff plans to match the lower resource use on these networks. This means the average revenue per machine will be lower. In addition, the increase in traffic upside for connected devices may be smaller.

Security threats in the Internet of Things

The differences in the IoT are giving rise to new security threats. The extended longevity of devices will present new threats over time. These threats will be further compounded by ownership and network operator changes that may take place during this time. Moreover, as the age of a device increases, so does the likelihood that its components will be damaged or break altogether. All of these factors can make devices less secure and more vulnerable to security challenges.

Security threats can take many forms. For instance, a breach could occur if someone tampers with the device by reflashing firmware, changing hardware, or manipulating it with probes or other peripherals. Hacking of devices that collect and transmit sensitive data can result in the loss of this data, either on a device-basis or by probing the data transmission.

Some security threats stem from the limitations of emerging network technologies. One such example is illustrated by the Sigfox network [3], which is a bi-directional low power wide area network. Sigfox is not fully encrypted and has a very low downstream data rate of 32 bytes per day. If traffic on this network was encrypted to ensure privacy, the client and server would have to agree on what type of encoding to use. It would not be possible to use the current Transport Layer Security (TLS) or Datagram Transport Layer Security (DTLS) techniques because there would not be enough bandwidth to support negotiation of the cipher suites. Careful compromises would be required to ensure security.

Other threats can arise from unauthorized use or misuse of devices, which can lead to denial of service. In some instances, denial of service can occur when a device's power source is inhibited. Consider, for example, a constrained gas or water meter with a battery designed to last 20 years when data is transmitted only once a day (a total of 7,300 downloads). If the device is compromised, data may be downloaded 7,300 times in a significantly shorter time period, discharging the battery completely.

Unauthorized use or misuse of devices can also occur through networks. This type of improper use can result in significant costs for network operators and breach boundaries set in tariff plans. Problems can also occur in instances where device use deviates from the operator's subscription terms – for example, if a SIM card is moved to another device.

Unentitled configurations may conflict with the application provider's business model and lead to misuse of components. Similarly, there are concerns about unauthorized use of a device in a particular location. Many companies operate regionally and may be restricted to using specific devices in certain locations to remain in line with regulations.

Effective decommissioning is another threat to consider. It is not only a concern as it relates to user data, but also in terms of protecting the integrity of the business model and other environmental considerations, such as ending the use of power and locating the device for disposal. As many of these devices are located remotely, this can prove particularly challenging.

NOKIA

Lightweight M2M device management

Lightweight M2M (LWM2M) is a secure device management protocol applied across device types within IoT as well as devices such as mobile phones, set top boxes, and enterprise tablets.

The protocol, which is prescribed by the Open Mobile Alliance [4][5], uses the off-the-shelf security model provided by DLTS [6]. It offers several different ways to secure devices, including X.509 certificates, raw private and public keys, and shared secrets.

The LWM2M protocol is lightweight in that its overhead and handshaking are optimized to target as many device types as possible, including very low bandwidth devices that run on Low Power Networks. This optimization has been achieved by using User Datagram Protocol (UDP) [7] instead of Transmission Control Protocol (TCP) [8] and compressing the protocol into tokens.

LWM2M has been tested in a number of use cases, including:

- Registration: Register a device and view it as part of an inventory.
- Provisioning: Set up the device at registration with information relating to its deployment.
- Updates: Change the firmware or software on the device to fix bugs, enhance functionality, or thwart security threats.
- Re-configuration: Change the configuration of the device to enhance functionality or thwart security threats.
- Monitoring: Track the status of different aspects of the device and receive alerts in important situations.
- Diagnosis: Offer tools to analyze a device and rectify issues. Diagnosis can range from simple actions such as rebooting a device to complex actions such as analyzing log files.
- Collection: Gather, filter, and aggregate information from a device. The information can be passed upstream for visualization and analytics.
- Decommissioning: Prepare the device for decommissioning, which can involve wiping the device's data and restoring its factory settings.

The LWM2M protocol is built on the Constrained Application Protocol (CoAP) using a RESTful message protocol. Its simple data model is made up of objects composed of name–value pairs. It is conceptually very simple, and it offers low interoperability and integration costs. Freeware clients are available for download as part of the Leshan project [9]. The LWM2M client is being incorporated into the ARM® mbed™ [10] operating system, which will run on a variety of ARM microcontrollers. Other commercial implementations are also available.

NOKIA

LWM2M and security

Figure 1 outlines the security architecture of IoT devices in which an independent LWM2M client sits alongside the core device application. The LWM2M client does not necessarily need to be integrated with the application. Rather, it can monitor application and device behavior and act as directed by the stakeholders. This approach can help reduce coordination costs.

Figure 1. IoT device security architecture with independent LWM2M client

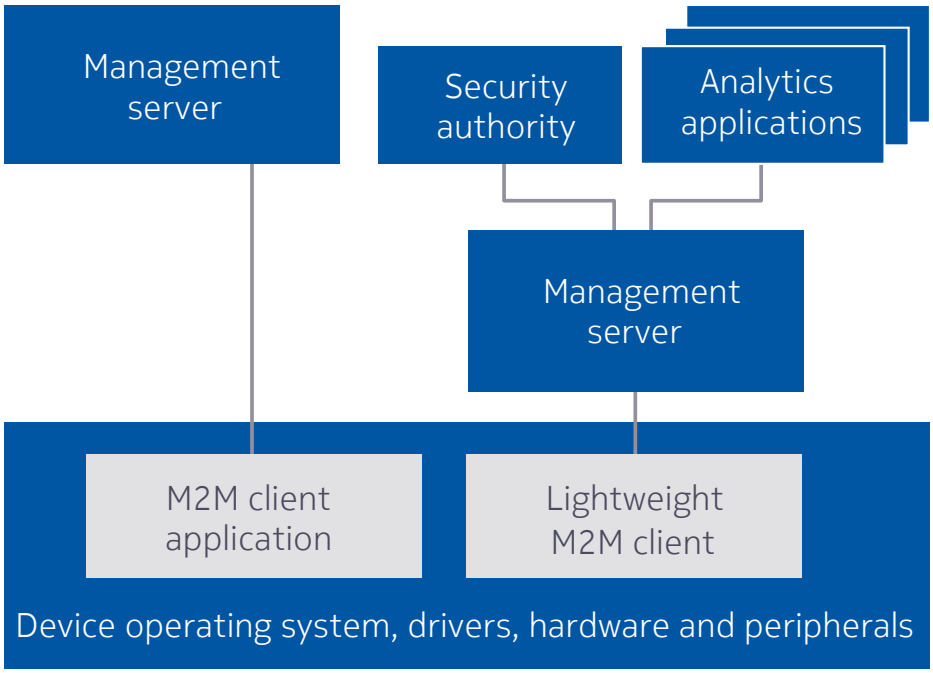

The architecture aims to address security threats in a cost-effective manner. Security policies can be applied to the client to allow for monitoring of security concerns such as tampering, feeding location or usage to detect fraud. In addition to monitoring the policy, the protocol invokes rules based on monitoring criteria. This can be as simple as notifying the server or, in situations where the security threat requires immediate action, acting locally within the device.

Security policies can be put in place to monitor and control firmware and software that is being applied to the device, either through the management server or inserted locally. This control helps ensure that new firmware of software does not compromise the device. A similar approach can be applied to ensure that configuration changes do not create security problems.

NOKIA

The LWM2M protocol can also be used to monitor a device's power and network use. Monitoring can be performed by way of counters and thresholds. Any discrepancy or change in behavior will prompt the LWM2M client to act. The client will send a report of the values it is monitoring. In more advanced situations, the LWM2M client will quarantine the application and suspend traffic, with the exception of that needed by the client itself.

A device's policy can also be linked to a particular subscription, usually identified by the IMSI. This linkage ensures that the operator's subscriptions are the only ones used with the device. If required, it can prevent the device from using particular operators.

In a similar vein, the policy can restrict the device so that it works only in a specific location or group of locations. The location can be defined at a very granular level (for example, a specific retail outlet) or more broadly aligned with a state or country (for example, California but not Nevada, England but not Scotland). If it encounters a breach in these rules, the LWM2M client can quarantine the application, disable cellular access, or generate an alarm on the server.

When a device comes to the end of its service life it will need to be effectively decommissioned. This can be achieved in a number of ways, including wiping the data from the device, restoring the device to its factory state, removing the application, or, in some cases, triggering self-destruction. These actions ensure that data remains protected and that the business model associated with the device maintains its integrity.

Benefits of LWM2M

The LWM2M client offers several benefits. For example:

- It provides a level of policing that application providers may not have the capability or capacity to provide.
- It offers a different view that allows security concerns to be addressed in way that complements the application provider.
- It captures and analyzes large amounts of data that can be used to enhance security and predict insecure scenarios.

The use of standardized protocols is the key to making deployment easy and cost effective. In addition, it offers a more consistent way to monitor and manage security threats. The very nature of a standardized solution means that it can be applied across numerous vendors, devices, and industries. This consistency offers a model of security that can be analyzed, understood, and improved.

NOKIA

Conclusion

Connected devices within the IoT are different from mobile phones or personal computers. There are billions of IoT devices connecting across geographic borders. They are often unmanned and embedded in equipment with a long lifespan. They will likely be upgraded with software and hardware multiple times to keep up with technological advances and changes in requirements. And they will have been assembled using different components and will likely have different owners over time. Recognizing these differences, cellular operators are creating separate networks and tariff plans that offer characteristics specific to the needs of connected devices.

The IoT has created security threats that, when combined with the expected scale of the IoT, could lead to cost damaging unmanageability if they are not properly addressed with each device rollout. The use of a distinctly separate LWM2M client enables the software and hardware on the device to be secured in a cost effective way. It also offers a standard model that can be used to monitor a large number of devices, and in case of a breach, provide effective remedies that align with the policies of the relevant security authority.

Abbreviations

DTLS	Datagram Transport Layer Security
IoT	Internet of Things
LWM2M	Lightweight Machine-to-Machine
M2M	machine-to-machine
REST	Representational State Transfer
SIM	subscriber identity module
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol

NOKIA

Sources

1. Gartner Says 4.9 Billion connected “Things” Will Be in Use in 2015 (2014, 11 11). Retrieved 09 07, 2015, from Gartner: http://www.gartner.com/newsroom/id/2905717
2. Machina Research. (2015, 09 07). M2M Forecasts. Retrieved 09 07, 2015, from Machina Research: https://machinaresearch.com/forecasts
3. Sigfox. (2014, 06 01). Sigfox Technoloty. Retrieved 09 07, 2015, from Sigfox.com: http://www.sigfox.com/en/#!/technology
4. Open Mobile Alliance. (2014, 06 20). M2M Enablers.Retrieved 09 07, 2015, from M2M Enablers: http://openmobilealliance.org/about-oma/work-program/m2m-enablers/
5. Open Mobile Alliance. (2014, 11 26). OMA LightweigtM2M v1.0. Retrieved 09 07, 2015, from OMA LightweightM2M v1.0: http://technical.openmobilealliance.org/Technical/technical-information/release-program/current-releases/oma-lightweightm2m-v1-0
6. Internet Engineering Task Force. (2006, 04 01). Datagram Transport Layer Security. Retrieved 09 07, 2015, from Datagram Transport Layer Security: https://tools.ietf.org/html/rfc4347
7. Internet Engineering Task Force. (1980, 08 01). User Datagram Protocal. Retrieved 09 07, 2015, from User Datagram Protocol: https://www.ietf.org/rfc/rfc768.txt
8. Internet Engineering Task Force. (1981, 09 01). Transmission Control Protocal. Retrieved 09 07, 2015, from Transmission Control Protocol: https://tools.ietf.org/html/rfc793
9. Leshan Project. (2015, 08 16). Leshan is an OMA Lightweight M2M (LWM2M) implementation in Java. Retrieved 09 07, 2015, from Github: https://github.com/eclipse/leshan
10. ARM. (2015, 04 01). Announcing our Plans for mbed v3.0. Retrieved 09 07, 2015, from ARM mbed Developer Site: https://developer.mbedorg/blog/entry/Announcing-our-plans-for-mbed-v30/

Nokia Oyj
Karaportti 3
FI-02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Product code: PR1510015182EN