WHITEPAPER

Securing the Internet of Things

Aeris COMMUNICATIONS

Table of Contents

INTRODUCTION

The topic of security is a complex issue. There are many factors that influence security in general, not just in the rapidly growing market for Machine to Machine ("M2M") communications and the Internet of Things ("IoT").

This white paper addresses the nature of security in the context of M2M and IoT. It describes some factors to take into account when designing an M2M or IoT application that transmits data from a remote device and send control messages to it.

EXPECTATIONS OF SECURITY RISK

However, we must set expectations correctly first. It is important to note that security risks can be:

- ... recognized and understood,
- ... detected and resolved,
- ... managed and control,
- ... but never completely eliminated!

When deploying an IoT application and designing for security, it is essential to understand at what points in the IoT communications and application operations path that there is an opportunity for risk management.

A design of a security plan for an M2M or IoT device and application can recognize and understand what possible risks the application may encounter once deployed. These design methods require understanding of the purpose of the application and the information flow chain that can be a target for security breaches of the overall IoT and M2M application.

Security breaches may be detected by understanding what breaches may occur in a given application, if that is possible, and design mechanisms to detect the breach when it occurs. When detected, a fix for the problem – whether it is dealing with compromised data or a compromised device – must be deployed rapidly.

The risk of security breaches can be managed and controlled – processes and tools to isolate problems, perhaps to a few devices, should be designed into the device firmware.

Ultimately, however, security risks cannot be completely eliminated.

It is simply not possible to understand every mechanism by which a particular M2M or IoT application can be compromised. Even if we can determine all possible threat mechanisms, the cost of designing a preventative measure to counter each threat might be prohibitive for the application.

COMPLEXITY

As mentioned earlier, security is a complex issue that includes many sub-topics.

For example, security is related to content privacy. Data sent from a remote device – for example, for medical M2M or IoT applications – must not be disclosed to unintended recipients. Laws and regulations may require specific actions to protect the privacy of patients under medical care.

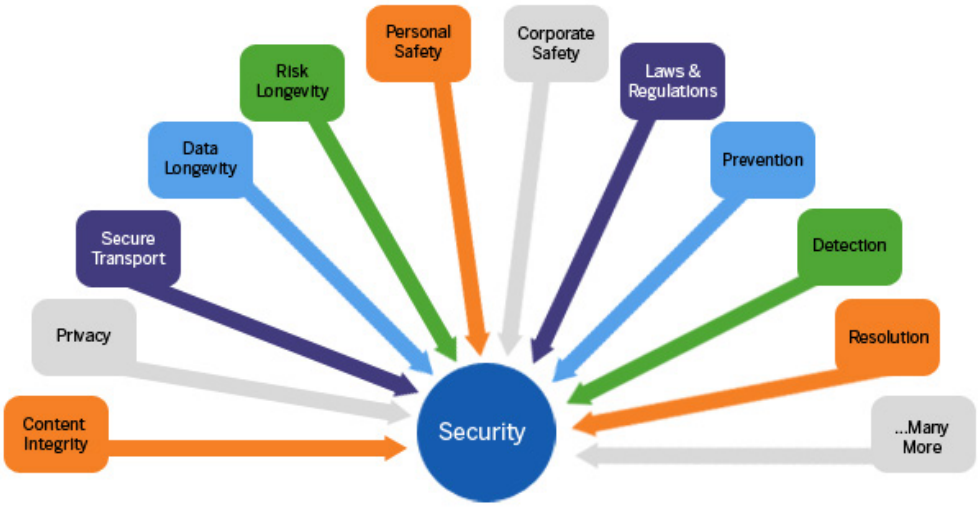

FIGURE 1. SECURITY COMPLEXITY

The implementation of breach detection mechanisms requires an understanding of what an M2M or IoT application may experience once it is deployed. If a breach goes undetected, and the data from a device is compromised, this may have consequences for the M2M or IoT application.

If personal safety (for individuals and groups) or corporate safety (for proprietary and confidential information) is vital, the IoT application firmware and data host systems must implement security solutions that are appropriate to the task.

For M2M and IoT applications, the concept of data longevity – i.e., long-term storage in archives – may require extended care to prevent information from falling into the wrong hands after its original purpose is completed. Companies, who record information and keep it properly private while in business, may decide to release that confidential data if the company fails, since there is financial value. For example, there are cases where bankruptcy executors have chosen to sell lists, and confidential information on users, from failed companies – effectively bypassing the confidentiality agreements that may have existed between the failed company and its users.

WHAT IS SECURE ENOUGH?

There are four key security objectives to achieve in any communications application:

- Authenticated sender & receiver
- Sender and receiver are accessible
- Trust in the data content
- Confidentiality of the information

Authenticated Sender & Receiver

In any communication where security is important, the sender and receiver of that transmission must be authenticated. The implementation must ensure that the correct devices are transmitting and that the correct servers are receiving the transmission.

Generally, the devices authenticate to the transport networks, the network infrastructure and remote servers, and the servers authenticate to the transport networks and remote devices.

Accessible Sender and Receiver

In any communications path, access to the network and infrastructure to transmit and receive data when needed, is important. If a device needs to reach a server to transmit data, that server must be available. If a server needs to send a control message to a device, that device must be on-line – or made on-line – to receive the message, sufficiently quickly to complete the action of the control transmission.

Trusted Content

This is the most important objective. The sender of the information must ensure that the data is sent correctly – the receiver must ensure that the data is received correctly and the information was not compromised while in the communications path.

Without trust, it is difficult to make necessary application decisions from the content – whether it is data received by a server from a remote device, or a control message received by a device from the server.

If trust is broken due to a security compromise (real or perceived), it is difficult for a server (or process or human) to make a decision on possibly erroneous data, or for a device to act on faulty control instructions.

Confidential Information

If confidentiality is important – whether related to privacy or to avoid providing information that may enable security breaches – only the intended recipient of the data must access it. The data from the device to the server (or control messages from the server to the device) must be secure to ensure that other elements in the communications path, that may have access to the transmitted bits, are unable to read the content or obtain value from the content.

Often, data encryption is required, although this may increase the amount of memory or processing power required in the devices.

SECURITY VS. OBJECTIVES AND RISKS

There is a balance needed between the cost and effort to implement security objectives and the amount of security that is required and achieved.

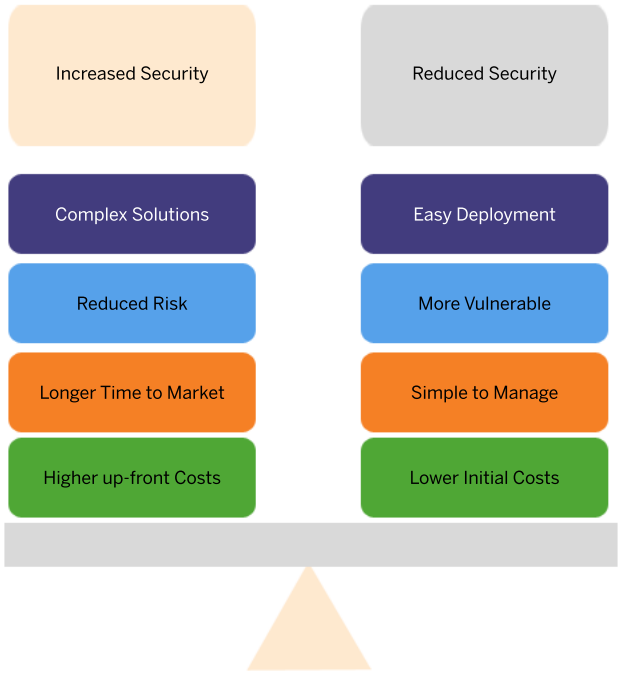

FIGURE 2. SECURITY TRADE-OFFS

For example, if security and data compromise is not a concern for a given application, it is easier to deploy the devices for that application, faster to get to market, simpler to manage and has lower initial costs. It is, of course, more vulnerable to security breaches and this may ultimately be a higher cost, if a breach has serious consequences, or causes failure of a device.

However, if effort is spent on security design during the development because it is believed to be an issue for an application, this generally means higher initial cost, a more complex design and takes longer to deploy into the market. It is, of course, less likely to suffer from a security breach and potentially prevents the application or device from failing.

These issues, and security implementation decisions, should be considered during the design of an IoT application – as they can have financial ramifications during the life of the deployed devices.

IOT RISK MANAGEMENT

When deploying an IoT application and designing for security, it is essential to understand at what points in the IoT communications and application operations path that there is an opportunity for risk management.

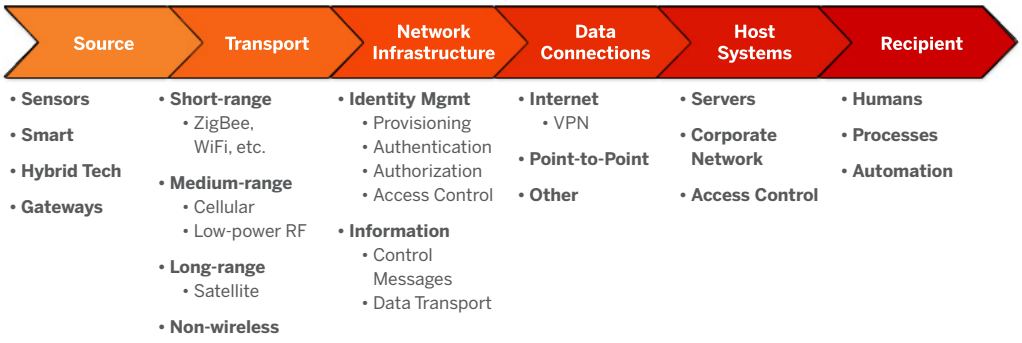

FIGURE 3. IOT RISK MANAGEMENT ACCESS POINTS

In a typical application, the Source of IoT Risk Management Access Points (see Figure 3) is usually a sensor or device, perhaps connected through a transport gateway. This source provides the data measured for the application, for eventual delivery to the Recipient – shown at the right of the diagram. Data may also be traversing from the Recipient to the Source – control messages sent to devices in the field allow them to take specific actions.

Each of these elements in the chain is a source of risk; security may be breached at that point in the chain. It is important to understand the specific risks associated with the elements.

- The Source of the data – the sensors and devices,
- The Transport of the bits and bytes from and to the Source,
- The Network that provides and supports the transport technology,
- The Data Connections from the network to send the bits and bytes on to
- The Customer Host systems and IT servers, and ultimately, to
- The Recipient of the data ... whether it is a human or an automated process that interprets data to make decisions and take actions.

It is important to understand all the access points in the chain, and evaluate if a breach in any given element of the communications chain could lead to issues at other elements in the chain.

For example, if a sensor at the left end of the chain transmits bad data, it could lead to faulty business decisions all the way at the other end of the chain by the Recipient – indeed, potentially without affecting any element in between the two end-points.

If a customer host system is compromised by an intruder from the Internet, the confidentiality of the data becomes suspect. This could affect decisions taken by the Recipient, but might not have any impact on the data transport chain or the sensor.

Each M2M or IoT application follows the basic flow in the chain shown, but the degree and type of concern of any given risk issue varies. Some applications may be tolerant of risk of compromise at the sensor – perhaps it is a relatively benign data collection function – but others may be specifically harmful.

For example, medical applications that monitor and report fitness measurements may not be a risk issue for the person wearing the sensors – erroneous data is simply suspect and ignorable. However, an automated insulin pump implanted in a diabetic must be highly secure to avoid serious harm – possibly fatal – if its security is compromised.

WHO IS AFFECTED? AND HOW MUCH?

Not surprisingly, the cost of reducing risk and increasing security may be quite significant. How much effort – financial and technical – to expend in search of implementing a security solution depends on many factors.

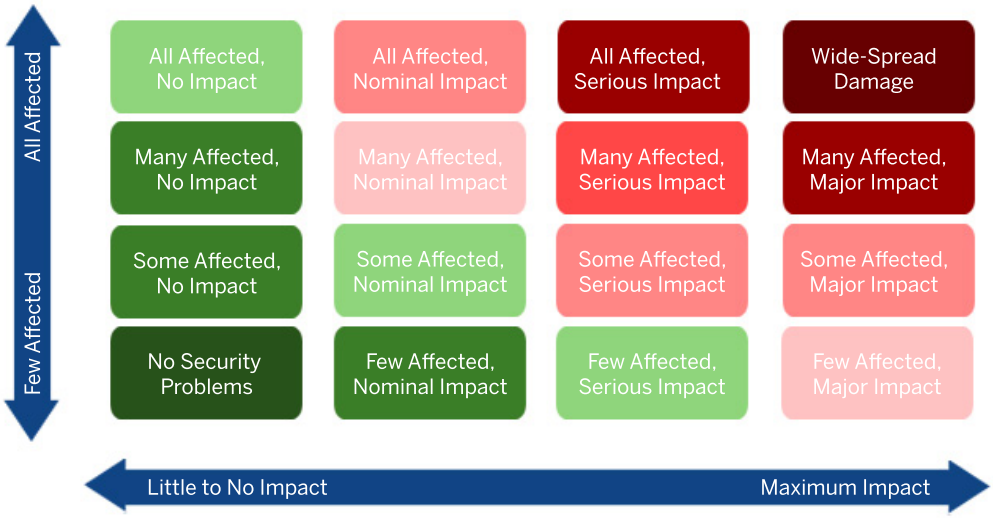

FIGURE 4. ASSESSING IMPACT AND AFFECTED POPULATION

Figure 4 shows one possible comparison of two factors in an IoT application: how many people are affected by a possible breach, and what is the impact to those who are affected. In this figure, the bottom left are applications where a security breach affects very few people and has little to no impact in a compromised data scenario. For example, a bad temperature sensor that reports incorrect data in a Home Automation application is not likely to affect more than a single home-owner and has no lasting impact once it is repaired.

In this example above, it may not be cost-effective to design a solution to avoid bad data from individual sensors – perhaps it may be quite sufficient to report the bad data to solve a product quality problem with a sensor or device.

The upper right on the chart however, may be an application where a security event could affect many people in a region and have a major impact. For example, an alarm at a water distribution center that fails to report an intruder with malicious intent, could seriously damage the well-being of a large population center. This example may require stringent and significant effort – perhaps with redundant monitoring devices and multiple transport technologies, including regular system checks, etc.

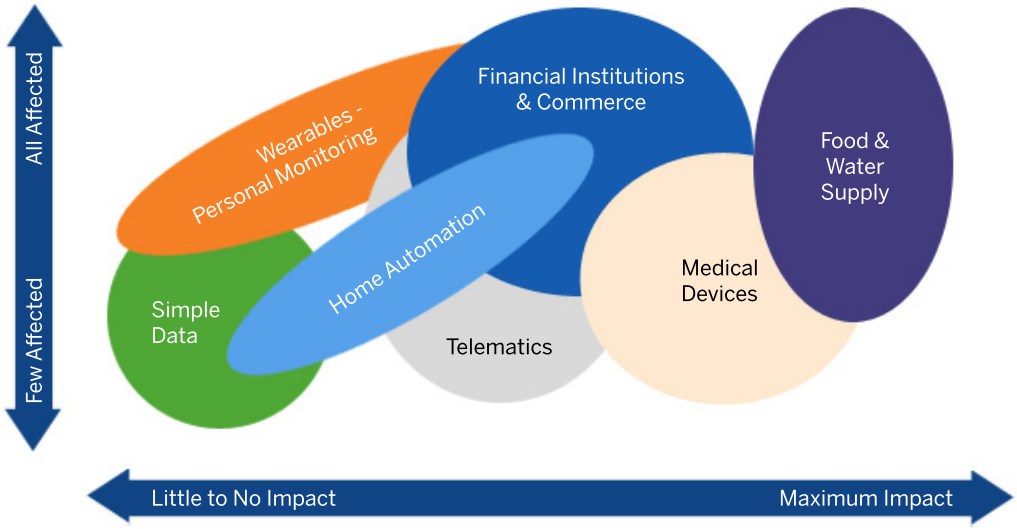

FIGURE 5. IOT APPLICATION IMPACT OF COMPROMISED DATA

Figure 5 - IoT Application Impact shows where some typical IoT categories of applications lie in this risk assessment. Usually, the assessment is not a single, or a simple, answer for each IoT category. Some feature in an application may put it more at risk (for example, for legal liabilities) than other applications, and a better security implementation may be needed to avoid significant damage.

Simple data gathering functions tend not to have a major impact of many people. For example, a home temperature monitor that transmits information to an individual homeowner may not have much impact if it is compromised – the owner can ignore erroneous readings. On the other hand, food and water supply chain controls and management systems are vital to have excellent security controls, since many people in a region could be affected by compromises and breaches.

Other M2M and IoT applications fall into a range between these axis extremes, with varied impact and varied number of people or systems that are affected. For example, Telematics device tend to have a medium level of impact if compromised. Usually, only a few people are affected by a compromised device, but the impact could be relatively low to medium, depending on the specific issue with the application.

Since an application within any of these categories can have a different set of included applications, it is important to check and assess the functions within each application and determine the level of security protection required.

UNIQUE M2M AND IOT CHALLENGES

Deploying an IoT application has unique security challenges. This section describes some of those unique issues that should be addressed in any deployment.

Scaling

The industry has predicted that billions of M2M and IoT devices will be deployed in the next five to ten years. Design for proper security and risk management is essential – the sheer number of compromised devices and data could overwhelm the system otherwise.

If an issue with a single application and its device is not of great significance or impact, it could add up to a large problem if that application is successful and a large number of its devices are deployed in the field.

Longevity

Unlike cellular handsets and computers and many other large-scale product deployments in the communications industry, M2M and IoT applications are generally deployed for long periods – devices may stay operational for years.

For example, in Aeris' experience with AMPS analog cellular, deployed units stayed operational for more than a decade, and were still fully functional – performing their M2M application functions – until they were removed from service due to the Analog system shutdown in 2008.

In many M2M and IoT applications, devices may also be relatively inaccessible – deployed in remote locations and inside other systems. Unlike cellular handsets, for example, that could be replaced by requesting it of their human users. Thus, a compromised device could remain in service, transmitting its compromised data, for a long time. Remote and inaccessible devices, or even ones that require significant cost and effort to replace (such as inside the console of a vehicle) could impact the overall application and system due to a security breach.

There is a balance needed between the cost and effort to implement security objectives and the amount of security that is required and achieved.

Automation

Remote M2M and IoT devices often send information that trigger alerts and automated actions. If a device, or a group of devices, is compromised, a simplistic automated response could cascade into a more impactful problem. For example, if a device cannot access the network or the remote data server, it should build in intelligent retry algorithms. A simplistic design of automated, repetitive retry attempts could overwhelm a network or server designed for handling the usual peak traffic, but which is unable to handle unusual traffic.

Regulations

At the moment, the legal regulations governing security for M2M and IoT applications, and data communications, are not yet fully in place, or are still playing catch-up. For example, in the US, the HIPAA rules for patient information privacy were designed in an era when the common transport for medical information between medical providers was postal mail or fax transmissions. After receiving permission from the patient, doctors and medical facilities could communicate information using these transports.

Sending the same information via e-mail or allowing access via a web site however, was not legal until recently – the HIPAA rules had to be amended to allow e-mail as a transport, as long as the data was encrypted.

Other M2M and IoT applications, such as data transmissions from medical units in the mHealth industry, will also require legal regulations that are not yet in place. Although there is a growing awareness that these regulations are needed, the work is in progress.

Standards

Standards for security in M2M and IoT implementations are in development. These standards would provide guidance for developers and ensure a common understanding of the needs and requirements for secure communication in this field. In particular, the architecture specification effort of the OneM2M project (see www.onem2m.org) for incorporating security in the design of the M2M and IoT applications – right from the start – is highly commendable. However, this architecture specification has not yet been officially released.

CONCLUSION

In this relatively new market of M2M and IoT, it is not obvious to everyone that security issues exist with these applications and devices. These security challenges must be addressed early and while the number of deployed devices are still relatively low.

This white paper addressed the complex issue of security in a rapidly growing market, and the issues and factors that M2M and IoT application developers should address during the design of the devices, the transport and the overall application.

For more information, call 1-888-GO-AERIS (1-888-462-3747) in North America or +44 118 925 3202 in Europe. You can also write to info@aeris.net.

ABOUT AERIS COMMUNICATIONS

Aeris is a pioneer and leader in the market of the Internet of Things – as an operator of end-to-end M2M services and as a technology provider enabling other operators to deliver profitable M2M services. Among our customers are the most demanding users of M2M services today, including Hyundai, Acura, Rand McNally, Leica, and Sprint. Through our "Made for Machines" technology and services, we strive to fundamentally improve their businesses – by dramatically reducing costs, improving operational efficiency, reducing time-to-market, and enabling new revenue streams.