

ACE Working Group
Internet-Draft
Intended status: Informational
Expires: April 30, 2015

L. Seitz, Ed.
SICS Swedish ICT AB
S. Gerdes, Ed.
Universitaet Bremen TZI
G. Selander
Ericsson
M. Mani
Itron
S. Kumar
Philips Research
October 27, 2014

ACE use cases
draft-seitz-ace-usecases-02

Abstract

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

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 lifecycle of a constrained device and are intended to provide a guideline for developing a comprehensive authentication and access control solution for this class of scenarios.

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

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 30, 2015.

Copyright Notice

Copyright (c) 2014 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	3
1.1.	Terminology	4
2.	Use Cases	4
2.1.	Container monitoring	4
2.1.1.	Bananas for Munich	5
2.1.2.	Authorization Problems Summary	5
2.2.	Home Automation	6
2.2.1.	Controlling the Smart Home Infrastructure	6
2.2.2.	Seamless Authorization	7
2.2.3.	Remotely letting in a visitor	7
2.2.4.	Authorization Problems Summary	7
2.3.	Personal Health Monitoring	8
2.3.1.	John and the heart rate monitor	9
2.3.2.	Authorization Problems Summary	10
2.4.	Building Automation	10
2.4.1.	Device Lifecycle	11
2.4.2.	Authorization Problems Summary	13
2.5.	Smart Metering	13
2.5.1.	Drive-by metering	14
2.5.2.	Meshed Topology	14
2.5.3.	Advanced Metering Infrastructure	14
2.5.4.	Authorization Problems Summary	15
2.6.	Sports and Entertainment	16
2.6.1.	Dynamically Connecting Smart Sports Equipment	16

2.6.2. Authorization Problems Summary	17
2.7. Industrial Control Systems	17
2.7.1. Oil Platform Control	17
2.7.2. Authorization Problems Summary	18
3. Security Considerations	18
3.1. Attacks	19
3.2. Configuration of Access Permissions	20
3.3. Design Considerations for Authorization Solutions	20
3.4. Proxies	21
4. Privacy Considerations	21
5. Acknowledgments	22
6. IANA Considerations	22
7. Informative References	22
Authors' Addresses	23

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, due to the devices' limitations, commonly used security protocols are not always easily applicable. As the devices are expected to be integrated in all aspects of everyday life, the application of adequate security mechanisms is required to prevent attackers from gaining control over data or functions important to our lives.

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 lifecycle of a constrained device.

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

Resource Server (RS): The constrained device which hosts resources the Client wants to access.

Client (C): A device which wants to access a resource on the Resource Server.
This could also be a constrained device.

Resource Owner (RO): The subject who owns the resource and controls its access permissions.

2. Use Cases

This section lists use cases involving constrained devices with certain authorization problems to be solved. Each use case first presents a general description of the application area, then one or more specific use cases, and finally a summary of the authorization-related problems device owners need to be solved.

There are various reasons for assigning a function (client or resource server) to a device, e.g. which device initiates the conversation, how do devices find each other, etc. 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 devices 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. The use of such sensor systems makes it possible to continuously track and transmit specific characteristics such as temperature, humidity and gas content during the transportation and storage of goods.

The proper handling of the sensors in this scenario is not easy to accomplish. They have to be associated to the appropriate pallet of the respective container. Moreover, the goods and the corresponding sensors belong to specific customers.

During the shipment to their destination the goods often pass stops where they are transloaded to other means of transportation, e.g. from ship transport to road transport.

The transportation and storage of perishable goods is especially challenging since they have to be stored at a constant temperature and with proper ventilation. Additionally, it is very important for

the vendors to be informed about irregularities in the temperature and ventilation of fruits to avoid the delivery of decomposed fruits to their customers. The need for a constant monitoring of perishable goods has led to projects such as The Intelligent Container (<http://www.intelligentcontainer.com>).

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 with their own company trucks.

The fruit vendor's quality management wants to assure the quality of their products and thus equips the banana boxes with sensors. The state of the goods is monitored consistently during shipment and ripening and abnormal sensor values are recorded. Additionally, the sensor values are used to control the climate within the cargo containers. Since a wrong sensor value leads to a wrong temperature and thus to spoiled goods, the integrity of the sensor data must be assured.

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. Those relaying nodes might belong to different owners. The sensors in the banana boxes cannot always reach the internet during the journey.

The personnel that transloads the goods must be able to locate the goods meant for a specific customer. 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.

When the goods arrive at the supermarket in Munich, the supermarket conducts its own quality check. If no anomalies occurred during the transport, the bananas are admitted for sale.

2.1.2. Authorization Problems Summary

- o U1.1 The device owner wants to grant different access rights to a resource to different parties.
- o U1.2 The device owner wants to control which devices are allowed to present data to the device.

- o U1.3 The device owner wants to grant different access rights for different resources on a device.
- o U1.4 The device owner requires the integrity of sensor data.
- o U1.5 The device owner requires the confidentiality of sensor data.
- o U1.6 The device owner is not always present at the time of access and cannot manually intervene in the authorization process.
- o U1.7 The device owner wants to grant temporary access permissions to a party.
- o U1.8 Messages between client and resource server might need to be forwarded over multiple hops.
- o U1.9 The constrained device might not always be able to reach the internet.

2.2. Home Automation

Automation of the home has the potential to become a big future market for the Internet of Things. A home automation system connects devices in a house to the Internet and thus makes them accessible and manageable remotely. Such devices might control for example heating, ventilation, lighting, home entertainment or home security.

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

Jane and her husband George own a flat which is equipped with home automation devices such as HVAC and shutter control, and they have a motion sensor in the corridor which controls the light bulbs there.

Jane and George can control the shutters and the temperature in each room using either wall-mounted touch panels or their smartphones. Since Jane and George both have a full-time job, 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.

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.

2.2.2. Seamless Authorization

Jane buys a new light bulb for the corridor and integrates it into the home network (how she does that is not in scope). George is not at home, but Jane wants him to be able to control the new device with his smart phone without the need for additional administration effort.

2.2.3. Remotely letting in a visitor

Jane and George 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.

Jane and George have invited Jane's parents over for dinner, but are stuck in traffic and can not arrive in time, while Jane's parents who use the subway will arrive punctually. Jane calls her parents and offers to let them in remotely, so they can make themselves comfortable while waiting.

Jane's parents download an application that lets them communicate with Jane's door-lock and alarm system. Then Jane sets temporary permissions that allow them to open the door, and shut down the alarm when they arrive.

The security system controlling the door-locks and alarm system needs to be at least as secure as for a comparable unautomated home.

2.2.4. Authorization Problems Summary

- o U2.1 A home owner wants to spontaneously provision authorization means to visitors.
- o U2.2 A home owner wants to spontaneously change the home's access control policies.
- o U2.3 A home owner wants to apply different access rights for different users.
- o U2.4 A home owner wants to apply context-based conditions (presence, time) to authorizations, and the devices need to be able to verify these conditions.

- o U2.5 The smart home devices need to be able to communicate with different control devices (e.g. wall-mounted touch panels, smartphones, electronic key fobs).
- o U2.6 The access control configuration of the automated home needs to be secure by default.
- o U2.7 The access control policies need to be easy to edit, even remotely and it needs to be easy to get access with correct authorization.
- o U2.8 The owners of the automated home wants to prevent eavesdroppers from being able to deduce behavioral profiles from the home network.
- o U2.9 Usability is particularly important in this scenario since administrative tasks such as installation, configuration and decommissioning of devices likely need to be performed by the home owners who in most cases have little knowledge of security.
- o U2.10 Home Owners want their devices to seamlessly (and in some cases even unnoticeably) fulfill their purpose. The administration effort needs to be kept at a minimum.

2.3. Personal Health Monitoring

The use of wearable health monitoring technology is expected to grow strongly, as a multitude of novel devices are developed and marketed. The need for open industry standards to ensure interoperability between products has lead to initiatives such as Continua Alliance (continuaalliance.org) and Personal Connected Health Alliance (pchalliance.org). Personal health devices are typically battery driven, and located physically on the user. They monitor some bodily function, such as e.g. temperature, blood pressure, or pulse. They are connected to the Internet through an intermediary base-station, using wireless technologies. Through this connection they report the monitored data to some entity, which may either be the user herself, or some medical personnel in charge of the user.

Medical data has always been considered as 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.

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 position. In case of a cardiac arrest it automatically sends an alarm to an emergency service, transmitting John's current location. The HeartGuard also broadcasts emergency information in the neighborhood to notify doctors or people with certain skills who have been enrolled in an emergency program, e.g. people who got training in heart and lung rescue. For doctors, medical information or diagnosis can be provided with the notification to improve immediate treatment.

The device includes some smart logic, with which it identifies its owner John and allows him to configure the device's settings, including access control.

This prevents situation where someone else wearing that device can act as the owner and mess up the access control and security settings.

John can configure additional persons 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.

However John is a rather private person, and is worried that Jill might use HeartGuard to monitor his location while 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 big a risk for them.

NOTE: Monitoring of some state parameter (e.g. an alarm button) and the position of a person also fits well into an elderly care service. 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 this case it is not the patient that decides about access.

2.3.2. Authorization Problems Summary

- o U3.1 A device owner wants to pre-configure access rights to specific data for persons or groups, in the context of an emergency.
- o U3.2 A device owner wants to selectively allow different persons or groups to access medical data.
- o U3.3 A device owner wants to block access to specific persons in an otherwise allowed group (e.g. doctors in an emergency), if he mistrusts them.
- o U3.4 The security measures could affect battery lifetime of the devices and should changes of battery are highly inconvenient.
- o U3.5 Devices are often used with default access control settings.
- o U3.6 Device users are often not trained in computer use and especially computer security.
- o U3.7 Security mechanisms themselves could provide opportunities for denial of service attacks on the device.
- o U3.8 The device provides a service that can be fatal for the device owner if it fails. Accordingly, the device owner wants a security mechanism to provide a high level of security.

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.

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 light 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 service company.

2.4.1. Device Lifecycle

2.4.1.1. Installation and Commissioning

A building is hired out to different companies for office space. This building features various automated systems, such as a fire alarm system, which is triggered by several smoke detectors which are spread out across the building. It also has automated HVAC, lighting and physical access control systems.

A vacant area of the building has been recently leased to company A. Before moving into its new office, Company A wishes to replace the lighting with a more energy efficient and a 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.

Company C gets the necessary authorization from the service company to interact with the existing Building and Lighting Management System (BLMS). 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. After installation (wiring) of the new lighting devices, the commissioner adds the devices into the 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. For this, the commissioner creates the authorization rules on the BLMS which define which lights form a group and which sensors /switches/controllers are allowed to control which groups. 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.

2.4.1.2. Operational

Company A's staff move into the newly furnished office space. Most lighting is controlled by presence sensors which control the lighting of 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 office 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 down 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 with company A. On a really hot day James who works for company A turns on the air condition in his office. Lucy who works for company B wants to make tea using an electric kettle. After she turned it on she goes outside to talk to a colleague until the water is boiling. Unfortunately, her kettle has a malfunction which 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 broadcaster multicast) to alert the staff of both companies. 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 although he turned on his air condition because the fire alarm overrides the manual setting by sending commands (broadcast or multicast) to switch off all the air conditioning.

The fire department is notified of the fire automatically and arrives within a short time. After inspecting the damage and extinguishing the smoldering fire a fire fighter resets the fire alarm because only the fire department is authorized to do that.

2.4.1.3. Maintenance

Company A's staff are annoyed that the lights switch off too often in their rooms if they work silently in front of their computer. Company A notifies the commissioning Company C about the issue and asks them to increase the delay before lights switch off.

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

2.4.1.4. 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.

Company C again gets the necessary authorization from the service company to interact with the BLMS. The commissioner now deletes any rules that allowed handheld controllers authorization to control the lighting. Additionally the commissioner instructs the BLMS to push these new rules to prevent cached rules at the end devices from being used.

2.4.2. Authorization Problems Summary

- o U4.1 Device owners want to be able to add a new device to their administrative domain (commissioning).
- o U4.2 Device owners want to be able to integrate a device that formerly belonged to a different administrative domain to their own administrative domain (handover).
- o U4.3 Device owner want to be able to remove a device from their administrative domain (decommissioning).
- o U4.4 Device owners want to be able to delegate selected administration tasks for their devices to others.
- o U4.5 The device owner wants to be able to define context-based Authorization rules.
- o U4.6 The device owner wants to be able to revoke granted permissions and delegations.
- o U4.7 The device owner wants to allow only authorized access to device resources (default deny).
- o U4.8 The device owner wants to be able to authorize a device to control several devices at the same time using a multicast protocol.
- o U4.9 Device owners 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 devices.

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 the 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.

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. The meters do so by sending data to a base station. 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.

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.

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 pre-defined schedules. Based on this technology, new services make it possible for consumers to control their utility consumption and reduce costs by supporting new tariff models from utility companies, and more accurate and timely billing.

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. Two possible intermediate levels are:

- o Head-End System (HES) which is hardware and software that receives the stream of meter data and exposes an interface to the MDM.
- o Data Collection (DC) units located in a local network communicating with a number of smart meters and with a backhaul interface communicating with the HES, e.g. using cellular communication.

For reasons of efficiency and cost end-to-end connectivity is not always feasible, so metering data is stored in batches in DC for some time before being forwarded to the HES, and in turn accessed by the MDM. The HES and the DC units may be operated by a third party service operator on behalf of the utility company. One responsibility of the service operator is to make sure that meter readings are performed and delivered to the HES. An example of a Service Level Agreement between the service operator and the utility company is e.g. "at least 95 % of the meters have readings recorded during the last 72 hours".

2.5.4. Authorization Problems Summary

- o U5.1 Devices are installed in hostile environments where they are physically accessible by attackers. Device owners want to make sure that an attacker cannot use a captured device to attack other parts of their infrastructure.
- o U5.2 Device owners want to restrict which entities are allowed to write data to the devices and thus ensure the integrity of the data on their devices.
- o U5.3 The device owner wants to control which entities are allowed to read data on the devices and protect such data in transfer.
- o U5.4 The devices may have intermittent Internet connectivity.

- o U5.5 The device owner is not always present at the time of access and cannot manually intervene in the authorization process.
- o U5.6 When authorization policies are updated it is impossible, or at least very inefficient to contact all affected devices directly.
- o U5.7 Messages between a client and the device may need to be stored and forwarded over multiple nodes.

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 functionalities can be integrated into fitness equipment, games and even clothes. Owners can carry their devices around with them at all times.

Usability is especially important in this area since owners 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 (preferably none at all).

The required level of security will in most cases be low since security breaches will likely have less severe consequences. The continuous monitoring of data might however enable an attacker to create behavioral or movement profiles. Moreover, the aggregation of data can seriously increase the impact on the privacy of device owners.

2.6.1. Dynamically Connecting Smart Sports Equipment

Jody is a 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 detect that the watch can be connected with Jody's shoes and then can additionally 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 doesn't want Jody to access her training plan. 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. Jody's shoes connect to Lynn's watch after only a press of a button because Jody already configured access rights for devices that belong to Lynn a while ago.

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

2.6.2. Authorization Problems Summary

- o U6.1 The owner of a device wants to be able to grant access rights dynamically when needed.
- o U6.2 The owner wants the configuration of access rights to work with very little effort.
- o U6.3 The device owner wants to be able to preconfigure access policies that grant certain access permissions to devices with certain attributes (e.g. devices of a certain user) without additional configuration effort at the time of access.
- o U6.4 Device owners wants to protect the confidentiality of their data for privacy reasons.
- o U6.5 Devices might not have an Internet connection at the time of access.

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 this kind of systems are, especially when connected to the Internet. The severity of these vulnerabilities are exacerbated by the fact that many ICS are used to control critical public infrastructure, such as power, water treatment of traffic control. Nevertheless the economical advantages of connecting such systems to the Internet can be significant if appropriate security measures are put in place.

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.

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.

The main interest of the operator is to ensure the integrity of control messages and sensor readings. The access to some resources needs to be restricted to certain clients, e.g. the operator wants wireless actuators only to accept commands by authorized control units.

The owner of the platform also wants to collect auditing information for liability reasons.

2.7.2. Authorization Problems Summary

- o U7.1 The device owner wants to ensure that only authorized clients can read data from sensors and sent commands to actuators.
- o U7.2 The device owner wants to ensure that data coming from sensors and commands sent to actuators are authentic.
- o U7.3 Some devices do not have direct Internet connection.
- o U7.4 Some devices have wired connection while other use wireless.
- o U7.5 The execution of unauthorized commands in an ICS can lead to significant financial damage, and threaten the availability of critical infrastructure services. Accordingly, the device owner wants a security solution that provides a very high level of security.

3. Security Considerations

As the use cases listed in this document demonstrate, constrained devices are used in various application areas. The appeal of these devices is that they are small and inexpensive. That makes it easy to integrate them into many aspects of everyday life. Therefore, the devices will be entrusted with vast amounts of valuable data or even control functions, that need to be protected from unauthorized access.

Moreover, the aggregation of data must be considered: attackers might not only collect data from a single device but from many devices, thus increasing the potential damage.

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 gained 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 and hence 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 owners 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.

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 ACE, but should be kept in mind by developers of authorization solutions.

Denial of service (DoS) attacks threaten the availability of services a device provides. E.g., an attacker can induce a device to perform steps of a heavy weight security protocol (e.g. Datagram Transport Layer Security (DTLS) [RFC6347]) before authentication and authorization can be verified, thus exhausting the device's system resources. This leads to a temporary or - e.g. if the batteries are drained - permanent failure of the service. For some services of constrained devices, availability is especially important (see Section 2.3). Because of their limitations, constrained devices are especially vulnerable to denial of service attacks. Solution designers must be particularly careful to consider these limitations in every part of the protocol. This includes:

- o Battery usage
- o Number of message exchanges required by security measures

- o Size of data that is transmitted (e.g. authentication and access control data)
- o Size of code required to run the protocol
- o Size of RAM memory and stack required to run the protocol

Another category of attacks that needs to be considered by solution developers is session interception and hijacking.

3.2. Configuration of Access Permissions

- o The access control policies of the Resource Owner need to be enforced (all use cases): The access control policies set by the Resource Owner need to be provisioned 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 e.g. Section 2.2, Section 2.3, Section 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. Design Considerations for Authorization Solutions

- o Devices need to be enabled to enforce the owner's authorization policies without the owner's intervention at the time of the access request (see e.g. Section 2.1, Section 2.2, Section 2.4, Section 2.5).
- o Authorization solutions need to consider that constrained devices might not have internet access at the time of the access request (see e.g. Section 2.1, Section 2.3, Section 2.5, Section 2.6).
- o It should be possible to update access control policies without manually re-provisioning individual devices (see e.g. Section 2.2, Section 2.3, Section 2.5, Section 2.6).

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

- o Owners might define authorization policies for 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.
- o It must be possible to dynamically revoke authorizations (see e.g. Section 2.4).
- o The authentication and access control protocol can put undue burden on the constrained resources of a device participating in the protocol. An authorization solutions must take the limitations of the constrained devices into account (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).
- 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 home scenarios where access control policies have to be configured by users that are typically not trained in security (see Section 2.2, Section 2.6).

3.4. Proxies

In some cases, the traffic between Client and Resource Server 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 Client and Resource Server with DTLS might not be possible (see Section 2.5).

4. Privacy Considerations

Many of the devices that are in focus of this document register data from the physical world (sensors) or affect processes in the physical world (actuators), which may involve data or processes belonging to individuals. To make matters worse the sensor data may be recorded

continuously thus allowing to gather significant information about an individual subject to the sensor readings. Therefore privacy protection is especially important, and Authentication and Access control are important tools for this, since they make it possible to control who gets access to private data.

Privacy protection can also be weighted in when evaluating the need for end-to-end confidentiality, since otherwise intermediary nodes will learn the content of potentially sensitive messages sent between a client and a resource server and thereby endanger the privacy of the individual that may be subject of this data.

In some cases, even the possession of a certain type of device can be confidential, e.g. owners might not want to others to know that they are wearing a certain medical device (see Section 2.3).

The personal health monitoring use case (see Section 2.3) indicates the need for secure audit logs which impose specific requirements on a solution. Auditing is not in the scope of ACE. However, if an authorization solution provides means for audit logs, it must consider the impact of logged data for the privacy of the owner and other parties involved.

Suitable measures for protecting and purging the logs must be taken during operation, maintenance and decommissioning of the device.

5. 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, and Andreas Backman 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.

6. IANA Considerations

This document has no IANA actions.

7. 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.

- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012.
- [RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, May 2014.
- [RFC7252] Shelby, Z., Hartke, K., and C. Bormann, "The Constrained Application Protocol (CoAP)", RFC 7252, June 2014.
- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, May 2014.

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