

Internet Engineering Task Force (IETF)
Request for Comments: 7744
Category: Informational
ISSN: 2070-1721

L. Seitz, Ed.
SICS Swedish ICT AB
S. Gerdes, Ed.
Universitaet Bremen TZI
G. Selander
Ericsson
M. Mani
Itron
S. Kumar
Philips Research
January 2016

Use Cases for Authentication and Authorization in Constrained Environments

Abstract

Constrained devices are nodes with limited processing power, storage space, and transmission capacities. In many cases, these devices do not provide user interfaces, and they are often intended to interact without human intervention.

This document includes a collection of representative use cases for authentication and authorization in constrained environments. These use cases aim at identifying authorization problems that arise during the life cycle of a constrained device and are intended to provide a guideline for developing a comprehensive authentication and authorization solution for this class of scenarios.

Where specific details are relevant, it is assumed that the devices use the Constrained Application Protocol (CoAP) as a communication protocol. However, most conclusions apply generally.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7744>.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Terminology	4
2. Use Cases	5
2.1. Container Monitoring	5
2.1.1. Bananas for Munich	6
2.1.2. Authorization Problems Summary	7
2.2. Home Automation	8
2.2.1. Controlling the Smart Home Infrastructure	8
2.2.2. Seamless Authorization	8
2.2.3. Remotely Letting in a Visitor	9
2.2.4. Selling the House	9
2.2.5. Authorization Problems Summary	9
2.3. Personal Health Monitoring	10
2.3.1. John and the Heart Rate Monitor	11
2.3.2. Authorization Problems Summary	12
2.4. Building Automation	13
2.4.1. Device Life Cycle	13
2.4.1.1. Installation and Commissioning	13
2.4.1.2. Operational	14
2.4.1.3. Maintenance	15
2.4.1.4. Recommissioning	16
2.4.1.5. Decommissioning	16
2.4.2. Public Safety	17
2.4.2.1. A Fire Breaks Out	17
2.4.3. Authorization Problems Summary	18
2.5. Smart Metering	19
2.5.1. Drive-By Metering	19
2.5.2. Meshed Topology	20
2.5.3. Advanced Metering Infrastructure	20
2.5.4. Authorization Problems Summary	21
2.6. Sports and Entertainment	22
2.6.1. Dynamically Connecting Smart Sports Equipment	22
2.6.2. Authorization Problems Summary	23
2.7. Industrial Control Systems	23
2.7.1. Oil Platform Control	23
2.7.2. Authorization Problems Summary	24
3. Security Considerations	24
3.1. Attacks	25
3.2. Configuration of Access Permissions	26
3.3. Authorization Considerations	26
3.4. Proxies	28
4. Privacy Considerations	28
5. Informative References	28
Acknowledgments	29
Authors' Addresses	30

1. Introduction

Constrained devices [RFC7228] are nodes with limited processing power, storage space, and transmission capacities. These devices are often battery-powered and in many cases do not provide user interfaces.

Constrained devices benefit from being interconnected using Internet protocols. However, deploying common security protocols can sometimes be difficult because of device or network limitations. Regardless, adequate security mechanisms are required to protect these constrained devices, which are expected to be integrated in all aspects of everyday life, from attackers wishing to gain control over the device's data or functions.

This document comprises a collection of representative use cases for the application of authentication and authorization in constrained environments. These use cases aim at identifying authorization problems that arise during the life cycle of a constrained device. Note that this document does not aim at collecting all possible use cases.

We assume that the communication between the devices is based on the Representational State Transfer (REST) architectural style, i.e., a device acts as a server that offers resources such as sensor data and actuators. The resources can be accessed by clients, sometimes without human intervention (M2M). In some situations, the communication will happen through intermediaries (e.g., gateways, proxies).

Where specific detail is necessary, it is assumed that the devices communicate using CoAP [RFC7252], although most conclusions are generic.

1.1. Terminology

Readers are required to be familiar with the terms defined in [RFC7228].

2. Use Cases

This section includes the use cases; each use case first presents a general description of the application environment, then one or more specific use cases, and finally a summary of the authorization-related problems to be solved. The document aims at listing the relevant authorization problems and not to provide an exhaustive list. It might not be possible to address all of the listed problems with a single solution; there might be conflicting goals within or among some requirements.

There are various reasons for assigning a function (client or server) to a device. The function may even change over time; e.g., the device that initiates a conversation is temporarily assigned the role of client, but could act as a server in another context. The definition of the function of a device in a certain use case is not in scope of this document. Readers should be aware that there might be reasons for each setting and that endpoints might even have different functions at different times.

2.1. Container Monitoring

The ability of sensors to communicate environmental data wirelessly opens up new application areas. Sensor systems make it possible to continuously track and transmit characteristics such as temperature, humidity, and gas content while goods are transported and stored.

Sensors in this scenario have to be associated with the appropriate pallet of the respective container. Sensors, as well as the goods, belong to specific customers.

While in transit, goods often pass stops where they are transloaded to other means of transportation, e.g., from ship transport to road transport.

Perishable goods need to be stored at a constant temperature and with proper ventilation. Real-time information on the state of the goods is needed by both the transporter and the vendor. Transporters want to prioritize goods that will expire soon. Vendors want to react when goods are spoiled to continue to fulfill delivery obligations.

The Intelligent Container <<http://www.intelligentcontainer.com>> is an example project that explores solutions to continuously monitor perishable goods.

2.1.1. Bananas for Munich

A fruit vendor grows bananas in Costa Rica for the German market. It instructs a transport company to deliver the goods via ship to Rotterdam where they are picked up by trucks and transported to a ripening facility. A Munich supermarket chain buys ripened bananas from the fruit vendor and transports them from the ripening facility to the individual markets with their own company's trucks.

The fruit vendor's quality management wants to assure the quality of their products; thus, it equips the banana boxes with sensors. The state of the goods is monitored consistently during shipment and ripening, and abnormal sensor values are recorded (U1.2). Additionally, the sensor values are used to control the climate within the cargo containers (U1.1, U1.5, U1.7). Therefore, the sensors need to communicate with the climate-control system. Since an incorrect sensor value leads to a wrong temperature, and thus to spoiled goods, the integrity of the sensor data must be assured (U1.2, U1.3). The banana boxes within a container will, in most cases, belong to the same owner. Adjacent containers might contain goods and sensors of different owners (U1.1).

The personnel that transloads the goods must be able to locate the goods meant for a specific customer (U1.1, U1.6, U1.7). However, the fruit vendor does not want to disclose sensor information pertaining to the condition of the goods to other companies and therefore wants to assure the confidentiality of this data (U1.4). Thus, the transloading personnel is only allowed to access logistic information (U1.1). Moreover, the transloading personnel is only allowed to access the data for the time of the transloading (U1.8).

Due to the high water content of the fruits, the propagation of radio waves is hindered, thus often inhibiting direct communication between nodes [Jedermann14]. Instead, messages are forwarded over multiple hops (U1.9). The sensors in the banana boxes cannot always reach the Internet during the journey (U1.10). Sensors may need to use relay stations owned by the transport company to connect to endpoints on the Internet.

In the ripening facility bananas are stored until they are ready to be sold. The banana box sensors are used to control the ventilation system and to monitor the degree of ripeness of the bananas. Ripe bananas need to be identified and sold before they spoil (U1.2, U1.8).

The supermarket chain gains ownership of the banana boxes when the bananas have ripened and are ready to leave the ripening facility.

2.1.2. Authorization Problems Summary

- U1.1: Fruit vendors and container owners want to grant different authorizations for their resources and/or endpoints to different parties.
- U1.2: The fruit vendor requires the integrity and authenticity of the sensor data that pertains to the state of the goods for climate control and to ensure the quality of the monitored recordings.
- U1.3: The container owner requires the integrity and authenticity of the sensor data that is used for climate control.
- U1.4: The fruit vendor requires the confidentiality of the sensor data that pertains the state of the goods and the confidentiality of location data, e.g., to protect them from targeted attacks from competitors.
- U1.5: The fruit vendor may need different protection for several different types of data on the same endpoint, e.g., sensor data and the data used for logistics.
- U1.6: The fruit vendor and the transloading personnel require the authenticity and integrity of the data that is used to locate the goods, in order to ensure that the goods are correctly treated and delivered.
- U1.7: The container owner and the fruit vendor may not be present at the time of access and cannot manually intervene in the authorization process.
- U1.8: The fruit vendor, container owner, and transloading company want to grant temporary access permissions to a party, in order to avoid giving permanent access to parties that are no longer involved in processing the bananas.
- U1.9: The fruit vendor, container owner, and transloading company want their security objectives to be achieved, even if the messages between the endpoints need to be forwarded over multiple hops.
- U1.10: The constrained devices might not always be able to reach the Internet but still need to enact the authorization policies of their principals.
- U1.11: Fruit vendors and container owners want to be able to revoke authorization on a malfunctioning sensor.

2.2. Home Automation

One application of the Internet of Things is home automation systems. Such a system can connect household devices that control, for example, heating, ventilation, lighting, home entertainment, and home security to the Internet making them remotely accessible and manageable.

Such a system needs to accommodate a number of regular users (inhabitants, close friends, cleaning personnel) as well as a heterogeneous group of dynamically varying users (visitors, repairmen, delivery men).

As the users are not typically trained in security (or even computer use), the configuration must use secure default settings, and the interface must be well adapted to novice users.

2.2.1. Controlling the Smart Home Infrastructure

Alice and Bob own a flat that is equipped with home automation devices such as HVAC and shutter control, and they have a motion sensor in the corridor that controls the light bulbs there (U2.5).

Alice and Bob can control the shutters and the temperature in each room using either wall-mounted touch panels or an Internet connected device (e.g., a smartphone). Since Alice and Bob both have full-time jobs, they want to be able to change settings remotely, e.g., turn up the heating on a cold day if they will be home earlier than expected (U2.5).

The couple does not want people in radio range of their devices, e.g., their neighbors, to be able to control them without authorization. Moreover, they don't want burglars to be able to deduce behavioral patterns from eavesdropping on the network (U2.8).

2.2.2. Seamless Authorization

Alice buys a new light bulb for the corridor and integrates it into the home network, i.e., makes resources known to other devices in the network. Alice makes sure that the new light bulb and her other devices in the network get to know the authorization policies for the new device. Bob is not at home, but Alice wants him to be able to control the new device with his devices (e.g., his smartphone) without the need for additional administration effort (U2.7). She provides the necessary configurations for that (U2.9, U2.10).

2.2.3. Remotely Letting in a Visitor

Alice and Bob have equipped their home with automated connected door-locks and an alarm system at the door and the windows. The couple can control this system remotely.

Alice and Bob have invited Alice's parents over for dinner, but are stuck in traffic and cannot arrive in time; whereas Alice's parents are using the subway and will arrive punctually. Alice calls her parents and offers to let them in remotely, so they can make themselves comfortable while waiting (U2.1, U2.6). Then, Alice sets temporary permissions that allow them to open the door and shut down the alarm (U2.2). She wants these permissions to be only valid for the evening since she does not like it if her parents are able to enter the house as they see fit (U2.3, U2.4).

When Alice's parents arrive at Alice and Bob's home, they use their smartphone to communicate with the door-lock and alarm system (U2.5, U2.9). The permissions Alice issued to her parents only allow limited access to the house (e.g., opening the door, turning on the lights). Certain other functions, such as checking the footage from the surveillance cameras, are not accessible to them (U2.3).

Alice and Bob also issue similarly restricted permissions to e.g., cleaners, repairmen, or their nanny (U2.3).

2.2.4. Selling the House

Alice and Bob have to move because Alice is starting a new job. They therefore decide to sell the house and transfer control of all automated services to the new owners (U2.11). Before doing so, they want to erase privacy-relevant data from the logs of the automated systems, while the new owner is interested to keep some historic data e.g., pertaining to the behavior of the heating system (U2.12). At the time of transfer of ownership of the house, the new owners also want to make sure that permissions issued by the previous owners to access the house or connected devices (in the case where device management may have separate permissions from house access) are no longer valid (U2.13).

2.2.5. Authorization Problems Summary

U2.1: A home owner (Alice and Bob in the example above) wants to spontaneously provision authorization means to visitors.

U2.2: A home owner wants to spontaneously change the home's access control policies.

- U2.3: A home owner wants to apply different access rights for different users (including other inhabitants).
- U2.4: The home owners want to grant access permissions to someone during a specified time frame.
- U2.5: The smart home devices need to be able to securely communicate with different control devices (e.g., wall-mounted touch panels, smartphones, electronic key fobs, and device gateways).
- U2.6: The home owner wants to be able to configure authorization policies remotely.
- U2.7: Authorized users want to be able to obtain access with little effort.
- U2.8: The owners of the automated home want to prevent unauthorized entities from being able to deduce behavioral profiles from devices in the home network.
- U2.9: Usability is particularly important in this scenario since the necessary authorization related tasks in the life cycle of the device (commissioning, operation, maintenance, and decommissioning) likely need to be performed by the home owners who, in most cases, have little knowledge of security.
- U2.10: Home owners want their devices to seamlessly (and in some cases even unnoticeably) fulfill their purpose. Therefore, the authorization administration effort needs to be kept at a minimum.
- U2.11: Home owners want to be able to transfer ownership of their automated systems when they sell the house.
- U2.12: Home owners want to be able to sanitize the logs of the automated systems when transferring ownership without deleting important operational data.
- U2.13: When a transfer of ownership occurs, the new owner wants to make sure that access rights created by the previous owner are no longer valid.

2.3. Personal Health Monitoring

Personal health monitoring devices, i.e., eHealth devices, are typically battery-driven and located physically on or in the user to monitor some bodily function, such as temperature, blood pressure, or

pulse rate. These devices typically connect to the Internet through an intermediary base station, using wireless technologies and through this connection they report the monitored data to some entity, which may either be the user or a medical caregiver.

Medical data has always been considered very sensitive, and therefore requires good protection against unauthorized disclosure. A frequent, conflicting requirement is the capability for medical personnel to gain emergency access, even if no specific access rights exist. As a result, the importance of secure audit logs increases in such scenarios.

Since the users are not typically trained in security (or even computer use), the configuration must use secure default settings, and the interface must be well adapted to novice users. Parts of the system must operate with minimal maintenance. Especially frequent changes of battery are unacceptable.

There is a plethora of wearable health monitoring technology and the need for open industry standards to ensure interoperability between products has lead to initiatives such as Continua Alliance <<http://continuaalliance.org>> and Personal Connected Health Alliance <<http://www.pchalliance.org>>.

2.3.1. John and the Heart Rate Monitor

John has a heart condition that can result in sudden cardiac arrests. He therefore uses a device called "HeartGuard" that monitors his heart rate and his location (U3.7). In the event of a cardiac arrest, it automatically sends an alarm to an emergency service, transmitting John's current location (U3.1). Either the device has long-range connectivity itself (e.g., via GSM) or it uses some intermediary, nearby device (e.g., John's smartphone) to transmit such an alarm. To ensure John's safety, the device is expected to be in constant operation (U3.3, U3.6).

The device includes an authentication mechanism to prevent other persons who get physical access to it from acting as the owner and altering the access control and security settings (U3.8).

John can configure a list of people that get notified in an emergency, for example his daughter Jill. Furthermore, the device stores data on John's heart rate, which can later be accessed by a physician to assess the condition of John's heart (U3.2).

However, John is a privacy-conscious person and is worried that Jill might use HeartGuard to monitor his location even when there is no emergency. Furthermore, he doesn't want his health insurance to get

access to the HeartGuard data, or even to the fact that he is wearing a HeartGuard, since they might refuse to renew his insurance if they decided he was too great of a risk for them (U3.8).

Finally, John, while being comfortable with modern technology and able to operate it reasonably well, is not trained in computer security. Therefore, he needs an interface for the configuration of the HeartGuard security that is easy to understand and use (U3.5). If John does not understand the meaning of a setting, he tends to leave it alone, assuming that the manufacturer has initialized the device to secure settings (U3.4).

Note: Monitoring of some state parameter (e.g., an alarm button) and the position of a person also fits well into a nursing service context. This is particularly useful for people suffering from dementia, where the relatives or caregivers need to be notified of the whereabouts of the person under certain conditions. In that case, it is not the patient that decides about access.

2.3.2. Authorization Problems Summary

- U3.1: The wearer of an eHealth device (John in the example above) wants to preconfigure special access rights in the context of an emergency.
- U3.2: The wearer of an eHealth device wants to selectively allow different persons or groups access to medical data.
- U3.3: Battery changes are very inconvenient and sometimes impractical, so battery life impacts on the authorization mechanisms need to be minimized.
- U3.4: Devices are often used with default access control settings that might threaten the security objectives of the device's users.
- U3.5: Wearers of eHealth devices are often not trained in computer use, especially computer security.
- U3.6: Security mechanisms themselves could provide opportunities for denial-of-service (DoS) attacks, especially on the constrained devices.
- U3.7: The device provides a service that can be fatal for the wearer if it fails. Accordingly, the wearer wants the device to have a high degree of resistance against attacks that may cause the device to fail to operate partially or completely.

U3.8: The wearer of an eHealth device requires the integrity and confidentiality of the data measured by the device.

2.4. Building Automation

Buildings for commercial use such as shopping malls or office buildings nowadays are equipped increasingly with semi-automatic components to enhance the overall living quality and to save energy where possible. This includes for example heating, ventilation and air condition (HVAC) as well as illumination and security systems such as fire alarms. These components are being increasingly managed centrally in a Building and Lighting Management System (BLMS) by a facility manager.

Different areas of these buildings are often exclusively leased to different companies. However, they also share some of the common areas of the building. Accordingly, a company must be able to control the lighting and HVAC system of its own part of the building and must not have access to control rooms that belong to other companies.

Some parts of the building automation system such as entrance illumination and fire-alarm systems are controlled either by all parties together or by a facility-management company.

2.4.1. Device Life Cycle

2.4.1.1. Installation and Commissioning

Installation of the building automation components often start even before the construction work is completed. Lighting is one of the first components to be installed in new buildings. A lighting plan created by a lighting designer provides the necessary information related to the kind of lighting devices (luminaires, sensors, and switches) to be installed along with their expected behavior. The physical installation of the correct lighting devices at the right locations are done by electricians based on the lighting plan. They ensure that the electrical wiring is performed according to local regulations and lighting devices, which may be from multiple manufacturers, are connected to the electrical power supply properly. After the installation, lighting can be used in a default out-of-box mode, e.g., at full brightness when powered on. After this step (or in parallel in a different section of the building), a lighting commissioner adds the devices to the building domain (U4.1) and performs the proper configuration of the lights as prescribed in the lighting plan. This involves, for example, grouping to ensure that light points react together, more or less synchronously (U4.8) and defining lighting scenes for particular areas of the building. The

commissioning is often done in phases, either by one or more commissioners, on different floors. The building lighting network at this stage may be in different network islands with no connectivity between them due to lack of the IT infrastructure.

After this, other building components, like HVAC and security systems, are similarly installed by electricians and later commissioned by their respective domain professionals. Similar configurations related to grouping (U4.8) are required to ensure, e.g., HVAC equipment is controlled by the closest temperature sensor.

For the building IT systems, the Ethernet wiring is initially laid out in the building according to the IT plan. The IT network is often commissioned after the construction is completed to avoid any damage to sensitive networking and computing equipment. The commissioning is performed by an IT engineer with additional switches (wired and/or wireless), IP routers, and computing devices. Direct Internet connectivity for all installed/commissioned devices in the building is only available at this point. The BLMS that monitors and controls the various building automation components is only connected to the field devices at this stage. The different network islands (for lighting and HVAC) are also joined together without any further involvement of domain specialists, such as lighting or HVAC commissioners.

2.4.1.2. Operational

The building automation system is now finally ready, and the operational access is transferred to the facility management company of the building (U4.2). The facility manager is responsible for monitoring and ensuring that the building automation system meets the needs of the building occupants. If changes are needed, the facility-management company hires an external installation and commissioning company to perform the changes.

Different parts of the building are rented out to different companies for office space. The tenants are provided access to use the automated HVAC, lighting, and physical access control systems deployed. The safety of the occupants is also managed using automated systems, such as a fire-alarm system, which is triggered by several smoke detectors that are spread out across the building.

Company A's staff moves into the newly furnished office space. Most lighting is controlled by presence sensors that control the lighting of a specific group of lights based on the authorization rules in the BLMS. Additionally, employees are allowed to manually override the lighting brightness and color in their offices by using the switches

or handheld controllers. Such changes are allowed only if the authorization rules exist in the BLMS. For example, lighting in the corridors may not be manually adjustable.

At the end of the day, lighting is dimmed or switched off if no occupancy is detected, even if manually overridden during the day.

On a later date, Company B also moves into the same building, and shares some of the common spaces and associated building automation components with Company A (U4.2, U4.9).

2.4.1.3. Maintenance

Company A's staff is annoyed that the lighting switches off too often in their rooms if they work silently in front of their computers. Company A notifies the facility manager of the building to increase the delay before lights switch off. The facility manager can either configure the new values directly in the BLMS or, if additional changes are needed on the field devices, hire commissioning Company C to perform the needed changes (U4.4).

Company C gets the necessary authorization from the facility-management company to interact with the BLMS. The commissioner's tool gets the necessary authorization from the BLMS to send a configuration change to all lighting devices in Company A's offices to increase the delay before they switch off.

At some point, the facility-management company wants to update the firmware of lighting devices in order to eliminate software bugs. Before accepting the new firmware, each device checks the authorization of the facility-management company to perform this update (U4.13).

A network-diagnostic tool of the BLMS detects that a luminaire in one of Company A's offices is no longer connected to the network. The BLMS alerts the facility manager to replace the luminaire. The facility manager replaces the old broken luminaire and informs the BLMS of the identity (e.g., the Media Access Control (MAC) address) of the newly added device. Then, the BLMS authorizes the new device in the system and seamlessly transfers all the permissions of the previous broken device to the replacement device (U4.12).

2.4.1.4. Recommissioning

A vacant area of the building has recently been leased to Company A. Before moving into its new office, Company A wishes to replace the lighting with more energy efficient and better light quality luminaries. They hire an installation and commissioning Company C to redo the illumination. Company C is instructed to integrate the new lighting devices, which may be from multiple manufacturers, into the existing lighting infrastructure of the building, which includes presence sensors, switches, controllers, etc. (U4.1).

Company C gets the necessary authorization from the facility-management company to interact with the existing BLMS (U4.4). To prevent disturbance to other occupants of the building, Company C is provided authorization to perform the commissioning only during non-office hours and only to modify configuration on devices belonging to the domain of Company A's space (U4.5). Before removing existing devices, all security and configuration material that belongs to the domain is deleted and the devices are set back to factory state (U4.3). This ensures that these devices may be reused at other installations or in other parts of the same building without affecting future operations. After installation (wiring) of the new lighting devices, the commissioner adds the devices into Company A's lighting domain.

Once the devices are in the correct domain, the commissioner authorizes the interaction rules between the new lighting devices and existing devices, like presence sensors (U4.7). For this, the commissioner creates the authorization rules on the BLMS that define which lights form a group and which sensors/switches/controllers are allowed to control which groups (U4.8). These authorization rules may be context based, like time of the day (office or non-office hours) or location of the handheld lighting controller, etc. (U4.5).

2.4.1.5. Decommissioning

Company A has noticed that the handheld controllers are often misplaced and hard to find when needed. So most of the time, staff use the existing wall switches for manual control. Company A decides it would be better to completely remove handheld controllers and asks Company C to decommission them from the lighting system (U4.4).

Company C again gets the necessary authorization from the facility-management company to interact with the BLMS. The commissioner now deletes any rules that allowed handheld controllers authorization to control the lighting (U4.3, U4.6). Additionally, the commissioner instructs the BLMS to push these new rules to prevent cached rules at

the end devices from being used. Any cryptographic key material belonging to the site in the handheld controllers is also removed, and they are set to the factory state (U4.3).

2.4.2. Public Safety

As part of the building safety code, the fire department requires that the building have sensors that sense the level of smoke, heat, etc., when a fire breaks out. These sensors report metrics that are then used by a back-end server to map safe areas and unsafe areas within a building and possibly the structural integrity of the building before firefighters may enter it.

Sensors may also be used to track where human/animal activity is within the building. This will allow people stuck in the building to be guided to safer areas and allow the suggestion of possible actions that they may take (e.g., using a client application on their phones or giving loudspeaker directions) in order to bring them to safety. In certain cases, other organizations such as the police, ambulance, and federal organizations are also involved and therefore the coordination of tasks between the various entities have to be carried out using efficient messaging and authorization mechanisms.

2.4.2.1. A Fire Breaks Out

James, who works for Company A, turns on the air conditioning in his office on a really hot day. Lucy, who works for Company B, wants to make tea using an electric kettle. After she turns it on, she goes outside to talk to a colleague until the water is boiling. Unfortunately, her kettle has a malfunction that causes overheating and results in a smoldering fire of the kettle's plastic case.

Due to the smoke coming from the kettle, the fire alarm is triggered. Alarm sirens throughout the building are switched on simultaneously (using a group communication scheme) to alert the staff of both companies (U4.8). Additionally, the ventilation system of the whole building is closed off to prevent the smoke from spreading and to withdraw oxygen from the fire. The smoke cannot get into James' office, even though he turned on his air conditioning, because the fire alarm overrides the manual setting by sending commands (using group communication) to switch off all the air conditioning (U4.10).

The fire department is notified of the fire automatically and arrives within a short time. They automatically get access to all parts of the building according to an emergency authorization policy (U4.4, U4.5). After inspecting the damage and extinguishing the smoldering fire, a firefighter resets the fire alarm because only the fire department is authorized to do that (U4.4, U4.11).

2.4.3. Authorization Problems Summary

- U4.1: During commissioning, the building owner or the companies add new devices to their administrative domain. Access control should then apply to these devices seamlessly.
- U4.2: During a handover, the building owner or the companies integrate devices that formerly belonged to a different administrative domain to their own administrative domain. Access control of the old domain should then cease to apply, with access control of the new domain taking over.
- U4.3: During decommissioning, the building owner or the companies remove devices from their administrative domain. Access control should cease to apply to these devices and relevant credentials need to be erased from the devices.
- U4.4: The building owner and the companies want to be able to delegate specific access rights for their devices to others.
- U4.5: The building owner and the companies want to be able to define context-based authorization rules.
- U4.6: The building owner and the companies want to be able to revoke granted permissions and delegations.
- U4.7: The building owner and the companies want to allow authorized entities to send data to their endpoints (default deny).
- U4.8: The building owner and the companies want to be able to authorize a device to control several devices at the same time using a group communication scheme.
- U4.9: The companies want to be able to interconnect their own subsystems with those from a different operational domain while keeping the control over the authorizations (e.g., granting and revoking permissions) for their endpoints and devices.
- U4.10: The authorization mechanisms must be able to cope with extremely time-sensitive operations that have to be carried out quickly.
- U4.11: The building owner and the public safety authorities want to be able to perform data origin authentication on messages sent and received by some of the systems in the building.

- U4.12: The building owner should be allowed to replace an existing device with a new device providing the same functionality within their administrative domain. Access control from the replaced device should then apply to these new devices seamlessly.
- U4.13: When software on a device is updated, this update needs to be authenticated and authorized.

2.5. Smart Metering

Automated measuring of customer consumption is an established technology for electricity, water, and gas providers. Increasingly, these systems also feature networking capability to allow for remote management. Such systems are in use for commercial, industrial, and residential customers and require a certain level of security, in order to avoid economic loss to the providers, vulnerability of the distribution system, as well as disruption of services for the customers.

The smart metering equipment for gas and water solutions is battery-driven and communication should be used sparingly due to battery consumption. Therefore, these types of meters sleep most of the time, and only wake up every minute/hour to check for incoming instructions. Furthermore, they wake up a few times a day (based on their configuration) to upload their measured metering data.

Different networking topologies exist for smart metering solutions. Based on environment, regulatory rules, and expected cost, one or a mixture of these topologies may be deployed to collect the metering information. Drive-by metering is one of the most current solutions deployed for collection of gas and water meters.

Various stakeholders have a claim on the metering data. Utility companies need the data for accounting, the metering equipment may be operated by a third-party service operator who needs to maintain it, and the equipment is installed in the premises of the consumers, measuring their consumption, which entails privacy questions.

2.5.1. Drive-By Metering

A service operator offers smart metering infrastructures and related services to various utility companies. Among these is a water provider, who in turn supplies several residential complexes in a city. The smart meters are installed in the end customer's homes to measure water consumption and thus generate billing data for the utility company. They can also be used to shut off the water if the bills are not paid (U5.1, U5.3). The meters do this by sending and

receiving data to and from a base station (U5.2). Several base stations are installed around the city to collect the metering data. However, in the denser urban areas, the base stations would have to be installed very close to the meters. This would require a high number of base stations and expose this more expensive equipment to manipulation or sabotage. The service operator has therefore chosen another approach, which is to drive around with a mobile base station and let the meters connect to that in regular intervals in order to gather metering data (U5.4, U5.6, U5.8).

2.5.2. Meshed Topology

In another deployment, the water meters are installed in a building that already has power meters installed, the latter are mains powered, and are therefore not subject to the same power saving restrictions. The water meters can therefore use the power meters as proxies, in order to achieve better connectivity. This requires the security measures on the water meters to work through intermediaries (U5.9).

2.5.3. Advanced Metering Infrastructure

A utility company is updating its old utility distribution network with advanced meters and new communication systems, known as an Advanced Metering Infrastructure (AMI). AMI refers to a system that measures, collects, and analyzes usage, and interacts with metering devices such as electricity meters, gas meters, heat meters, and water meters, through various communication media either on request (on-demand) or on predefined schedules. Based on this technology, new services make it possible for consumers to control their utility consumption (U5.2, U5.7) and reduce costs by supporting new tariff models from utility companies, and more accurate and timely billing. However, the end consumers do not want unauthorized persons to gain access to this data. Furthermore, the fine-grained measurement of consumption data may induce privacy concerns, since it may allow others to create behavioral profiles (U5.5, U5.10).

The technical solution is based on levels of data aggregation between smart meters located at the consumer premises and the Meter Data Management (MDM) system located at the utility company (U5.9). For reasons of efficiency and cost, end-to-end connectivity is not always feasible, so metering data is stored and aggregated in various intermediate devices before being forwarded to the utility company, and in turn accessed by the MDM. The intermediate devices may be operated by a third-party service operator on behalf of the utility company (U5.7). One responsibility of the service operator is to make sure that meter readings are performed and delivered in a regular, timely manner. An example of a Service Level Agreement

between the service operator and the utility company is, for example, at least 95% of the meters have readings recorded during the last 72 hours.

2.5.4. Authorization Problems Summary

- U5.1: Devices are installed in hostile environments where they are physically accessible by attackers (including dishonest customers). The service operator and the utility company want to make sure that an attacker cannot use data from a captured device to attack other parts of their infrastructure.
- U5.2: The utility company wants to control which entities are allowed to send data to, and read data from, their endpoints.
- U5.3: The utility company wants to ensure the integrity of the data stored on their endpoints.
- U5.4: The utility company wants to protect such data transfers to and from their endpoints.
- U5.5: Consumers want to access their own usage information and also prevent unauthorized access by others.
- U5.6: The devices may have intermittent Internet connectivity but still need to enact the authorization policies of their principals.
- U5.7: Neither the service operator nor the utility company are always present at the time of access and cannot manually intervene in the authorization process.
- U5.8: When authorization policies are updated it is impossible, or at least very inefficient to contact all affected endpoints directly.
- U5.9: Authorization and authentication must work even if messages between endpoints are stored and forwarded over multiple nodes.
- U5.10: Consumers may not want the service operator, the utility company or others to have access to a fine-grained level of consumption data that allows the creation of behavioral profiles.

2.6. Sports and Entertainment

In the area of leisure-time activities, applications can benefit from the small size and weight of constrained devices. Sensors and actuators with various functions can be integrated into fitness equipment, games, and even clothes. Users can carry their devices around with them at all times.

Usability is especially important in this area since users will often want to spontaneously interconnect their devices with others. Therefore, the configuration of access permissions must be simple and fast and not require much effort at the time of access.

Continuously monitoring allows authorized users to create behavioral or movement profiles, that correspond to the devices' intended use, and unauthorized access to the collected data would allow an attacker to create the same profiles. Moreover, the aggregation of data can seriously increase the impact on the privacy of the users.

2.6.1. Dynamically Connecting Smart Sports Equipment

Jody is an enthusiastic runner. To keep track of her training progress, she has smart running shoes that measure the pressure at various points beneath her feet to count her steps, detect irregularities in her stride, and help her to improve her posture and running style. On a sunny afternoon, she goes to the Finnbahn track near her home to work out. She meets her friend Lynn, who shows her the smart fitness watch she bought a few days ago. The watch can measure the wearer's pulse, show speed and distance, and keep track of the configured training program. The girls realize that the watch can be connected with Jody's shoes and can display the information the shoes provide.

Jody asks Lynn to let her try the watch and lend it to her for the afternoon. Lynn agrees, but she doesn't want Jody to access her training plan (U6.4). She configures the access policies for the watch so that Jody's shoes are allowed to access the display and measuring features but cannot read or add training data (U6.1, U6.2). Jody's shoes connect to Lynn's watch at the press of a button, because Jody already configured access rights for devices that belong to Lynn a while ago (U6.3). Jody wants the device to report the data back to her fitness account while she borrows it, so she allows it to access her account temporarily.

After an hour, Jody gives the watch back and both girls terminate the connection between their devices.

2.6.2. Authorization Problems Summary

- U6.1: Sports equipment owners want to be able to grant access rights dynamically when needed.
- U6.2: Sports equipment owners want the configuration of access rights to work with very little effort.
- U6.3: Sports equipment owners want to be able to preconfigure access policies that grant certain access permissions to endpoints with certain attributes (e.g., endpoints of a certain user) without additional configuration effort at the time of access.
- U6.4: Sports equipment owners want to protect the confidentiality of their data for privacy reasons.

2.7. Industrial Control Systems

Industrial control systems (ICS) and especially supervisory control and data acquisition systems (SCADA) use a multitude of sensors and actuators in order to monitor and control industrial processes in the physical world. Example processes include manufacturing, power generation, and refining of raw materials.

Since the advent of the Stuxnet worm, it has become obvious to the general public how vulnerable these kind of systems are, especially when connected to the Internet [Karnouskos11]. The severity of these vulnerabilities are exacerbated by the fact that many ICS are used to control critical public infrastructure, such as nuclear power, water treatment, or traffic control. Nevertheless, the economical advantages of connecting such systems to the Internet can be significant if appropriate security measures are put in place (U7.5).

2.7.1. Oil Platform Control

An oil platform uses an industrial control system to monitor data and control equipment. The purpose of this system is to gather and process data from a large number of sensors and control actuators such as valves and switches to steer the oil extraction process on the platform. Raw data, alarms, reports, and other information are also available to the operators, who can intervene with manual commands. Many of the sensors are connected to the controlling units by direct wire, but the operator is slowly replacing these units by wireless ones, since this makes maintenance easier (U7.4).

Some of the controlling units are connected to the Internet, to allow for remote administration, since it is expensive and inconvenient to fly in a technician to the platform (U7.3).

The main interest of the operator is to ensure the integrity of control messages and sensor readings (U7.1). Access in some cases needs to be restricted, e.g., the operator wants wireless actuators only to accept commands by authorized control units (U7.2).

The owner of the platform also wants to collect auditing information for liability reasons (U7.1).

Different levels of access apply e.g., for regular operators vs. maintenance technician vs. auditors of the platform (U7.6).

2.7.2. Authorization Problems Summary

- U7.1: The operator of the platform wants to ensure the integrity and confidentiality of sensor and actuator data.
- U7.2: The operator wants to ensure that data coming from sensors and commands sent to actuators are authentic.
- U7.3: Some devices do not have direct Internet connection, but they still need to implement current authorization policies.
- U7.4: Devices need to authenticate the controlling units, especially those using a wireless connection.
- U7.5: The execution of unauthorized commands or the failure to execute an authorized command in an ICS can lead to significant financial damage and threaten the availability of critical infrastructure services. Accordingly, the operator wants authentication and authorization mechanisms that provide a very high level of security.
- U7.6: Different users should have different levels of access to the control system (e.g., operator vs. auditor).

3. Security Considerations

As the use cases listed in this document demonstrate, constrained devices are used in various environments. These devices are small and inexpensive and this makes it easy to integrate them into many aspects of everyday life. With access to vast amounts of valuable data and possible control of important functions, these devices need to be protected from unauthorized access. Protecting seemingly innocuous data and functions will lessen the possible effects of aggregation; attackers collecting data or functions from several sources can gain insights or a level of control not immediately obvious from each of these sources on its own.

Not only the data on the constrained devices themselves is threatened, the devices might also be abused as an intrusion point to infiltrate a network. Once an attacker gains control over the device, it can be used to attack other devices as well. Due to their limited capabilities, constrained devices appear as the weakest link in the network; hence, they pose an attractive target for attackers.

This section summarizes the security problems highlighted by the use cases above and provides guidelines for the design of protocols for authentication and authorization in constrained RESTful environments.

3.1. Attacks

This document lists security problems that users of constrained devices want to solve. Further analysis of attack scenarios is not in scope of the document. However, there are attacks that must be considered by solution developers.

Because of the expected large number of devices and their ubiquity, constrained devices increase the danger from Pervasive Monitoring [RFC7258] attacks. Solution Designers should consider this in the design of their security solution and provide for protection against this type of attack. In particular, messages containing sensitive data that are sent over unprotected channels should be encrypted if possible.

Attacks aimed at altering data in transit (e.g., to perpetrate fraud) are a problem that is addressed in many web security protocols such as TLS or IPsec. Developers need to consider these types of attacks, and make sure that the protection measures they implement are adapted to the constrained environment.

As some of the use cases indicate, constrained devices may be installed in hostile environments where they are physically accessible (see Section 2.5). Protection from physical attacks is not in the scope of this document, but it should be kept in mind by developers of authorization solutions.

Denial-of-service (DoS) attacks threaten the availability of services a device provides and constrained devices are especially vulnerable to these types of attacks because of their limitations. Attackers can illicit a temporary or, if the battery is drained, permanent failure in a service simply by repeatedly flooding the device with connection attempts; for some services (see Section 2.3), availability is especially important. Solution designers must be particularly careful to consider the following limitations in every part of the authorization solution:

- o Battery usage
- o Number of required message exchanges
- o Size of data that is transmitted (e.g., authentication and access control data)
- o Size of code required to run the protocols
- o Size of RAM memory and stack required to run the protocols
- o Resources blocked by partially completed exchanges (e.g., while one party is waiting for a transaction time to run out)

Solution developers also need to consider whether the session should be protected from information disclosure and tampering.

3.2. Configuration of Access Permissions

- o The access control policies need to be enforced (all use cases):
The information that is needed to implement the access control policies needs to be provided to the device that enforces the authorization and applied to every incoming request.
- o A single resource might have different access rights for different requesting entities (all use cases).

Rationale: In some cases, different types of users need different access rights, as opposed to a binary approach where the same access permissions are granted to all authenticated users.

- o A device might host several resources where each resource has its own access control policy (all use cases).
- o The device that makes the policy decisions should be able to evaluate context-based permissions such as location or time of access (see Sections 2.2, 2.3, and 2.4). Access may depend on local conditions, e.g., access to health data in an emergency. The device that makes the policy decisions should be able to take such conditions into account.

3.3. Authorization Considerations

- o Devices need to be enabled to enforce authorization policies without human intervention at the time of the access request (see Sections 2.1, 2.2, 2.4, and 2.5).

- o Authorization solutions need to consider that constrained devices might not have Internet access at the time of the access request (see Sections 2.1, 2.3, 2.5, and 2.6).
- o It should be possible to update access control policies without manually re-provisioning individual devices (see Sections 2.2, 2.3, 2.5, and 2.6).

Rationale: Peers can change rapidly which makes manual re-provisioning unreasonably expensive.

- o Authorization policies may be defined to apply to a large number of devices that might only have intermittent connectivity. Distributing policy updates to every device for every update might not be a feasible solution (see Section 2.5).
- o It must be possible to dynamically revoke authorizations (see Section 2.4 for example).
- o The authentication and access control protocol can put undue burden on the constrained system resources of a device participating in the protocol. An authorization solution must take the limitations of the constrained devices into account (all use cases, see also Section 3.1).
- o Secure default settings are needed for the initial state of the authentication and authorization protocols (all use cases).

Rationale: Many attacks exploit insecure default settings, and experience shows that default settings are frequently left unchanged by the end users.

- o Access to resources on other devices should only be permitted if a rule exists that explicitly allows this access (default deny) (see Section 2.4 for example).
- o Usability is important for all use cases. The configuration of authorization policies as well as the gaining access to devices must be simple for the users of the devices. Special care needs to be taken for scenarios where access control policies have to be configured by users that are typically not trained in security (see Sections 2.2, 2.3, and 2.6).
- o Software updates are an important operation for which correct authorization is crucial. Additionally, authenticating the receiver of a software update is also important, for example, to make sure that the update has been received by the intended device.

3.4. Proxies

In some cases, the traffic between endpoints might go through intermediary nodes (e.g., proxies, gateways). This might affect the function or the security model of authentication and access control protocols e.g., end-to-end security between endpoints with Datagram Transport Layer Security (DTLS) might not be possible (see Section 2.5).

4. Privacy Considerations

The constrained devices in focus of this document either collect data from the physical world via sensors or affect their surroundings via actuators. The collected and processed data often can be associated with individuals. Since sensor data may be collected and distributed on a regular interval, a significant amount of information about an individual can be collected and used as input for learning algorithms as part of big data analysis and used in an automated decision making process.

Offering privacy protection for individuals is important to guarantee that only authorized entities are allowed to access collected data, to trigger actions, to obtain consent prior to the sharing of data, and to deal with other privacy-related threats outlined in RFC 6973.

RFC 6973 was written as guidance for engineers designing technical solutions. For a short description about the deployment-related aspects of privacy and further references relevant for the Internet of Things sector, please see Section 7 of RFC 7452.

5. Informative References

[Jedermann14]

Jedermann, R., Poetsch, T., and C. Lloyd, "Communication techniques and challenges for wireless food quality monitoring", Philosophical Transactions of the Royal Society A Mathematical, Physical and Engineering Sciences, May 2014, <<http://rsta.royalsocietypublishing.org/content/372/2017/20130304.short>>.

[Karnouskos11]

Karnouskos, S., "Stuxnet Worm Impact on Industrial Cyber-Physical System Security", IECON 2011 - 37th Annual Conference on IEEE Industrial Electronics Society, pp. 4490-4494 10.1109/econ.2011.612.0048, November 2011, <<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6120048>>.

- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <<http://www.rfc-editor.org/info/rfc7228>>.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, DOI 10.17487/RFC7252, June 2014, <<http://www.rfc-editor.org/info/rfc7252>>.
- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <<http://www.rfc-editor.org/info/rfc7258>>.

Acknowledgments

The authors would like to thank Olaf Bergmann, Sumit Singhal, John Mattson, Mohit Sethi, Carsten Bormann, Martin Murillo, Corinna Schmitt, Hannes Tschofenig, Erik Wahlstroem, Andreas Baeckman, Samuel Erdtman, Steve Moore, Thomas Hardjono, Kepeng Li, Jim Schaad, Prashant Jhingran, Kathleen Moriarty, and Sean Turner for reviewing and/or contributing to the document. Also, thanks to Markus Becker, Thomas Poetsch, and Koojana Kuladinithi for their input on the container monitoring use case. Furthermore, the authors thank Akbar Rahman, Chonggang Wang, Vinod Choyi, and Abhinav Somaraju who contributed to the building automation use case.

Ludwig Seitz and Goeran Selander worked on this document as part of EIT-ICT Labs activity PST-14056; and as part of the CelticPlus project CyberWI, with funding from Vinnova.

Authors' Addresses

Ludwig Seitz (editor)
SICS Swedish ICT AB
Scheelevaegen 17
Lund 223 70
Sweden

Email: ludwig@sics.se

Stefanie Gerdes (editor)
Universitaet Bremen TZI
Postfach 330440
Bremen 28359
Germany

Phone: +49-421-218-63906
Email: gerdes@tzi.org

Goeran Selander
Ericsson
Faroegatan 6
Kista 164 80
Sweden

Email: goran.selander@ericsson.com

Mehdi Mani
Itron
52, rue Camille Desmoulins
Issy-les-Moulineaux 92130
France

Email: Mehdi.Mani@itron.com

Sandeep S. Kumar
Philips Research
High Tech Campus
Eindhoven 5656 AA
The Netherlands

Email: sandeep.kumar@philips.com

