



IoT Security Guidelines Overview Document

February 2016



IoT Security Guidelines Overview Document

Version 1.0

08 February 2016

This is a Non-binding Permanent Reference Document of the GSMA

Security Classification: Non-confidential

Access to and distribution of this document is restricted to the persons permitted by the security classification. This document is confidential to the Association and is subject to copyright protection. This document is to be used only for the purposes for which it has been supplied and information contained in it must not be disclosed or in any other way made available, in whole or in part, to persons other than those permitted under the security classification without the prior written approval of the Association.

Copyright Notice

Copyright © 2016 GSM Association

Disclaimer

The GSM Association ("Association") makes no representation, warranty or undertaking (express or implied) with respect to and does not accept any responsibility for, and hereby disclaims liability for the accuracy or completeness or timeliness of the information contained in this document. The information contained in this document may be subject to change without prior notice.

Antitrust Notice

The information contain herein is in full compliance with the GSM Association's antitrust compliance policy.

Table of Contents

1	Introduction	4
1.1	Executive Overview	4
1.2	GSMA IoT Security Guideline Document Set	4
1.3	Document Purpose	5
1.4	Intended Audience	5
1.5	Definitions	6
1.6	Abbreviations	7
1.7	References	8
2	The Challenges Created by the Internet of Things	9
2.1	The Availability Challenge	9
2.2	The Identity Challenge	10
2.3	The Privacy Challenge	10
2.4	The Security Challenge	11
3	The Mobile Solution	12
3.1	Addressing the Challenge of Availability	12
3.2	Addressing the Challenge of Identity	12
3.3	Addressing the Challenge of Privacy and Security	13
4	The IoT Model	14
4.1	Service Ecosystem	14
4.2	Endpoint Ecosystem	14
5	Risk Assessments	15
5.1	Goal	16
5.2	Risk Model References	16
6	Privacy Considerations	16
7	Using This Guide Effectively	18
7.1	Evaluating the Technical Model	18
7.2	Review the Current Security Model	19
7.3	Review and Evaluate Recommendations	19
7.4	Implementation and Review	20
7.5	Ongoing Lifecycle	21
8	Example – Wearable Heart Rate Monitor	21
8.1	The Endpoint Overview	21
8.2	The Service Overview	22
8.3	The Use Case	22
8.4	The Security Model	23
8.5	The Result	24
8.6	Summary	25
9	Example – Personal Drone	25
9.1	The Endpoint Overview	25
9.2	The Service Overview	26
9.3	The Use Case	26
9.4	The Security Model	27

9.5	The Result	28
9.6	Summary	28
10	Example – Vehicle Sensor Network	29
10.1	The Endpoint Overview	29
10.2	The Service Overview	30
10.3	The Use Case	30
10.4	The Security Model	31
10.5	The Result	32
10.6	Summary	32
Annex A	Recommended Privacy Considerations for IoT Service providers	33
Annex B	Example based upon Automotive Tracking System	37
B.1	Evaluating the Technical Model	37
B.2	Review the Security Model	37
B.3	Review and Assign Security Tasks	38
B.4	Review Recommendations	39
B.5	Review Component Risk	39
B.6	Implementation and Review	39
B.7	Ongoing Lifecycle	40
Annex C	Document Management	41
C.1	Document History	41
C.2	Other Information	41

1 Introduction

1.1 Executive Overview

The emergence of the Internet of Things (IoT) is creating new service providers who are looking to develop new, innovative, connected products and services. Analysts have predicted that hundreds of thousands of new IoT services will connect billions of new IoT devices over the next decade. This rapid growth of the Internet of Things represents a major opportunity for all members of the new ecosystem to expand their service offerings and to increase their customer base.

Analysts have indicated that security issues are a significant inhibitor to the deployment of many new IoT services and, at the same time, the provision of wide area connectivity to an ever-widening variety of IoT services will increase the whole ecosystem's exposure to fraud and attack. There is already much evidence to show that attackers are beginning to show ever greater interest in this area.

As these new service providers develop new and innovative services for particular market segments, they may be unaware of the threats their service may face. In some cases, the service provider may not have developed a service that has connected to a communications network or the internet before and they may not have access to the skills and expertise to mitigate the risks posed by enabling internet connectivity within their devices. In contrast, their adversaries understand the technology and security weaknesses, quickly taking advantage if vulnerabilities are exposed.

Whilst many service providers, such as those in automotive, healthcare, consumer electronics and municipal services, may see their particular security requirements as being unique to their market, this is generally not the case. Almost all IoT services are built using endpoint device and service platform components that contain similar technologies to many other communications, computing and IT solutions. In addition to this, the threats these different services face, and the potential solutions to mitigate these threats, are usually very similar, even if the attacker's motivation and the impact of successful security breaches may vary.

The telecommunications industry, which the GSMA represents, has a long history of providing secure products and services to their customers. To help ensure that the new IoT services coming to market are secure, the network operators together with their network, service and device equipment partners would like to share their security expertise with service providers who are looking to develop IoT services.

The GSMA has therefore created this set of security guidelines for the benefit of service providers who are looking to develop new IoT services.

1.2 GSMA IoT Security Guideline Document Set

This document is the first part of a set of GSMA security guideline documents that are intended to help the nascent "Internet of Things" industry establish a common understanding of IoT security issues. The set of guideline documents promotes a methodology for developing secure IoT Services to ensure security best practices are implemented throughout the life cycle of the service. The documents provide recommendations on how to mitigate common security threats and weaknesses within IoT Services.

The structure of the GSMA security guideline document set is shown below. It is recommended that this document, (i.e. the overview document) is read as a primer before reading the supporting documents.

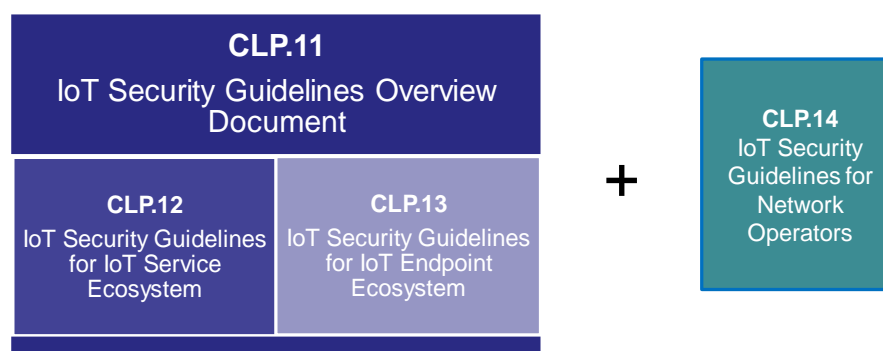


Figure 1 - **GSMA IoT Security Guidelines Document Structure**

Network Operators, IoT Service Providers and other partners in the IoT ecosystem are advised to read GSMA document CLP.14 “IoT Security Guidelines for Network Operators” [13] which provides top-level security guidelines for Network Operators who intend to provide services to IoT Service Providers to ensure system security and data privacy.

1.3 Document Purpose

The goal of the Internet of Things Security Guidelines document set is to provide the implementer of an IoT technology or service with a set of design guidelines for building a secure product. To accomplish this task, this document will serve as an overarching model for interpreting what aspects of a technology or service are relevant to the implementer. Once these aspects, or components, are identified, the implementer can evaluate the risks associated with each component, and determine how to compensate for them. Each component can be broken down into *sub-components*, where more granular risks will be described. Each risk shall be assigned a priority, to assist the implementer in determining the cost of the attack, as well as the cost of remediation, and the cost, if any, of not addressing the risk.

The scope of this document is limited to recommendations pertaining to the design and implementation of IoT services.

This document is not intended to drive the creation of new IoT specifications or standards, but will refer to currently available solutions, standards and best practice.

This document is not intended to accelerate the obsolescence of existing IoT Services.

It is noted that adherence to national laws and regulations for a particular territory may, where necessary, overrule the guidelines stated in this document.

1.4 Intended Audience

The primary audience for this document are:

- IoT Service Providers - enterprises or organisations who are looking to develop new and innovative connected products and services. Some of the many fields IoT

Service Providers operate in include smart homes, smart cities, automotive, transport, health, utilities and consumer electronics.

- IoT Device Manufacturers - providers of IoT Devices to IoT Service Providers to enable IoT Services.
- IoT Developers - build IoT Services on behalf of IoT Service Providers.
- Network Operators who are themselves IoT Service Providers or build IoT Services on behalf of IoT Service Providers.

1.5 Definitions

Term	Description
Access Point Name	Identifier of a network connection point to which an endpoint device attaches. They are associated with different service types, and in many cases are configured per network operator.
Attacker	A hacker, threat agent, threat actor, fraudster or other malicious threat to an IoT Service typically with the intent of retrieving, destroying, restricting or falsifying information. This threat could come from an individual criminal, organised crime, terrorism, hostile governments and their agencies, industrial espionage, hacking groups, political activists, 'hobbyist' hackers, researchers, as well as unintentional security and privacy breaches.
Cloud	A network of remote servers on the internet that host, store, manage, and process applications and their data.
Complex Endpoint	This Endpoint model has a persistent connection to a back-end server over a long-distance communications link such as cellular, satellite, or a hardwired connection such as Ethernet. See CLP.13 [4] for further information.
Components	Refers to the components contained in documents CLP.12 [3] and CLP.13 [4]
Embedded SIM	A SIM which is not intended to be removed or replaced in the device, and enables the secure changing of profiles as per GSMA SGP.01 [2].
Endpoint	A generic term for a lightweight endpoint, Complex Endpoint, gateway or other connected device. See CLP.13 [4] for further information.
Endpoint Ecosystem	Any configuration of low complexity devices, rich devices, and gateways that connect the physical world to the digital world in novel ways. See section 4.2 for further information.
Internet of Things	The Internet of Things (IoT) describes the coordination of multiple machines, devices and appliances connected to the Internet through multiple networks. These devices include everyday objects such as tablets and consumer electronics, and other machines such as vehicles, monitors and sensors equipped with communication capabilities that allow them to send and receive data.
IoT Service	Any computer program that leverages data from IoT devices to perform the service.
IoT Service Provider	Enterprises or organisations who are looking to develop new and innovative connected products and services.
Network Operator	The operator and owner of the communication network that connects the IoT Endpoint Device to the IoT Service Ecosystem.
Organizational Root of Trust	A set of cryptographic policies and procedures that govern how identities, applications, and communications can and should be cryptographically secured.

Term	Description
Recommendations	Refers to the recommendations contained in documents CLP.12 [3] and CLP.13 [4]
Risk	Refers to the risks contained in documents CLP.12 [3] and CLP.13 [4]
Security Tasks	Refers to the security tasks contained in documents CLP.12 [3] and CLP.13 [4]
Service Access Point	A point of entry into an IoT Service's back end infrastructure via a communications network.
IoT Service Ecosystem	The set of services, platforms, protocols, and other technologies required to provide capabilities and collect data from Endpoints deployed in the field. See section 3.1 for further information.
Subscriber Identity Module (SIM)	The smart card used by a mobile network to authenticate devices for connection to the mobile network and access to network services.
UICC	A Secure Element Platform specified in ETSI TS 102 221 that can support multiple standardized network or service authentication applications in cryptographically separated security domains. It may be embodied in embedded form factors specified in ETSI TS 102 671.

1.6 Abbreviations

Term	Description
3GPP	3 rd Generation Project Partnership
API	Application Program Interface
APN	Access Point Name
CERT	Computer Emergency Response Team
CLP	GSMA's Connected Living Programme
CPU	Central Processing Unit
EAP	Extensible Authentication Protocol
EEPROM	Electrically Erasable Programmable Read-Only Memory
GBA	Generic Bootstrapping Architecture
GPS	Global Positioning System
GSMA	GSM Association
GUI	Graphic User Interface
HIPAA	Health Insurance Portability and Accountability Act
IoT	Internet of Things
LPWA	Low Power Wide Area
NIST	National Institute of Standards and Technology
OBD	On Board Diagnostics
OCTAVE	Operationally Critical Threat, Asset, and Vulnerability Evaluation
OMA	Open Mobile Alliance
PIA	Privacy Impact Assessment
PII	Personally Identifiable Information
RAM	Random Access Memory

Term	Description
SIM	Subscriber Identity Module

1.7 References

Ref	Doc Number	Title
[1]	n/a	"The Mobile Economy 2015" http://www.gsma.com/mobileeconomy/
[2]	SGP.01	"Embedded SIM Remote Provisioning Architecture" http://www.gsma.com/connectedliving/embedded-sim/
[3]	CLP.12	IoT Security Guidelines for IoT Service Ecosystem www.gsma.com/connectedliving
[4]	CLP.13	IoT Security Guidelines for IoT Endpoint Ecosystem www.gsma.com/connectedliving
[5]	n/a	NIST Risk Management Framework http://csrc.nist.gov/groups/SMA/fisma/framework.html
[6]	CMU/SEI-2007-TR-012	Introducing OCTAVE Allegro: Improving the Information Security Risk Assessment Process http://www.cert.org/resilience/products-services/octave/
[7]	Not Used	Not Used
[8]	TS 33.220	Generic Authentication Architecture (GAA); Generic Bootstrapping Architecture (GBA) www.3gpp.org
[9]	RFC 4186	Extensible Authentication Protocol Method for Global System for Mobile Communications (GSM) Subscriber Identity Modules (EAP-SIM) www.ietf.org
[10]	n/a	Conducting privacy impact assessments code of practice https://ico.org.uk/media/for-organisations/documents/1595/pia-code-of-practice.pdf
[11]	n/a	Open Mobile Alliance http://openmobilealliance.org/
[12]	n/a	oneM2M http://www.onem2m.org/
[13]	CLP.14	IoT Security Guidelines for Network Operators www.gsma.com/connectedliving
[14]	GE.11-13201	Report of the Special Rapporteur on the promotion and protection of the right to freedom of opinion and expression, Frank La Rue* www.ohchr.org/english/bodies/hrcouncil/docs/17session/A.HRC.17.27_en.pdf
[15]	n/a	Right to Internet Access https://en.wikipedia.org/wiki/Right_to_Internet_access

2 The Challenges Created by the Internet of Things

Several years ago a United Nation's special report recommended that the Internet is a basic human right, and that all people of the world should have access to broadband services [14]. More recently laws are being adopted in countries such as France, Greece, Spain and others [15], to ensure that Internet access is broadly available and/or to prevent the state from unreasonably restricting an individual's access to information and the Internet.

These declarations are the result of the rapid social and technological changes that have stemmed from the growth of the Internet. This has resulted in the Internet becoming a way of life, one of the primary sources of all classes of information, and the most common method for maintaining connectivity to loved ones and peers. The Internet is not simply a technology, it has become a *part* of us.

In concert with the growing desire to maintain connectivity, a technological explosion has occurred over the past few years. While technologists have declared "The Internet of Things is coming!" for over a decade, the interest in ubiquitous access to information and the cost model required to do so had not yet combined into a practical business model until the past five years. At this point, component costs sharply decreased, while access to wireless services and the speed of those services have dramatically increased. Protocols, battery life, and even business models have all evolved to accommodate our ever increasing demand for information and connectivity.

And that, in essence, is what the Internet of Things is all about. It isn't really about things. It's about Us. The Internet of Us. The human and digital experiences no longer sit side-by-side, they are bound ever tighter by this new way of life.

And because the human physical experience is bound more to the digital world than ever before, it must be protected, as digital security now directly impacts the physical world more than ever. The Internet of Things is an excellent opportunity for the world to move forward together, in order to create ever greater databases of knowledge, shared experiences, and explosions of innovation. But, for this to work effectively, the technologies that drive this connectivity must be secured, to enforce the privacy, reliability, and quality of services necessary to ensure that this great utility, this imperative basic need, is kept available to all those that require it.

For the Internet of Things to evolve effectively, we must resolve the security challenges inherent to its growth. These challenges are:

- Availability: Ensuring constant connectivity between Endpoints and their respective services
- Identity: Authenticating Endpoints, services, and the customer or end-user operating the Endpoint
- Privacy: Reducing the potential for harm to individual end-users
- Security: Ensuring that system integrity can be verified, tracked, and monitored

2.1 The Availability Challenge

For the Internet of Things to evolve at its expected pace, Endpoint devices must be able to constantly communicate with each other, end-users, and back-end services. To accomplish this, new technologies need to be designed that allow persistent connectivity. This dovetails

well with the challenge of ubiquitous Internet access for the modern world. For this to succeed, several questions must be answered:

- How can Low Power Wide Area Networks (LPWAN) be implemented with a similar level of security to modern cellular systems?
- How can multiple mobile operators support the same level of security as new IoT Endpoints migrate across network boundaries?
- How can network trust be *forwarded* to capillary Endpoints that rely on Gateway Endpoints for communication?
- How can the power constraints of Lightweight Endpoints be addressed in secure communications environments?

2.2 The Identity Challenge

In order for an Endpoint to function within an IoT product or service ecosystem, it must be capable of securely identifying itself to its peers and services. This critical and fundamental aspect of IoT technology ensures that services and peers are able to guarantee to what – and to whom – data is being delivered. Access to information and services isn't the only issue directly tied to identity. We also must ask the questions:

- Can the user operating the Endpoint be strongly associated with the Endpoint's identity?
- How can services and peers verify the identity of the end-user by verifying the identity of the Endpoint?
- Will Endpoint security technology be capable of securely authenticating peers and services?
- Can rogue services and peers impersonate authorized services and peers?
- How is the identity of a device secured from tampering or manipulation?

2.3 The Privacy Challenge

Privacy can no longer be seen as an add-on to existing products and services. Because the physical world is directly affected by actions taken in the digital world, privacy must be designed into products from the ground up, to ensure that every action is authorized and every identity is verified while guaranteeing that these actions and the associated meta-data are not exposed to unauthorized parties. This can only be achieved by defining the proper architecture for a product or service, and is exceptionally difficult and expensive to perform retroactively.

Medical devices, automotive solutions, industrial control systems, home automation, building and security systems, and more, all directly impact human physical lives. It is the duty of the engineers to uphold these products and services to the highest level of assurance possible, to reduce the potential for physical harm as well as the exposure of privacy relevant data.

Therefore, we must ask ourselves how privacy affects not only the end-user, but how IoT technologies are designed:

- Is the identity of an Endpoint exposed to unauthorized users?
- Can unique Endpoint or IoT Service identifiers allow an end-user or Endpoint to be physically monitored or tracked?

- Is data emanating from an Endpoint or IoT Service indicative of or directly associated with physical end-user attributes such as location, action, or a state, such as *sleeping* or *awake*?
- Is confidentiality and integrity employed with sufficient security to ensure that patterns in the resultant cipher-text cannot be observed?
- How does the product or service store or handle user-specific Personally Identifiable Information (PII)?
- Can the end-user control the storage or use of PII in the IoT Service or product?

2.4 The Security Challenge

While Internet security has drastically improved over the past several decades, there have been several significant gaps in the overall health of modern technology. These gaps have been most evident in embedded systems and in cloud services - the two primary components in IoT technology.

In order for IoT to evolve while not exposing massive groups of users and physical systems to risk, information security practices must be enforced on both Endpoints and IoT Services.

- Are security best practices incorporated into the product or service at the start of the project?
- Is the security life-cycle incorporated into the Software or Product Development Life Cycle?
- Is application security being applied to both services and applications running on the embedded system?
- Is a Trusted Computing Base (TCB) implemented in both the Endpoint and the Service Ecosystem?
- How does the TCB enforce self-verification of application images and services?
- Can the Endpoint or IoT Service detect if there is an anomaly in its configuration or application?
- How are Endpoints monitored for anomalies indicative of malicious behaviour?
- How is authentication and identity tied to the product or service security process?
- What incident response plan is defined for detected anomalies indicative of a compromise?
- How are services and resources segmented to ensure a compromise can be contained quickly and effectively?
- How are services and resources restored after a compromise?

3 The Mobile Solution

While there have been a myriad of technologies that offer connectivity solutions for IoT, none shape the future of IoT better than mobile networks. Mobile networks offered the first wireless services to consumers and industry over twenty years ago, and have been building reliable, available, secure, and cost effective services ever since. The mobile industry has extensive experience in network availability due to the volatile nature of wireless radio networks managed over long distances. Network identity has been a challenge that has spawned numerous standards, device technologies, protocols and analytics models. Privacy and security are constant concerns of the mobile industry, who have worked to decrease the potential for abuses, identity theft, and fraud in all mobile technology.

3.1 Addressing the Challenge of Availability

According the GSMA's "The Mobile Economy 2015" report [1]:

- The mobile industry continues to scale rapidly, with a total of 3.6 billion unique mobile subscribers at the end of 2014. Half of the world's population now has a mobile subscription—up from just one in five 10 years ago and an additional one billion subscribers are predicted by 2020, taking the global penetration rate to approximately 60%. There were 7.1 billion global SIM connections at the end of 2014, and a further 243 million machine-to-machine connections.
- There is an accelerating technology shift to mobile broadband networks across the world. Mobile broadband connections (i.e. 3G and 4G technologies) accounted for just under 40% of total connections at the end of 2014, but by 2020 will increase to almost 70% of the total.
- While 2G remains the dominant network technology globally today, its position has already declined materially. 2G connections accounted for 90% of the total in 2008, but this had fallen to around 60% at the end of 2014. In absolute terms, the number of 2G connections peaked in 2013 and fell by 6% during 2014.
- The ongoing technology migration to higher speed networks is also facilitated by significant operator investments. Recent research from GSMA predicts that more than four out of five people will have access to 3G networks by 2020, up from 70% today. The report also highlights that 4G networks are being rolled out at a faster pace than was the case with 3G. While it took 10 years for 3G network coverage to reach half of the global population, it will take 4G networks just eight years after launch to reach the same milestone in 2017.

In the future we will also see the integration of Low-Power Wide-Area (LPWA) wireless technology into the cellular communications space covering the needs of IoT. This class of communications technology offers the wide area, wireless connectivity of today's mobile networks at a fraction of the power required to communicate effectively. Mobile operators will be integrating LPWA protocols and technologies into their offerings to provide services and solutions to businesses in the near future.

3.2 Addressing the Challenge of Identity

Identity management has been a challenge for decades and has strengthened the mobile industry's standards and technology offerings significantly. While the mobile industry is typically associated with the removable SIM card, the GSMA has created a SIM based

solution called the ‘Embedded SIM Remote Provisioning Architecture’ [2] which is appropriate for use in IoT to enable deeper component level integration into Endpoint devices, reduced production costs and the management of connectivity via Over-The-Air (OTA) platforms to enable the connectivity of the IoT Endpoint devices for their whole lifetime.

Identity technologies, such as the Embedded SIM, are designed as trust anchors that integrate security by default. They are manufactured to withstand attacks such as:

- Glitching
- Side-channel analysis
- Passive data interception
- Physical tampering
- Identity theft

An excellent advancement to this already security hardened technology is that new generations of these trust anchors incorporate an important addition to the IoT landscape. These technologies will be dual use. They won't simply be used to verify the security of the network, they will also be capable of securing application communications and the application itself, similar to traditional computing trust anchors.

This dual use capability will be further augmented by the integration of mobile industry security specifications such as those provided by 3GPP GBA [8], OMA [11], oneM2M [12] and others. These technologies will help to securely provision devices in the field, securely enable over-the-air firmware updates, and manage device capabilities and identity.

These technologies, when used together, will ease the currently complex engineering processes and combine it into one simple component. Instead of application engineers building complex technologies that they themselves have to manage, the network operator, who already manages the network identity, can perform this on behalf of the application. This not only reduces the engineering complexity, but the business's daily management requirements.

3.3 Addressing the Challenge of Privacy and Security

Along with the capabilities of the SIM, the mobile industry has developed resilient protocols, processes, and monitoring systems to enable security and reduce the potential for fraud and other malicious activities. For example, 3G and 4G technologies use mutual authentication to verify the identity of the Endpoint and the network. This process helps ensure that adversaries are unable to intercept communications.

Furthermore, network technology can be secured through the use of the SIM and technologies such as GBA [8] or EAP-SIM [9]. By using these technologies, the SIM can be provisioned with a session security key that can be used in communications with application network peers over well-known protocols. This process can diminish the potential for adversaries to manipulate the application protocol to compromise the devices or service. Thus, it is possible to secure both the network and the application with this model.

4 The IoT Model

The figure below shows the standard IoT model used throughout these documents is depicted as components of the service and endpoint ecosystems. Each component is composed of sub-components, which are detailed in a document that focuses solely on the primary component. For example, the Endpoint component, and its respective risks, are outlined in the Endpoint Ecosystem document [3] provided within this document set and the Service components are outlined in the Service Ecosystem document [4].




Figure 2 – Example IoT Model

In almost all modern IoT service or product models, this diagram defines the primary components that are required when deploying a production-ready technology.

Communications network components are inherent to IoT and, for the purposes of this model, provide the connection between the two ecosystems with each 'end' of the communication link discussed within the appropriate Endpoint Ecosystem and Service Ecosystem document.

Specific network security guideline recommendations for Network Operators can be found in the GSMA's "IoT Security Guidelines for Network Operators" [13].

4.1 Service Ecosystem

The Service Ecosystem represents the set of services, platforms, protocols, and other technologies required to provide capabilities and collect data from Endpoints deployed in the field. This ecosystem typically gathers data from Endpoints and stores them within its server environment. This data can be rendered to the user by handing elegant visual depictions of the data to various user interfaces. This data, often in the form of metrics, parameters or commands, can also be handed off to authorized third parties via an API originating at the service infrastructure, which is commonly how IoT Service Providers monetize the service.

The Service Ecosystem security guidelines to be used in conjunction with the process described in this overview document can be found in CLP.12 IoT Security Guidelines for IoT Service Ecosystem [4]

4.2 Endpoint Ecosystem

The Endpoint Ecosystem [4] consists of low complexity devices, rich devices and gateways that connect the physical world to the digital world in via several types of wired and wireless networks. Examples of common Endpoints are motion sensors, digital door-locks,

automotive telematics systems, sensor-driven industrial control systems, and more. Endpoints gather metrics from the physical environment around them, and push that data in different formats via a capillary or cellular network to the Service Ecosystem, often receiving instructions or actions in response. They may also include rich user interfaces that render data obtained either through the Endpoint itself, or from the Service Ecosystem.

The Endpoint Ecosystem security guidelines to be used in conjunction with the process described in this overview document can be found in CLP.13 IoT Security Guidelines for IoT Endpoint Ecosystem [13]

5 Risk Assessments

While the concept of a risk assessment has been around for many decades, many businesses are more familiar with applying the concept to general business risk than to information security. However, an information security risk assessment process is also imperative toward the secure operation and longevity of the technological side of a business. Obviously, in Internet of Things technology, where the engineering team is a critical component to the success of the business, the risk assessment process should be the first step the organization takes to building a security practice.

While every organization should create a granular perspective of technological risk, there are high level questions that function as starting points for the risk assessment process

- What assets (digital or physical) need to be protected?
- What groups of people (tangible or intangible) are potential threat actors?
- What is a threat to the organization?
- What is a vulnerability?
- What would the result be if a protected asset were compromised?
- What is the probability of the asset being compromised?
- What would the result be when put in context with different *groups* of attackers?
- What is the value of the asset to the organization and its partners?
- What is the safety impact of the asset being compromised?
- What can be done to remediate or mitigate the potential for vulnerability?
- How can new or evolving gaps in security be monitored?
- What risks cannot be resolved and what do they mean to the organization?
- What budget should be applied toward incident response, monitoring, and risk remediation?

These starting points will help the engineering and information technology teams work more effectively with the organization. The goal is to ensure that the technical side of the business agrees on the risks, values, and remediation plans with the executive side of the business. Forcing the teams to work together will help create a more realistic perspective of not only the risk to the business, but the value of assets. This will directly affect the budget that should be applied toward resolving outstanding gaps in security.

There are some risks that simply cannot be resolved. Some of these risks will be discussed in these guidelines. The organization should evaluate these risks and determine whether they are acceptable. This will provide the business with a realistic understanding of their limitations, the technology's limitations, and their ability to react to certain types of threats.

There is nothing more monetarily draining than presuming that all security gaps can be resolved in a cost-effective manner.

5.1 Goal

The goal of a risk assessment is to create (or update) a set of policies, procedures, and controls that remediate, monitor, and respond to gaps in security found in the technical part of the organization. The output of the risk assessment should help the business adjust not only its technology, but the way the technology is managed, designed, and deployed. Once the risk assessment output more adequately describes the value of the information and resources used by the organization, the overall business can be secured through the enhancement of its personnel, processes, and policies.

Remember, the core benefits to using the output of a risk assessment are:

- Informing personnel
- Enhancing processes
- Defining (or updating) policies
- Executing remediation
- Monitoring for new gaps
- Enhancing the product or service

This, essentially helps the organization enforce a base platform for personnel and process security. This platform then should be incorporated into a cycle that constantly assesses and refines the overall roles and responsibilities of the organization.

5.2 Risk Model References

Rather than attempt to define a risk assessment and threat modelling process here, please review the following references for an adequate depiction and walk-through of the risk assessment process:

- National Institute of Standards and Technology (NIST)'s Risk Management Framework [5]
- Computer Emergency Response Team (CERT)'s OCTAVE model [6]

6 Privacy Considerations

Many IoT services and products will be designed to create, collect, or share data. Some of this data may not be considered 'personal data' or impact a consumer's privacy, and therefore, not subject to data protection and privacy laws. This data could include information about the physical state of the machines, internal diagnostic data, or metrics regarding the state of the network.

However, many IoT services will involve data about or related to individual consumers and will be subject to general data protection and privacy laws. Where mobile operators provide IoT services they will also be subject to telecommunications-specific privacy and security rules. 'Consumer' focused IoT services are likely to involve the generation, distribution and use of detailed data that could impact an individuals' privacy. For example, drawing inferences about their health or developing profiles based on their shopping habits and

locations. As consumer IoT services gain in popularity, more consumer data will be created, analysed in real-time and shared between multiple parties across national borders.

Where data relates to specific individuals, this complex, 'connected' ecosystem may raise concerns from the consumer over:

- Who is collecting, sharing and using individuals' data?
- What specific data is being acquired?
- Where is the data being acquired from (what technologies or interfaces)?
- When is the data being collected?
- Why is the data being collected from the user?
- How the privacy (not just the security) of individuals' information is ensured?
- Are individuals in control over how their data is shared and how companies will use it?

All providers of IoT services that rely on consumer data – as well as any partner companies capturing or using such data – have an obligation to respect individuals' privacy and keep personally identifiable or privacy-invasive information secure.

A key challenge for IoT service providers is that there are multiple, and often-inconsistent, laws dealing with privacy and data protection. Different laws may apply in different countries, depending on the types of data involved, as well as the industry sector and services that the service provider is offering. This has implications for a number of consumer-oriented IoT service providers;

A connected vehicle, for example, can move between different countries, meaning the associated data transfers may be governed by several different legal jurisdictions. In-car sensors tracking the location of the car (static or dynamic) and its frequent destinations could be used to infer a number of insights about the driver's lifestyle, hobbies or religion, which the driver may consider personal information. Additionally, insights about driving habits through 'on-board diagnostics' sensors might be shared with insurance companies who might use those insights to impose a higher premium and therefore discriminate against the driver without their knowledge.

IoT services and devices (including connected cars) can also move between different sovereign territories and therefore different legal jurisdictions. In many cases, an individual's personal data may transit or reside in jurisdictions different from the individual. These are important issues that need to be considered before a multi-national IoT Service is deployed.

Another challenge is that most data protection laws require companies collecting consumers' data to get the affected consumer's (also known as the 'data subject') consent before processing certain categories of 'personal data' – such as health related data. Most laws define 'personal data' as any information that relates to an 'identified' or 'identifiable' living, natural person.

But as more and more devices are connected to the Internet, more and more data about individuals will be collected and analysed and possibly impact their privacy, without necessarily being considered 'personal' by law. The combination of massive data volumes, Cloud storage and predictive analytics can provide detailed profiles of users. In particular, it

may become challenging to truly anonymise information and personal information can be inferred from other data types.

The need to maintain the privacy of sensitive, health data records is well recognised, not least due to the potential for commercial abuse of such records. In the United States of America, the Health Insurance Portability and Accountability Act of 1996 (HIPAA) includes privacy and security requirements to mitigate the risks of unauthorised disclosure of health records.

HIPAA, like many other regulations such as those in the European Union, only applies if the health data is *personally identifiable*. The data stored in a blood monitoring device (which does not identify the user) would not be covered by these requirements, whereas that same data in a smartphone app or in a Cloud server is likely to be covered because it is able to be linked to an individual (in the case of a smartphone because the phone will almost certainly contain other data identifying the user and in a Cloud server because it will be associated with an identifiable user account). Policymakers around the world are realising that information and insights about people can impact their privacy even if they are not defined as 'personally identifiable'. They are therefore beginning to adopt more risk-based approaches to regulation but also considering the wider privacy implications of data use rather than focusing on legal definitions.

In order to build trust in the IoT ecosystem Governments should ensure data protection and privacy legislation is technology-neutral and that rules are applied consistently to all players in the internet ecosystem. Furthermore, in order for IoT Service Providers to minimise the need for formal regulatory intervention, we recommend that they follow the steps described in Annex A at the early development stages of their IoT service or product.

7 Using This Guide Effectively

While security is best implemented at the start of an engineering project, this guide can also assist in organizations that have already designed, fabricated, and even deployed an IoT product or service. Regardless of which stage the reader's product or service has reached, there is a useful process that should be followed to get the most benefit from this set of documents:

- Evaluate the technical model
- Review the current product or service's Security Model
- Review and evaluate Recommendations
- Implementation and Review
- Ongoing Lifecycle

7.1 Evaluating the Technical Model

The first and most important step in the process is understanding the organization's own IoT product or service. In order to perform a security review and risk assessment, the team should be familiarized with each component used in the organization's solution, how components interact, and how the components interact with their environment. Without a clear understanding of how the product or service was (or will be) built, a review will be incomplete.

Start by making a document describing each component used in the system. Identify how the component is sourced, how it is used, what privilege level it requires, and how it is integrated into the overall solution. Map each component to the technologies described in the Model section of each Endpoint Ecosystem [3] and Service Ecosystem [4] guidelines documents. It is acceptable if the document doesn't specifically match a component, as it should map the component's general class. Simply use the class of component, such as a microcontroller, communications module, or trust anchor, as the context. Consider the following questions:

- What components are used to build the product or service?
- What inputs and outputs are applicable to the given component?
- What security controls are already applied to these inputs and outputs?
- What privilege level is applied to the component?
- Who in the organization is responsible for implementing the component?
- Who in the organization is responsible for monitoring and managing the component?
- What process is in place to remediate risks observed in the component?

These questions, when answered, will provide an understanding of how the technical components interact with each other, and how the overall product or service is affected by each component.

This process corresponds with the first and second phases of the CERT OCTAVE risk assessment model [6], or the Frame stage of the NIST Risk Management Framework [5]. This assists in the development of a profile for each critical business asset, the development of security objectives, and establishes a foundation for how the company will assess, monitor, and respond to risk.

7.2 Review the Current Security Model

Next, read through the security model section of the Endpoint or Service being assessed. This section will help the reader understand the model that an Attacker will use to compromise a given technology. This model is based on years of experience performing security assessments on, reverse engineering, and designing embedded systems.

Once the security model has been reviewed, the reader should have a better understanding of what technologies are most vulnerable, or most desirable to the Attacker, in the product or service being developed. This information should be shared with the organization, to ensure that both engineers and leadership understand the risks and threats to the current model.

However, it should be noted that the organization should *not* take steps to adjust their security model at this time. It is too early to make concise architectural changes.

This process again corresponds to the first and second phases of the CERT OCTAVE model [6], or the Frame stage of the NIST Risk Management Framework [5]. Reviewing the security model helps enhance the technical model by identifying potential gaps in security and shining a spotlight on security objectives that should be prioritized.

7.3 Review and Evaluate Recommendations

The Recommendations section should be reviewed at this time to evaluate *how* Security Tasks can be resolved. This section will not only provide methodologies for implementing

recommendations, but will provide insight into the challenges involved in implementing the particular recommendation.

For each recommendation, a *Method* section is provided. This section will outline methodologies that assist in the remediation or mitigation of the corresponding security risk. These methods, while presented from a high level, will outline concepts that reduce risk from a holistic perspective, to ensure the greatest amount of gain is acquired from a reasonable and practical amount of effort.

An *Expense* section is provided to discuss, where applicable, extra financial expenses that the organization should prepare for when implementing a particular recommendation. While most expenses, such as engineering time and raw materials, are fairly obvious, less obvious expenses can alter the finances applied to products and services whose profit margins and budgetary limits have already been defined by the business leadership. While specific numbers are not provided, technologies and services are specified that may incur additional costs.

A *Risk* section is also provided so the reader understands the gaps in security that are likely to result from *not* implementing a particular recommendation. While the business may accept that some risks are within the business's operating guidelines, the reader should review each risk section to ensure that the business fully understands the side effects of not implementing (or not correctly implementing) a given recommendation. This may seem straight forward for recommendations such as "Encrypt Data", but the subtlety of some threats, such as replay attacks against messages that are not cryptographically unique, may be a surprise to the reader at a later date.

In some cases, *references* are provided for further review. While this document does not provide detailed information on every technology, risk, or remediation plan, other standards and time-proven strategies do. This set of documents will provide references to those materials, where applicable, within each recommendation.

The output from reviewing the Recommendations section should directly tie into the Security Tasks section. The Security Tasks should now be filled out with Recommendations that are appropriate for implementing the Security Tasks correctly. These Security Tasks will then tie back to specific Components assigned to members of the organization.

Evaluating recommendations corresponds to the Assess step of the NIST Risk Management Framework [5], and steps six, seven, and eight of the CERT OCTAVE methodology [6].

7.4 Implementation and Review

By this stage, clear Security Tasks have been outlined and the business will have a better comprehension of their security vulnerabilities, their value and their risk. The business shall now create a clear architectural model for each Component being adjusted, and use the Risk Assessment process chosen by the organization to develop a threat model of each Component, incorporating the Recommendations and Risks that are appropriate for each Component and Security Task. When the architectural model is completed, the organization can begin implementing each Recommendation in order to fulfil the Security Tasks.

When the implementation is complete, the organization should review the Risks in both the Recommendations subsection and the Component sections. The organization should ensure

that the implementation fulfils the requirements set forth by these sections. The organization should then ensure that the implementation solves security with regard to the context in which the Component is designed in the organization’s product or service, as these documents cannot fully address every product or service being designed in the field. If possible, have a third party consulting firm evaluate the implementation to ensure that it does indeed adhere to security best practices.

Implementation and review corresponds with the Respond component of the NIST Risk Management Framework [5], and step eight of the CERT OCTAVE model [6].

7.5 Ongoing Lifecycle

The security life cycle does not stop at this juncture. Rather, security is an inherent part of the overall engineering of a process. Endpoints and IoT Services have a lifetime, and must be continually serviced throughout that lifetime, just like a living organism.

Requirements change over time. Cryptographic algorithms become dated or deprecated. New protocols and radio technologies must interoperate with the product or service. This ever changing ecosystem our embedded products are deployed in must be constantly reviewed to ensure that confidentiality, integrity, availability, and authenticity are maintained.

Managing the ongoing security lifecycle corresponds with the Monitor and Frame components of the NIST Risk Management Framework [5], and steps one, four, and five of the CERT OCTAVE model [6].

8 Example – Wearable Heart Rate Monitor

In this example, a simple Heart Rate Monitor (HRM) design will be evaluated using this set of guidelines. The endpoint will be assessed using the Endpoint Ecosystem document, while the service side of the design will be assessed using the Service Ecosystem document.

8.1 The Endpoint Overview

First, let’s start by evaluating the hardware design of the endpoint.




Figure 3– Simple HRM and Primary Components

The HRM is composed of standard components for a simple wireless wearable device: an ambient light photo sensor and a Bluetooth Low Energy (BLE) transceiver enabled microcontroller. The sensor is used to capture pulse rate data, while the microcontroller

analyses the data emitting from the sensor and chooses what data to send over the built-in BLE transceiver. In this example, the BLE stack used is version 4.2.

A coin cell battery is used in this example to transmit data from the HRM to another device, such as a smart-phone or tablet. No other components are required for this device to function.

According to the Endpoint Ecosystem document, this device would fit into the Lightweight Endpoint class of devices.

8.2 The Service Overview

From a service perspective, the application on the smart-phone or tablet pushes metrics from the endpoint up to a back-end service over any available network connection. The back-end service for the application simply associates the device owner with the metrics being captured and stores them in a database local to the application server.

Visualization of the data can be achieved using the mobile application, or via the service's website. Users of the wearable technology can log into the service provider's website to perform more actions with the metrics captured by the endpoint.

This is a very simple and common service model with no custom or unnecessary complexities.

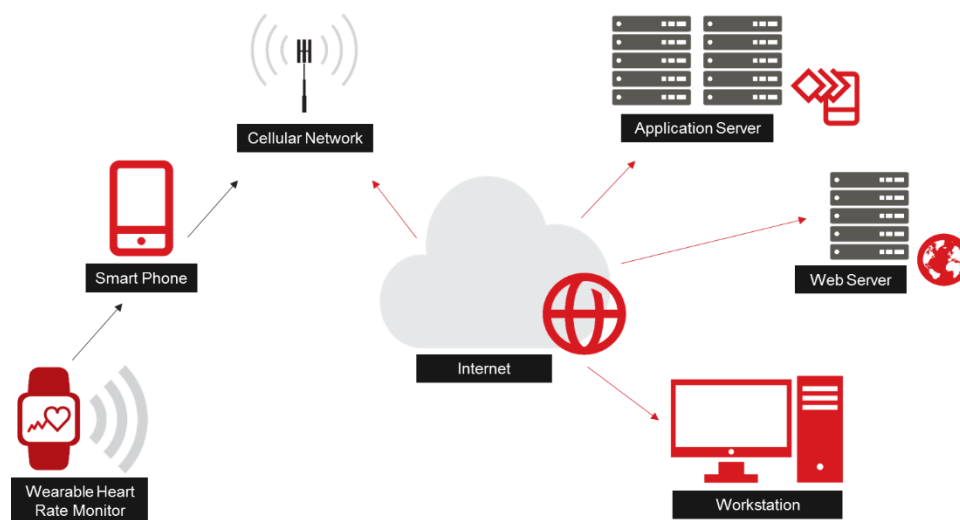


Figure 4– Flow of Data to Simple Back End Service

8.3 The Use Case

The business developing this technology intends the end user to track their pulse data throughout the day, storing it in both the application and the back-end database. The intention is to allow users to review their heart rate over time to track their overall health. Users can watch their health improve or worsen over time, depending on whether they are maintaining a healthy life style. This allows the users to incentivize themselves by evaluating both positive and negative trends in their HRM data.

The business intends to use this data to partner with medical device manufacturers, health care providers, and other organizations that can use these metrics to identify whether a consumer is more or less likely to incur a health-related event, such as a heart attack or a stroke.

8.4 The Security Model

The engineering team at this example business leveraged the Frequently Asked Security Question sections of the Endpoint and Service documents to determine what issues are most relevant to their product and service.

From an endpoint perspective, the team learned the following issues are of concern:

- Cloning
- Endpoint impersonation
- Service impersonation
- Ensuring privacy

From a service perspective, the team decided the following issues are of concern:

- Cloning
- Hacked services
- Identifying anomalous endpoint behaviour
- Limiting compromise
- Reducing data loss
- Reducing exploitation
- Managing user privacy
- Improving availability

The team reviewed the recommendations for each of the above issues, as suggested by each relevant Frequently Asked Security Question section. The team then chose to implement recommendations that were cost-effective improvements that ensured the greatest amount of security.

In this example model, the endpoint would not require a substantial change. Since the endpoint has very little functionality, minimal security can be employed on the endpoint for both application security and communication. Since the endpoint application is flashed on a single device, as long as the device firmware is *locked*, there is no real threat of attack against the endpoint within the given use case.

However, since privacy is an issue, the organization should employ at least a Personalized PSK version of a Trusted Computing Base (TCB). This would ensure that encryption tokens were unique to each endpoint, so that one compromised endpoint cannot compromise all endpoints. If the personalized (unique) keys were encoded into the locked microcontroller, it would be reasonable to believe that this use case were adequately secured from the threat of cloning, impersonation, and privacy issues. Review the IoT Service [3] and Endpoint [4] documents for a more complete discussion on what a Trusted Computing Base is within each ecosystem's context.

The server infrastructure, however, requires a significant amount of changes. The engineers realize that, according to the recommendations, they are at serious risk of abuse. The following issues are acknowledged:

- There is no security front-end diminishing the effects of a Denial of Service attack
- There are no ingress or egress controls limiting the flow of traffic to or from services
- There is no separation of duties between service tiers
- There is no separate secured database containing Personalized PSK tokens
- No adequate security measures are implemented in the service operating system
- There are no metrics taken to evaluate anomalous endpoint behaviour

8.5 The Result

After implementing the recommendations, the organization has a much better defined back-end service architecture that adequately addresses the risks identified through the guidelines.

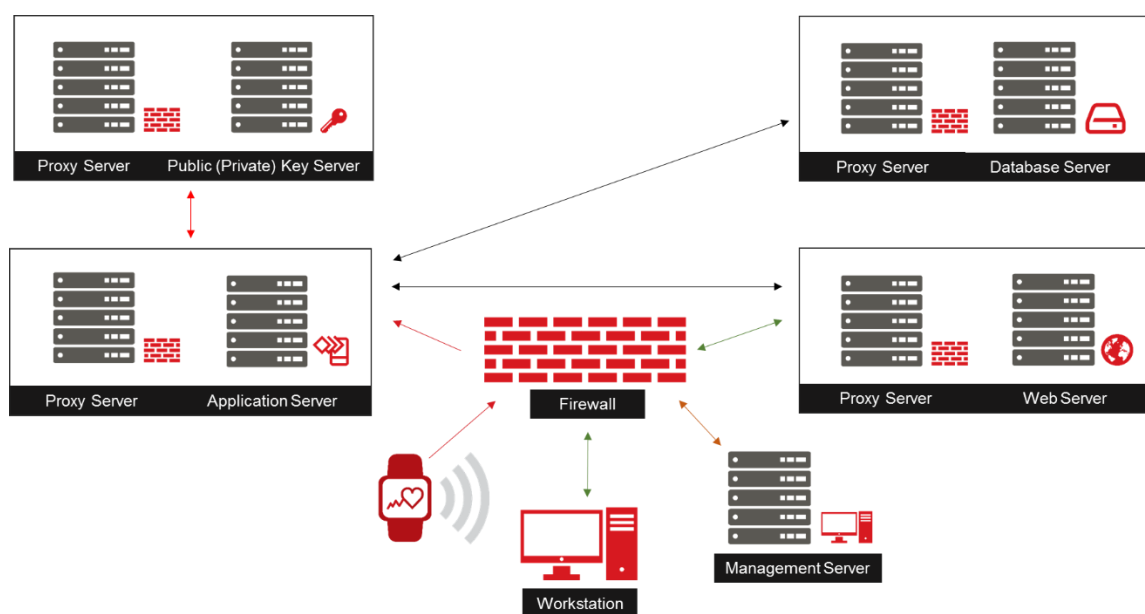


Figure 5– Resultant Service Ecosystem

In the above figure, the changes to the service ecosystem are easily observable. Each class of service has been broken into separate tiers to help secure *and* scale the technology easily in the event that demand spikes. Two additional tiers were added, a database tier and an authentication tier, to separate critical systems from services that directly interface with the outside world. A security front-end was implemented to help guard the internal network from multiple types of attacks, including DoS and DDoS attacks that reduce the overall availability of the system. Finally, an administrative model was defined to allow management secure access to the production environment. One component not depicted in the above diagram is the presence of an analytics model that observes when endpoint behaviour may be indicative of a compromise, or a flaw in the firmware or hardware design.

8.6 Summary

Overall, this simple technology could have been easily compromised had it been deployed “as is”. Yet, with a few fast, simple, and cost-effective changes made on the endpoint, the technology is assured to have years of longevity in the field without change to the architecture.

With the service ecosystem ramped up, there is far less of a threat to both users *and* the business. Cloning and impersonation is no longer a threat. Privacy is ensured by granting each endpoint unique cryptographic tokens. Systems that contain critical information are separated and secured from more heavily abused public-facing systems. This model, while slightly more complex, reduces the overall risk of the production environment.

9 Example – Personal Drone

In this example, a small personal drone device will be evaluated using this set of guidelines. The endpoint will be assessed using the Endpoint Ecosystem document, while the service side of the design will be assessed using the Service Ecosystem document.

9.1 The Endpoint Overview

First, let’s start by evaluating the hardware design of the endpoint.




Figure 6– A Drone and its Primary Components

This personal drone is composed of a robust set of components. The processing capabilities of the drone are high performance due to the multiple motors, sensors, and other equipment that all must function efficiently in parallel. This model uses an ARM Cortex-A8 CPU with the primary operating system (Linux) stored in NVRAM on a separate chip. An array of various sensors are required for detecting movement, light, speed, and more. A SD/MMC card is used to store video, sensor metrics, and metadata. A camera is used to allow the operator to see from the drone’s perspective. A cellular/GPS combination module is used to ensure the drone can maintain connectivity to its operator even when it is out of range of a proprietary protocol. GPS is also used for guidance, and for minimal automation.

A Lithium Polymer (LiPo) battery is used to drive the drone. Its fly time is approximately two hours before a new charge is required when all functions are active at once.

According to the Endpoint Ecosystem document, this device would fit into the Complex Endpoint class of devices. Even though it contains a cellular module, it is not considered a gateway as it does not route messages to or from other endpoints.

9.2 The Service Overview

From a service perspective, the back-end is only used for operator connectivity when loss is detected on the proprietary radio interface during flight. If the drone is in flight and the cellular connection can be enabled, it will attempt to wait for its operator to connect via the LTE network. If, however, it is unable to be controlled over LTE, it will attempt an automated landing at the location where it last lifted off.

However, as the drone has some light automation features, it can be given coordinates and a path to traverse while taking photos or short videos. These media files can be uploaded in real time over LTE to the back-end service to show the operator its course and viewpoint during automated execution.

Thus, a robust back-end service is required to ensure a high degree of service availability for each drone that might connect to the system. Availability is also necessary for the high bursts of network traffic required to transmit videos and high-resolution images over a cellular link. There must also be a web interface that allows the operator to view media uploads from a web browser.

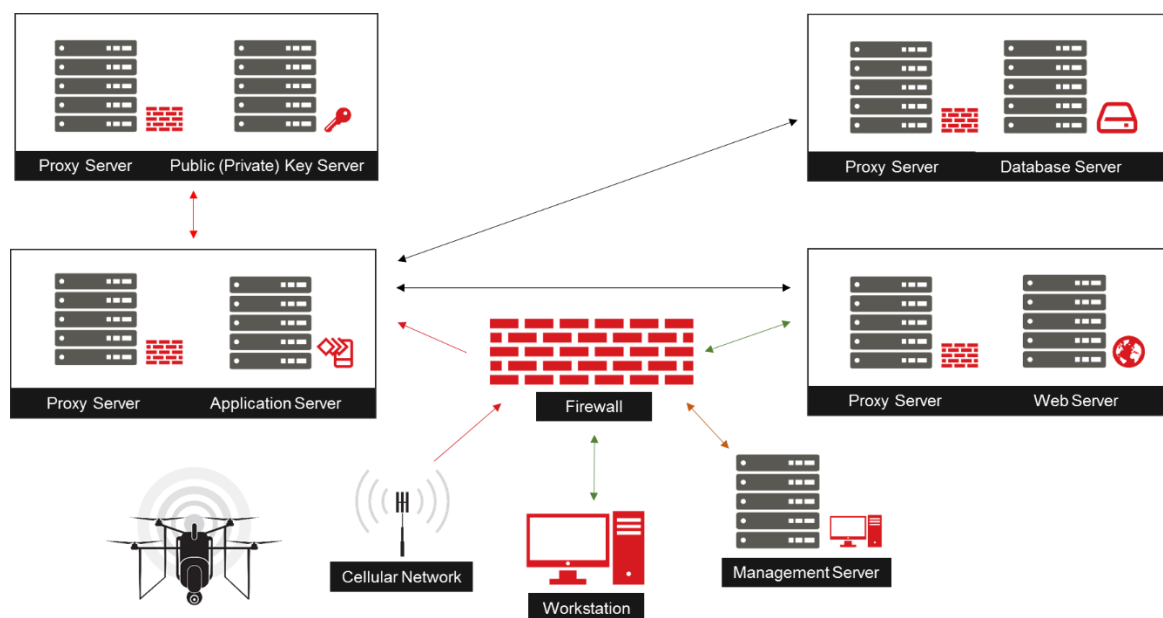


Figure 7– Flow of Data to Back End Services

9.3 The Use Case

The business developing this technology intends the end user to use the drone for filming in the wild. However, some of their customers have used the drone for filming scenes in

cinema, as the camera and stabilization capabilities of the drone are exceptional for the price point. As a result, the drone will be used in expensive filming projects where intellectual property and privacy are major concerns.

9.4 The Security Model

The engineering team at this example business leveraged the Frequently Asked Security Question sections of the Endpoint and Service documents to determine what issues are most relevant to their product and service.

From an endpoint perspective, the team learned the following issues are of concern:

- Endpoint identity
- Endpoint impersonation
- Trust anchor attacks
- Software and firmware tampering
- Secure remote management
- Detecting compromised endpoints
- Service impersonation
- Ensuring privacy

From a service perspective, the team decided the following issues are of concern:

- Managing user privacy
- Improving availability

The team reviewed the recommendations for each of the above issues, as suggested by each relevant Frequently Asked Security Question section. The team then chose to implement recommendations that were cost-effective improvements that ensured the greatest amount of security.

In this example, the service infrastructure does not require a substantial change. This is because the service infrastructure already had to be built out extensively to accommodate for the bursts of traffic required in servicing the endpoint product. The architecture already demanded a well formed and secure architecture simply to be able to scale effectively and maintain availability of resources even when some services were incurring temporary faults. However, the organization chose to investigate user privacy further as this has become a primary point of contention for the business's unexpected niche.

The endpoint infrastructure, however, requires a significant amount of changes. The engineers realize that, according to the recommendations, they are at serious risk of abuse. The following issues are acknowledged:

- The bootloader is not properly validating the application prior to executing the operating system kernel, leading to a risk of tampering
- There is no TCB used to manage the security of the application or communications
- Because there is no properly implemented TCB or trust anchor, endpoint impersonation is a problem, which may lead to data leakage
- Without a well implemented TCB, the endpoint can't properly authenticate services

- Without a well implemented TCB, the endpoint can't properly authenticate the operator over the proprietary radio interface
- The engineers have relied on the security of LTE to ensure the communications channel can't be compromised, but has not considered the threat of endpoint impersonation or Femtocell repurposing, both of which bypass the security of LTE to compromise weak service security

9.5 The Result

After implementing the recommendations for the issues cited above, the organization has a much better defined endpoint architecture that adequately addresses the risks identified through the guideline documents.

For the existing drone system already in production, the engineering team issues a firmware update that implements a Personalized Pubkey security model. The firmware update improves the bootloader as well to bake security into the core architecture. Since a Personalized Pubkey model was used, anyone attempting to abuse the initial lack of security in the endpoint to attempt to impersonate another user's endpoint would fail, as the engineers leveraged their existing user-to-endpoint mapping database to create personalized keys on a per-user basis. This way, no user without the appropriate web credentials can download and install another user's Personalized Pubkey update. While this process was complex and time consuming to implement, it will be worth the effort.

Future versions of the drone technology will implement an internal CPU trust anchor. This trust anchor will be tied to a Personalized Pubkey TCB, to ensure that each endpoint is uniquely seeded with exceptional security from the ground up.

Deploying strong cryptography in this fashion is imperative, as it also negates the potential for the other classes of attack the company identified as a concern. By leveraging the benefit of strong cryptography and a TCB for verification and authentication, the engineering team can easily identify whether rogue services are being made available to the drone. The drone, upon detecting rogue services, can simply land back at the original take-off site.

Any service that detects an improperly secured drone can also raise flags internally. The administration team, at that time, can determine how to deal with the potentially compromised drone. This provides a level of agility with regard to security events, and also gives the organization a way to evaluate if there are software or hardware problems that are causing abnormal behaviour on the endpoint.

9.6 Summary

While the engineering team obviously spent an exceptional amount of time creating a resilient architecture from a mechanical engineering and back-end services perspective, substantial work needed to be done to create secure endpoint technology. While this scenario did not pose a critical threat to the overall business, it was fortunate that there was a solution that worked *well enough* for their customer's needs. Had this been a more safety-critical technology, even the solution deployed here may have not been sufficient.

For more information on Trusted Computing Base variants, such as Personalized Pubkey TCB or Personalized PSK TCB, please review the IoT Service [3] and Endpoint [4] Ecosystem documents.

10 Example – Vehicle Sensor Network

In this example, a vehicle sensor network deployed in a new class of automobile will be evaluated using this set of guidelines. The endpoint will be assessed using the Endpoint Ecosystem document, while the service side of the design will be assessed using the Service Ecosystem document.

10.1 The Endpoint Overview

First, let's start by evaluating the hardware design of the endpoint.




Figure 8– Full Vehicle Sensor Network and Communications System

While the above model is too complex to properly depict in a simple diagram, the three high-level components involved are:

- A telematics uplink unit that manages the sensor network, makes complex decisions on behalf of the driver, and maintains a connection to the back-end system
- A vehicle-to-vehicle (V2V) system that detects and reacts to V2V events
- A general sensor network that provides metrics to the telematics uplink unit

In modern automotive systems, the telematics unit is a part of the automobile's computer network and makes decisions based on sensor data and back-end communications. This unit will make decisions with, or on behalf of, the consumer driving the vehicle. The unit ensures that the vehicle is operating properly, attempts to make intelligent decisions during emergencies, and takes commands from the back-end network.

The V2V sensor network identifies vehicles in the vicinity and makes decisions based on metrics gathered from sensors. While the telematics unit primarily makes decisions based on the state of components (such as brakes or tire pressure monitors) the V2V system makes decisions based on the presence of other vehicles, or sends out alerts to nearby vehicles in the case of a critical event.

The general sensor network is a series of components that provide data to the telematics unit, and sometimes the V2V unit. These units use the information gathered from the general sensor network to make accurate decisions during critical events.

According to the Endpoint Ecosystem document, this system has components that fit into every IoT endpoint class. The telematics uplink unit acts as a gateway. The V2V unit acts as a complex endpoint. The general sensor devices are effectively all lightweight endpoints.

10.2 The Service Overview

From a service perspective, the vehicle sensor network will provide metrics to the back-end environment. This data may or may not be provided to the consumer. Rather, the data could be stored by the manufacturer to observe or identify potential problems with components. This may trigger service warnings that are then issued to the consumer.

The system may also be augmented to provide the consumer with useful services, such as “remotely unlock door”, “start engine”, and similar features. In the near future, these systems may allow vehicles to be driven remotely through automated guidance systems.

While most critical decisions will be made in the processing units on the vehicle itself, it is reasonable to conjecture that some decisions will be made in the cloud, where more machine learning (ML) and artificial intelligence (AI) along with behavioural or statistical models can be leveraged to make more complex decisions.

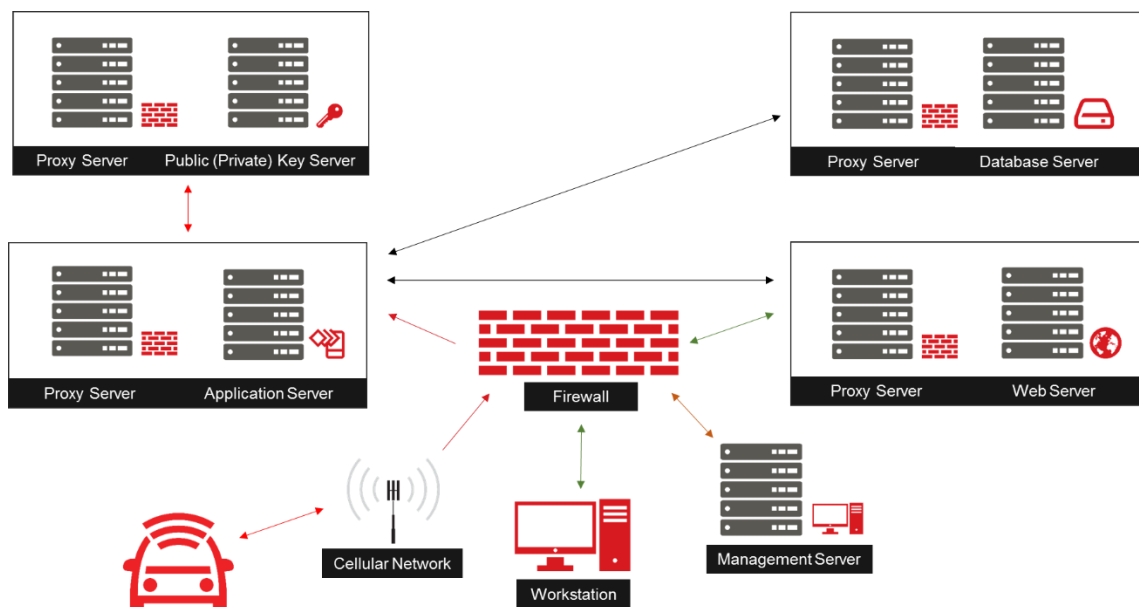


Figure 9– Flow of Data to Back End Services

10.3 The Use Case

The use case of this technology is obvious: to build smarter vehicles that can make complex decisions in safety-critical scenarios. The goal is to leverage the intelligence of as many sensors as possible to make critical decisions in very small windows of time. Automatic breaking, tire blow-out broadcast alerts, temporarily disabled operator warnings, and other

critical scenarios can potentially be resolved through the use of sensors and well designed computer systems.

One interesting feature of this technology is that it may be entirely transparent to the user. The user would not need to configure these computers to act in a certain fashion. Instead, they should be capable of negotiating the current landscape through the use of sensor metrics. This will allow the computers to behave correctly regardless of the environment.

10.4 The Security Model

The engineering team at this example business leveraged the Frequently Asked Security Question sections of the Endpoint and Service documents to determine what issues are most relevant to their product and service.

From an endpoint perspective, the team learned the following issues are of concern:

- Endpoint impersonation
- Service or Peer impersonation
- Side-channel attacks
- Detecting compromised endpoints
- Ensuring safety at the risk of security

From a service perspective, the team decided the following issues are of concern:

- Identifying anomalous endpoint behaviour
- Managing user privacy

The biggest risk to this environment that hasn't been discussed in previous examples is the risk of impersonation with regard to peers. One concern that engineers have in this type of environment is the risk that a computer will make critical decisions using data that is not properly authenticated.

Since sensor data in critical scenarios requires exceptionally fast processing times, it is theorized that it may not always be feasible to implement asymmetric cryptography or PKI based communications. However, this may not be an accurate assertion. Instead, an accurate security model should account ahead of time for time-critical scenarios and cache session keys for nearby Endpoints. For example, if two objects are approaching each other at a known rate, security applications in the Service Ecosystem can prepare session keys specific to these two Endpoints before they reach a distance where they can physically impact one another. This would ensure that secure communication between Endpoints and sensors can still be used in the event that there is no time to renegotiate an instantaneous secure session when the potential for a critical scenario (like an impending automotive crash) is detected.

Thus, an augmentation to the TCB implementation is required. One interesting solution is GBA, where the UICC used in the telematics uplink unit can distribute keys securely to endpoints throughout the system. This protocol will allow even rudimentary endpoints to be seeded with secure session keys that can be used in multiple critical scenarios. This way, the environment can always be seeded from a root of trust, even if lightweight endpoints are not capable of critical maths for public key session initialization.

Another critical issue in these environments is detecting compromised endpoints. For example, how can the environment recognize whether a simple sensor, such as a Tire Pressure Monitor (TPM) has been compromised? If the computer makes a critical decision based on the TPM signalling a tire has blown, a safety issue may arise. As a result, the behaviour of devices, and their trustworthiness, must be reassessed at every boot-up phase. All devices should have tamper resistance, and must be able to notify the network if there is a compromise. Inversely, there should be a way that other devices in the sensor network can evaluate the trustworthiness of peers in the network.

10.5 The Result

After implementing the recommendations, the vehicle sensor network is well guarded against attacks on the vehicle communications network. GBA is used to distribute keys to all endpoints in the system, and does so on every boot-up, ensuring that old keys are not reused. This, along with tamper resistance, a strong TCB in every endpoint, and an organizational root of trust, allows the environment to function with far less risk.

Yet, regardless of these changes, safety is still a critical factor. The engineering team and business leadership, along with the company's legal team and insurance brokers, should evaluate safety critical technology and determine whether security can be implemented without risking safety of the users. While security can often be implemented, even in safety-critical scenarios, with some architectural adjustments, there are times when safety must come before all other concerns.

10.6 Summary

Systems like these are often well engineered and take a large amount of effort to attack the ecosystem. However, subtle flaws in the communications architecture can lead to a compromised environment. In walled gardens, such as some CANbus networks, a single flawed endpoint can cause the entire system to become vulnerable. This, in safety-critical environments, is unacceptable.

Annex A Recommended Privacy Considerations for IoT Service providers

In order to build trust in the IoT ecosystem and minimise the need for formal regulatory intervention, the GSMA proposes the following high-level steps as a guide to minimising any privacy risks. We recommend that IoT Service Providers follow these steps and consider these questions at the early development stages of their IoT service or product.

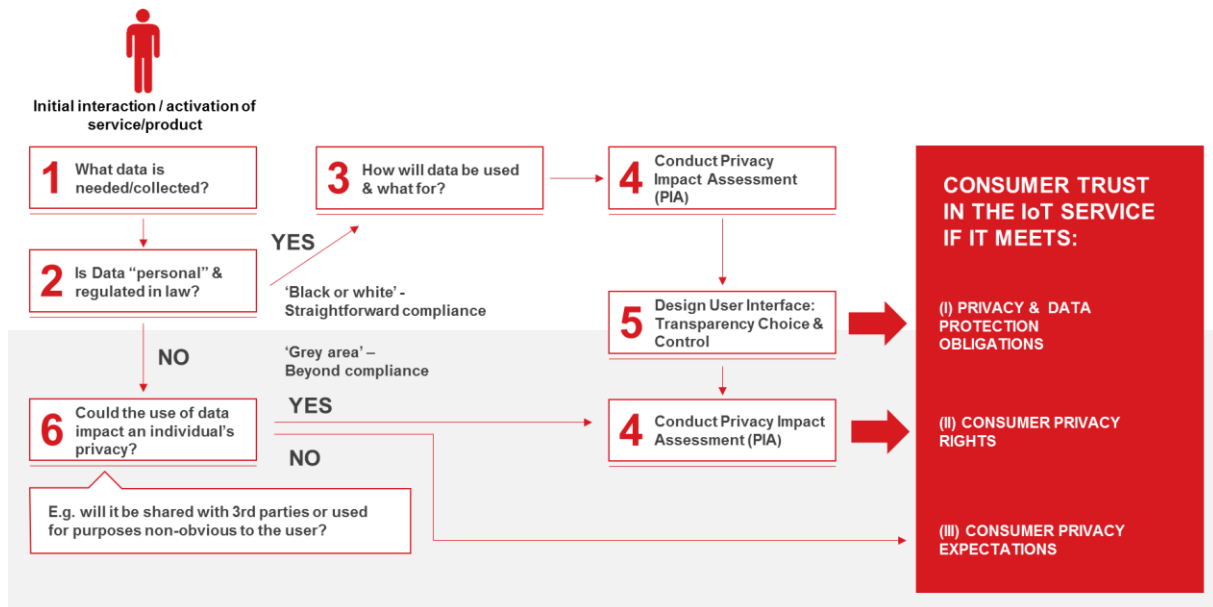


Figure 10– GSMA IoT Privacy by Design Decision Tree

Step	Consideration
Step 1	<p>What data do you need to collect from / about the user so that your IoT service or product can function properly?</p> <p>One of the first steps in any business model relying on data is to identify what information is actually required from or about the consumer, for the service or product to function properly. The types of data a service requires could be categorised as static – such as the consumer’s name or home address – and data that is dynamic, such as real-time location. So if you are offering, for example, a fitness wristband tracking someone’s steps and calories burned, then you would need to know the weight, age, gender, distance travelled and the heart rate of the individual wearing the wristband, but you would arguably not need the actual location of the individual.</p> <p>When assessing the types of data needed, it’s also important to decide whether the individuals’ consent is needed to use that data and how you would obtain their consent or indeed offer them options to control their privacy preferences. A smartphone could act as a medium for offering the user privacy options (e.g. mobile app or online dashboard) where the product itself has no screen.</p>

<p>Step 2</p>	<p>Is the data ‘personal’ and regulated in law?</p> <p>The next step should be to identify the data protection and privacy requirements that the law imposes on you. Questions to consider include:</p> <ul style="list-style-type: none"> • What is the definition of ‘personal’ data in the country/market concerned? • Is the data collected ‘personal’ & regulated in law? If so, have you identified the legal basis that allows you to process such data? • Are you subject to any privacy-related licence conditions (e.g. as a telecoms provider) • Are there any federal, state, local or sector-specific laws that apply in relation to your proposed data collection model, in addition to general data protection laws? e.g.: <ul style="list-style-type: none"> ○ Financial / payment services, healthcare regulations ○ Potential restrictions on cross-border data transfers
<p>Step 3</p>	<p>How will data be used and what for?</p> <p>Once you have established what your legal compliance requirements are, the next step is to map out how the data you collect will be used – and who they need to be shared with – to achieve intended outcomes as part of your service offering. The following questions should help you address both security and privacy considerations in relation to the treatment of the data:</p> <ul style="list-style-type: none"> • Is the data kept secure both when stored and transmitted? • Have you clearly set out the data flows? I.e. identify how the data will be used and shared across the value chain and for what purposes • Can you justify why each type of data collected is needed in the specific context of offering the intended service? • Have you defined/agreed privacy responsibilities with your partners from the outset (and does your product design reflect these responsibilities?) • Are there appropriate contractual agreements in place with the companies you are sharing consumers’ data with? (E.g. limiting the use of data by analytics providers for their own commercial purposes). Such agreements or restrictions can be bi-lateral or you could establish a code of conduct or guidelines and ask your partners to commit to them with defined consequences and liabilities if they fail to do so.

<p>Step 4</p>	<p>Conduct a Privacy Impact Assessment</p> <p>Conducting a Privacy Impact Assessment (PIA) is about:</p> <ul style="list-style-type: none"> • Identifying what, if any privacy risks your product or service raises for individuals. • Reducing the risk of harm to individuals that might arise from the possible misuse of their personal information • Designing a more efficient and effective process for handling data about individuals <p>PIA requirements are increasingly becoming common in data protection and privacy laws. There are a number of guides on how to conduct a PIA including those published by the UK's Information Commissioner's Office [10] and those by the International Association of Privacy Professionals.</p> <p>Typical questions to be addressed when conducting a PIA include:</p> <ul style="list-style-type: none"> • Will the project result in you/your partners making decisions or taking action against individuals in ways that can have a significant impact on them? • Is the information about individuals of a kind particularly likely to raise privacy concerns or expectations? For example, health records, criminal records or other information that people would consider to be private? • Will the project require you to contact individuals in ways that they may find intrusive?
<p>Step 5</p>	<p>Design Privacy into the User Interface</p> <p>After assessing the privacy risks to the consumers, you should consider how to raise those consumers' awareness of such risks and how to mitigate them as well as offer them options to express their privacy preferences. Ultimately, this step is about ensuring you offer a service that meets your legal obligations and the consumers' needs and expectations in a user friendly way. And it's about building their trust by reassuring them that they have more control over their privacy. Questions to consider include:</p> <ul style="list-style-type: none"> • How can consumers be made aware of any risks to their privacy and how can they make informed choices? • Have you obtained their consent, where legally required? Key elements of consent include: disclosure, comprehension, voluntariness, competence, and agreement) • Is data secured in transit and at rest? • Is there a set period for which you need to keep consumer data (and why)? • Does the consumer journey help gain their trust? For example: <ul style="list-style-type: none"> ○ Do they understand what data they are sharing in return for using the service? ○ Can consumers express their privacy preferences in simple steps e.g. via a web based 'permissions dashboard', 'just-in-time' prompts, a call centre, a mobile app, a voice activated command etc.

Step 6	<p>Could the use of data impact an individual's privacy?</p> <p>Your product or service may collect data that are not necessarily classified as 'personal' in law but may still have privacy implications to the consumer and should therefore be considered early on. To ascertain whether the relevant data could be used to impact a consumer's privacy consider the following:</p> <ul style="list-style-type: none">• Could (non-personal) data from your service/product be combined with other data from different sources to draw inferences about a consumer's personal life? For example inferences about his/her lifestyle, habits or religion that would:<ul style="list-style-type: none">○ Affect his/her ability to get health insurance?○ Be used by 3rd parties (retailers, insurance companies) to price discriminate against the specific consumer?• If your product or service is likely to change at any point in the future what are the likely privacy implications of any such change on the consumer. For example:<ul style="list-style-type: none">○ Does the change involve the collection of new data about the consumer (such as location data)?○ Are existing or new consumer data shared or sold to third parties (e.g. advertisers) who would start using consumer data for different purposes than those originally obtained for?• If any such changes occur you should:<ul style="list-style-type: none">○ Check the possible impact on your business if new laws are invoked as a result of the change○ Establish processes to inform the consumers and obtain their consent where necessary○ Provide the means for consumers to change their privacy preferences• Some additional considerations that we recommend IoT service providers consider are:<ul style="list-style-type: none">○ Make sure you have appropriate contractual agreements in place defining the responsibilities of each partner in the value chain○ Have a clear process of redress so that the consumers know who to turn to if things go wrong or if they suffer from a privacy breach
---------------	--

The following diagram presents one option of how the proposed steps above could be illustrated:

Annex B Example based upon Automotive Tracking System

In this example, an automotive tracking system will be evaluated from the perspective of the IoT Security Guidelines. The process will stem from section six of this Overview document – “Using This Guide Effectively”.

B.1 Evaluating the Technical Model

In the first step, “Evaluating the Technical Model”, the engineering team assesses how the device functions based on their product’s architecture. The engineering team creates a document that itemizes the technologies used in the solution in order to organize personnel, assign Security Tasks, and track progress.

For the sake of simplicity, our automotive tracking system will have the following capabilities:

- **Endpoint Ecosystem:**
 - A simple Graphic User Interface (GUI) that allows a user to:
 - Log in with a username and password
 - Disable tracking
 - Enable tracking
 - Identify and visualize current location
 - A cellular module for connecting to back-end services
 - A SIM card for the cellular module
 - A Lithium-Polymer battery for back-up power
 - A Central Processing Unit (CPU)
 - An embedded application in Non-Volatile RAM
 - RAM
 - EEPROM
- **Service Ecosystem:**
 - Cellular Data connectivity
 - Secure Private APN
 - Service Access Point
 - Cellular Modem OTA management service
 - SIM Card OTA management service

After marking down the information relevant to each technology, the team reviews the Model section of each Guideline document and identifies the appropriate technological model. This Endpoint is a Complex Endpoint. The Service and Network model is a standard mobile-enabled IoT service.

B.2 Review the Security Model

With the technical model outlined, the organization should now be ready to move forward with the review of the security model. In the security model, the team will evaluate how an adversary is likely to attack the solution.




Figure 11– Connected Car Attack Surfaces

In our example solution, there are only two threat surfaces that are relevant to an attack:

- The cellular network
- A localised attack on the vehicle

Since there is no local network connection, only a mobile network connection, an Attacker would have to either compromise the cellular network connection, enter the communications channel from the private APN or enter via the Service Access Point, cellular modem OTA management server or SIM card OTA management server.

Physical attacks are the only other way to compromise the device of which there are multiple entry points as shown in the above diagram, so in the case of this IoT service the Endpoint should be heavily focused on.

B.3 Review and Assign Security Tasks

With the security model evaluated it is now simple to assign Security Tasks. Each team should assign a specific person to each Component of the solution that needs evaluation. This should be evaluated not only from the high level perspective (Endpoint, Network, and Service) but from the subcomponent perspective. This means that the CPU should be assigned a worker, the operating system, the network service, and so forth.

Once each Component is assigned to an owner, the process can begin. This means, at this stage, the team understands:

- How the technology is composed
- What technologies affect security
- What engineering stakeholders own the given technology

B.4 Review Recommendations

In the recommendation review phase, each member of the team should read and understand *as many* of the recommendations as possible. This is by design. Instead of focusing solely on the recommendations affixed to a specific Component, engineers should take the time to understand as many recommendations as they are able, even if only at a high level, to gain a better view of how their Component affects the overall security of the product or service. This way, the group can engage in valuable discussion on what remediation or mitigation strategies will have the most balance from a cost effectiveness, longevity, and management perspective.

Once the recommendations are reviewed, the *Component owners* can determine whether a recommendation has already been applied, or mark a recommendation *pending*. This will allow the group to have a discussion regarding the applicability of a recommendation prior to its deployment. This is a better strategy to follow, as some recommendations may have side effects that impact the fulfilment of other recommendations, or existing controls.

In this example, the team would have determined that:

- An application trust base should be used
- An Organizational Root of Trust should be defined
- Device personalization should be implemented
- Tamper resistant casing should be implemented
- Endpoint password management should be enforced
- Endpoint communications security should be enforced
- Cryptographically signed images should be implemented
- Privacy management should be implemented
- Device power alerts should be integrated

B.5 Review Component Risk

Next, the Components section should be evaluated to identify the various risks involved in implementing or integrating a particular Component into the product or service. This section can generally be reviewed only by the Component owner to minimize work. Though, it is always beneficial to read as much as possible.

After reviewing Recommendations and the Component risk section, the following security gaps were identified:

- Secrets were stored unprotected in EEPROM
- Secrets were not processed in internal RAM
- User interface must protect passwords
- User privacy should be outlined for the user

B.6 Implementation and Review

Now the team can adjust the solution to adhere to the security requirements they agreed upon. The team re-implements components, where necessary, and adds security controls, where necessary.

In this particular instance, the team has identified that they are working with a GSMA member that is capable of provisioning an SIM card that contains application-capable trust anchor technology. They will resolve their need for a trust anchor by using the existing SIM card. This also resolves personalization, as each SIM can be personalized in the field using standard GSMA technology.

SIM technology can also help provision communication security keys over the air, resolving the need to implement communications authentication and privacy.

The SIM company-specific zone can be programmed with a trusted root base that enables the business to authenticate peers using a certificate chain. This resolves Organizational Root of Trust and peer authentication requirements.

The product encasing is updated with an appropriate tamper-resistant package.

The EEPROM is encoded with data that is encrypted with security keys stored in the SIM trust anchor.

The bootloader is altered to use the trust anchor for the authentication of the application image.

The Endpoint is reprogrammed to support secure password input from the user by blocking out password characters as they are typed.

A privacy management GUI is added so the user can view and control what information is being gathered by the business.

Secrets are processed only in internal memory of the same chip.

Once these implementations are defined, the team re-evaluates all security Recommendations and Risks, and reviews the Security Model to identify whether the changes have resolved their concerns.

B.7 Ongoing Lifecycle

Now that the team has achieved an approved configuration, they are ready to deploy their technology. However, security does not stop here. The team negotiates a methodology for monitoring Endpoints for security anomalies, and a methodology for identifying whether the technology they are using contains newly discovered security gaps.

The team will plan how each incident or gap is identified, remediated, and recovered from. This will ensure that, over time, the evolving technological and security landscape will not take the organization by surprise.

Annex C Document Management

C.1 Document History

Version	Date	Brief Description of Change	Approval Authority	Editor / Company
1.0	08-Feb-2016	New PRD CLP.11	PSMC	Ian Smith GSMA & Don A. Bailey Lab Mouse Security

C.2 Other Information

Type	Description
Document Owner	Connected Living Programme
Contact	Ian Smith - GSMA

It is our intention to provide a quality product for your use. If you find any errors or omissions, please contact us with your comments. You may notify us at prd@gsma.com

Your comments or suggestions & questions are always welcome.