FP7-SMARTCITIES-2013
Project number: 609132
http://www.e-balance-project.eu/

Co-funded by
the European Union

e-balance

Deliverable D4.2

Detailed security and privacy specification and implementation

Editor:	Krzysztof Piotrowski (IHP)
Dissemination level: (Confidentiality)	PU
Suggested readers:	Consortium/Experts/other reader groups
Version:	1.0
Total number of pages:	23
Keywords:	Security, Privacy, Smart Grid

Abstract

This deliverable describes the security and privacy solution for the e-balance system. The focus of this document is on the lower layers in the Communication Platform that provide the basic mechanisms necessary to implement the security and privacy protocols for the Energy Management Platform. However, this document provides also a holistic view on the security and privacy solution for the e-balance project.

Disclaimer

The information, documentation and figures available in this deliverable are written by the e-balance partners under EC co-financing (project number: 609132) and does not necessarily reflect the view of the European Commission.

Impressum

[Full project title] Balancing energy production and consumption in energy efficient smart neighbourhoods

[Short project title] e-balance

[Number and title of work-package] WP4, Communication Platform
[Document title] Detailed security and privacy specification and implementation
[Editor: Name, company] Krzysztof Piotrowski, IHP

[Work-package leader: Name, company] Daniel Garrido, UMA

Copyright notice

Executive Summary

The security aspects of the energy grid are gathering more and more attention in the recent years. This is mainly due to the fact that the energy grid is part of the critical infrastructure and it shall be protected from any interference that can disturb its working, especially if this interference is an intentional act. The security aspects actually become critical in the even more digitalized approach of the energy grid – the smart grid, where the exchanged data controls the state of the grid (sometimes in real-time) and it is sometimes even easier to interfere with this data exchange due to the large amount of the communicated data.

Our social studies confirmed that the acceptance of an approach such as e-balance highly depends on the privacy protection provided by the system. As the effect of the e-balance approach depends on the willingness of customers to share information on their energy consumption and energy production profile privacy protection is of utmost importance.

This document presents the holistic view on the e-balance security and privacy approach. In this document we name the identified security and privacy issues that need consideration and propose technical solutions to cope with them. We identify and describe the implementation of the different security modules that are applied in the e-balance solution to protect it. Here, the protection is, mainly, against attacks from outside the system, but it is not limited to external attack sources only.

Whenever possible, the implementation presented in this deliverable focuses on available and standard mechanisms and solutions that can be used to protect the devices and the communication. These mechanisms provide the background for and are completed by the security and privacy solutions we have developed within the project, like the security and privacy protocol described in the deliverable D5.4.

List of authors

Company	Author
IHP	Krzysztof Piotrowski Ievgen Kabin Peter Langendörfer
UMA	Daniel Garrido Eduardo Cañete Jaime Chen
EDP	Nuno Emanuel Pereira
EFA	Paulo Rodrigues Alberto Bernardo
ALLI	Marcel Geers

Table of Contents

List of Tables

List of Figures

Abbreviations

AES	Advanced Encryption Standard
APT	Advanced Packaging Tool
CaMyTs	Castellucia-Mykletun-Tsudik
CCM	Counter with CBC-MAC mode
CMU	Customer Management Unit
DH	Diffie-Hellman
DSO	Distribution System Operator
ECC	Elliptic Curve Cryptography
ECDSA	Elliptic Curve Digital Signature Algorithm
HMAC	Hash-based Message Authentication Codes
IT	Information Technology
LVGMU	Low Voltage Grid Management Unit
MAC	Message Authentication Codes
MVGMU	Medium Voltage Grid Management Unit
OFB	Output Feedback mode
PKI	Public Key Infrastructure
PLC	Power-line Communication
RSA	Rivest-Shamir-Adleman cryptosystem
SG	Smart Grid
SHA	Secure Hash Algorithm
SQL	Structured Query Language
SSH	Secure Shell
TLGMU	Top Level Grid Management Unit
TLS	Transport Layer Security

1 Introduction

The security aspects of the energy grid are gathering more and more attention in the recent years. This is mainly due to the fact that the energy grid is part of the critical infrastructure and it shall be protected from any interference that can disturb its working, especially if this interference is an intentional act. This problem actually became critical in the era of Smart Grid (SG), where the exchanged data controls the state of the grid and where it is sometimes even easier to interfere with this data exchange due to the large scale of the communication, compared to the energy grid several years ago.

Further, the great amount of information that is gathered in a smart grid system often causes privacy issues. The smart grid infrastructure providers require high quality data often allowing profiling their customers with a great detail. And since the customers are nowadays really concerned about their privacy (what was confirmed by our social studies) there is a serious conflict of interests that need to be addressed and resolved.

Due to the importance of the subject we address it in the e-balance project. In this document we provide the holistic view on our security approach, we identify the issues and propose technical solutions to cope with them. The solutions are represented in a form of modules constituting the security architecture of e-balance. The implementation presented in this deliverable focuses on the standard mechanisms and approaches that can be used to protect the devices and the data. Use of standards increases the applicability and acceptance of the proposed solution and also assures to some extend the protection level of the system as a whole. On the other hand, these mechanisms provide the background for the security and (mainly) privacy protocol described in deliverable D5.4 [2].

Figure 1 presents the position of the deliverable D4.2 within the e-balance project. It is part of the communication platform; it defines and provides all the security and privacy mechanisms necessary for the middleware and for the whole e-balance system to work securely and reliably.

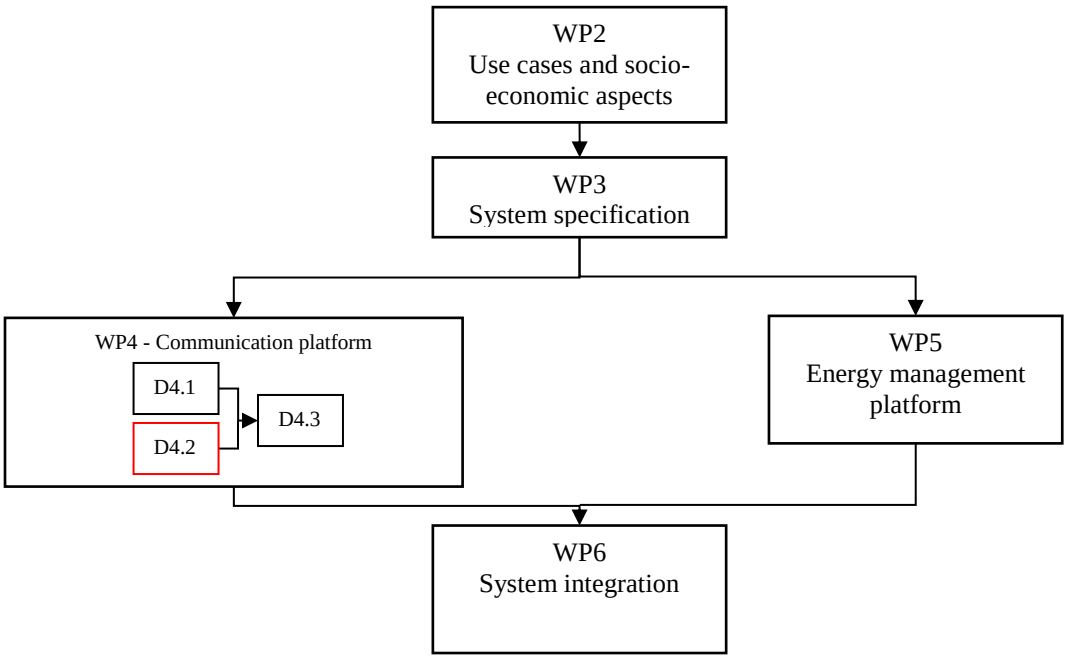

Figure 1: The position of the deliverable D4.2 within the e-balance project

The two deliverables on security and privacy, namely D4.2 and D5.4 are complementing each other. In the following paragraphs we describe the coverage of each of these documents, based on the security architecture, defined in deliverable D3.2 [1] and depicted in Figure 2.

The deliverable D5.4 covers mainly the security and privacy module in the energy management platform – providing more details of the implementation as well as the evaluation of the approach. The remaining security related modules (blue coloured boxes in Figure 2) are covered by deliverable D4.2 – that presents the implementation details of the remaining security related modules, mainly present in the communication platform. The data access control module is addressed in both documents.

There are four major categories of security and privacy modules within the communication platform that are covered by this document. They can be defined as follows.

- Mechanisms for privacy and security (the library base for the other mechanisms)
- Node protection and maintenance (code update, node parameterisation and node protection)
- Trust and group management (PKI related solutions for authentication, group and trust management)
- Network security (embedded in the networking protocols or added on top of them)

The following modules, related to security and privacy, are defined within the communication platform.

- Node maintenance, protection, parameterisation and reset
- Data access control
- Secure group management
- Secure storage
- Network stack security and privacy

The instantiations of these modules and their very specific low level functions depend very much on the hardware and software platform of the individual management unit. However, these modules should fulfil a specific set of requirements in order to provide their function properly. Additionally, the standard Information Technology (IT) security solutions change and are updated with respect to new security settings and bug-fixes. Thus, this document will mainly focus on the required functions rather than on specific security strength parameters, as these can be updated to keep the security level up-to-date or according to specific needs. However, we define the security strength parameters currently used in our implementation. And the specific solutions and libraries we have chosen for the prototype are also named.

It is suggested to read this deliverable document prior to the deliverable D5.4 document.

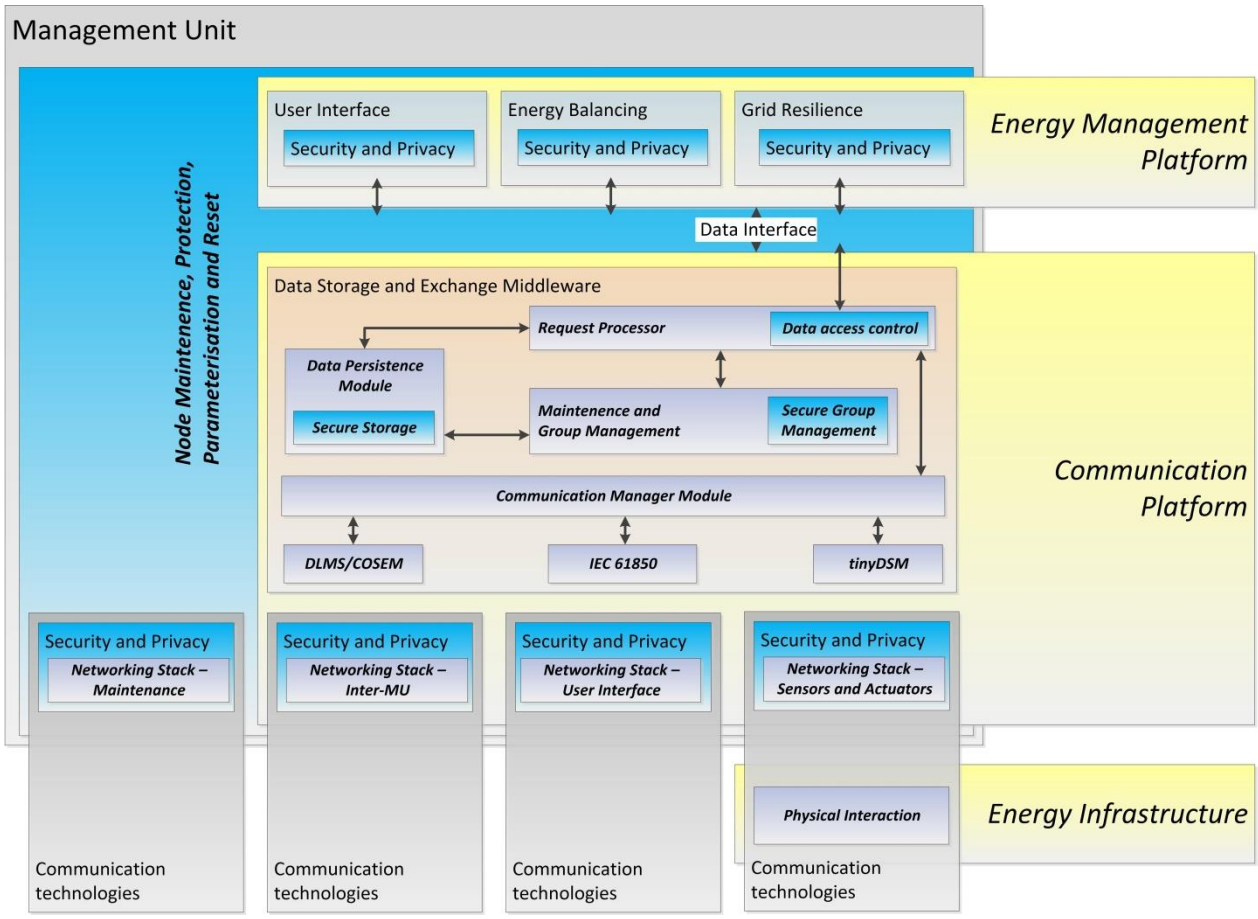

Figure 2: The security architecture of the e-balance management unit

1.1 Relation to the scientific and technical objectives

This document, together with deliverable D5.4 relates to the e-balance scientific and technical objectives: STO2, point d), and STO3, point d). They present the analysis and results towards providing privacy and security protection for the communication platform as well as for the energy management platform.

1.2 Beyond state-of-the-art contributions

Depending on the system level there are different considerations regarding the possible security solutions, e.g., due to different hardware characteristics and performance. Embedded systems and especially wireless sensor nodes can be easily attacked. This is due to the fact that they are normally unprotected by cryptographic means. This is due to the fact that both types of devices suffer from severe resource constraints, e.g., energy resources and processing power so that standard cryptographic approaches cannot be applied. Thus, there is a necessity of development of the lightweight cryptographic solutions, which take the above mentioned constraints into consideration and are able to ensure the needed level of the security.

The security design process for smart energy systems involves multiple layers and issues to be considered together, in order to provide a secure, dependable and efficient solution. These issues add to the total set of system functional requirements and may influence the choice of underlying system modules, e.g., the required processing power may induce the choice of more powerful nodes.

We are focusing on researching studies and prototypes on authentication and encryption methods, as well as integrity solutions, for the information flowing through the smart energy network. We suggest a solution aligned with the conventional and standard guidelines for a secure architecture, based on the assurance of the following security properties: confidentiality, integrity and authenticity. We believe that by guaranteeing such properties on an end-to-end basis (security in depth strategy) we will robust the infrastructure and develop a strong defence solution for a more resilient Smart Grid for a Smart City.

We argued in the DoW that in order to ensure high acceptance and low double work that we will apply standards whenever possible. In this document we propose a combination of such approaches with an extension of our own solutions providing an implementation that is conform to the suggestions by standards (IEC 62351).

Our approach extends IEC 62351 by a multi-party privacy protection protocol and by a low complexity and scalable privacy protecting mechanism.

We propose an approach for secure group management that helps to maintain the set of management units in the system in a secure and reliable way. This approach allows also the application of buffering proxies that help to cope with communication related limitations in the approach where the devices are distributed over zones with different security requirements and with rules defining the allowed communication patterns.

In our implementation of the security and privacy protocol for the energy management platform we also are using as a base an already available approach – the Castellucia-Mykletun-Tsudik (CaMyTs) homomorphic encryption scheme. However, we have modified the approach to reduce its complexity and apply it in a modified context, what adds the possibility to use data aggregators others than the grid operator to provide even more protection for the customer (and other stakeholder) privacy.

Further, we propose a privacy protecting protocol that allows the data owners (customers and other stakeholders) to define and control the allowed use of their data. The data owners can specify which other stakeholders can use the specific data and for what purpose (by defining the services of the stakeholder). The middleware implements the defined security and privacy policy. The data can be further encrypted using the modified CaMyTs encryption mechanism.

These approaches are described in detail in deliverable D5.4.

2 Security and privacy context

The smart grid solution defined by the e-balance system is a quite complex set of devices and communication means used for collecting, processing and exchanging the data handled by the system. In this section we will try to describe the context the system is working within, in order to define the security and privacy solution.

2.1 Problems and required system properties – the stakeholder's view

There are several security related problems/requirements a smart grid has to cope with. Most of them are similar to the problems known in the general Internet or office applications. But they get a new flavour in the context of smart grid and sometimes these two worlds have different priorities, making the smart grid problems more similar to those known from industry automation. Additionally, the individual problems, even in the same context, can also lead to contradicting goals. Thus, the priorities define the importance of the protection means, as well as the requirements to the underlying infrastructure, like the required computational power, the features of the communication technologies, etc.

Not satisfying these security requirements will enable actions that may harm the system, its components and the involved stakeholders. These actions may be performed by some attackers by means of physical interaction with the devices, but also by interacting with the (wirelessly) exchanged data. Of course, whenever possible, direct interaction with the system devices for not authorized persons shall not be possible. But even then, interfering with the communication may enable several kinds of attack.

The first requirement is data confidentiality. Providing confidentiality helps to prevent theft of information. This is the first and major point where the privacy may be harmed. The simplest form of the data theft is eavesdropping – passive and usually not detectable reception of unprotected data that can reveal both private and business information of the affected stakeholders. Thus, providing confidentiality covers mainly the protection of the data, but it also includes assuring that only authorized stakeholders get access to the data.

The second requirement is data integrity. Providing data integrity helps to prevent unauthorized modification of the data. Data manipulation is more complicated than eavesdropping and it already requires active interaction – it is not only about receiving information from the system, but it also includes active changing of the transmitted data or generating own data. Such manipulated data may cause system malfunction due to falsified system state or control command. Data manipulation can be realized as modifying the transmitted data bits by generating an interfering signal that changes the information received by the recipient. Due to the required active actions, data manipulation can be detected and this is the main aim of mechanisms providing data integrity – assuring that the data is not manipulated, or at least allowing detecting data manipulations.

The third requirement is data availability. The smart grid is a control system that may require some data to be exchanged or provided in real-time (within the defined allowed maximum delays) to work properly – to define correct actions based on the data. Interfering here with the data exchange and blocking it may cause system malfunction due to lack of data needed for taking some particular decisions – a denial of service. Providing availability may be a hard task if the underlying communication infrastructure is working on its limits with respect to throughput (or if the data sources can be easily switched off) and if there is no redundancy (alternative communication means). Additionally, it may contradict with providing data confidentiality and data integrity, if the data processing devices are working on their limits with respect to processing power. The allowed delay thresholds can be exceeded due to the time needed to perform the mechanisms needed for providing data confidentiality and integrity. Thus, in order to assure data availability the possible constraints with respect to computational power of the devices and to the networking throughput need to be identified and the system has to be defined in a way to avoid these bottle necks. Additionally, the devices and the networking can be protected against actions that may deactivate them (lock device or block communication). Data availability is also supported by redundancy - data replication, redundant devices and redundant communication means.

The fourth requirement is non-repudiation. An action related to the data should be performed on behalf of some stakeholder and the fact that it has been executed by that particular stakeholder has to be undeniable. Providing non-repudiation shall protect against cases, where a stakeholder denies performing an action she actually did or claims that she performed an action that was in fact not executed. Transaction logs providing

the history of authenticated (signed) requests and replies to these at each device allow reconstructing the sequence of events. This is necessary in case of misbehaving to detect who did what and also in case of conflicts (e.g. on business level) to track back the performed actions.

A further requirement is authentication. Providing authentication protects against masquerading – one stakeholder cannot claim she is another one and perform actions, e.g., access data, on behalf of the other one. All system participants have to be identifiable and the identities have to be provable. A very close requirement is to provide authorisation. Authorisation means to assure that some actions are available only for specific parties or stakeholders.

Modifying the data as it is transmitted can influence the state and the decisions taken by the system, but this effect can also be achieved by influencing the data processing directly. Thus, an important requirement is to provide means to protect the system against execution of non-authorized and possibly malicious program code. The smart grid system involves a diversity of stakeholders and these stakeholders process their own data, but also the data of their customers. The data processing code can be developed by the stakeholder and can be executed at the system device (management unit). In this case it is absolutely necessary to provide a secure execution environment for those applications, so that the applications of different stakeholders are unable to interact directly and interfere or distract each other.

The priorities of these requirements in the initial stage of smart grid were favouring data availability as it is critical for the responsiveness and reliability of the system. But recently confidentiality and data integrity are gaining more and more importance, since they are critical for several aspects in smart grids, e.g., dependability. First, they are important for the acceptance of the system by the customers and also other stakeholders, as they want to be absolutely sure that their private and business data is protected against unauthorized reception and modification. Further, as part of critical infrastructure, the smart grids shall be immune to attacks that can influence their working (modifying the data to influence the system actions). Nowadays, the technology (combined with social techniques) allows attacking almost every ICT system, but it should be hard (and expensive) enough to discourage the attackers from doing this. And the possible damage shall be limited to the minimum.

Since data integrity and data confidentiality come at the price of more computational power required, the smart grid systems have to be designed properly (with respect to computational power), not to affect data availability. Additional replication of data and redundant communication channels help to cope with denial of service attacks.

From the point of view of the different stakeholders involved in a smart grid system, there are different interests in the above mentioned security properties. Some of these security properties are interesting for all the parties, but some show contradicting interests. Confidentiality is important for all the parties. For the customers it means that their private data is protected, for other stakeholders it means that their business data is protected. Data integrity is mainly important for the grid operator, who is providing the infrastructure and may be the main party responsible for its dependability; however the customers and other stakeholders definitely require from the system that their data is not modified. Availability is crucial for the grid operator and other stakeholders running the energy control algorithms. Authentication is crucial while providing other stakeholders with access to own data. All these properties are important to all the stakeholders, even if they are not directly aware about the fact.

There is however one aspect that is dividing the stakeholder group – it is privacy. The stakeholders running the energy management control algorithms, like an energy supplier, want that the data available to these algorithms contains very detailed information. This can cause a conflict with the owners of the data, since they usually do not want to reveal detailed information that allows precise profiling, i.e., detailed analysis of their behaviour.

The goal of the e-balance approach is to provide the same and high level of protection for all involved stakeholders. All of them shall be assured that their data and devices are protected against misuse, of course as long as the stakeholders behave along the rules. Misbehaviour shall be detected and logged. Further, it is important that the system is transparent and allows auditing the internal procedures.

2.2 Events that can disrupt the system – the threat model

A smart grid system is a cyber-physical system with the power system part and the cyber equipment part that includes monitoring and controlling devices with networking. The aim of this complex system is to provide a

stable and reliable energy distribution with as few as possible and as short as possible interruptions in the delivery. Thus, the control system is usually already developed to be able to react to some events in the power domain in a more or less autonomous manner. Our control mechanisms that are developed within WP5 also provide/support such behaviour. But additionally, the control system itself has also to be developed in a way that it is able to cope with events that may disrupt its working and, as a result, possibly cause the power system failure. The events can be of natural kind, like fire, earthquake, etc. But they may also be an intentional act by parties, like terrorists, cyber-kids, business competitors, unhappy employees, etc. In e-balance we focus on means to prevent intentional attacks. These means provide inherently also protection against accidental or natural influences on the energy grid.

In the e-balance system the communication platform is the part that handles the data exchange and storage and therefore the basic security means are integrated here.

Attacker classes

We envision three main groups of attackers; these are driven by different goals and have different budgets available to execute their actions. These groups are: Kiddie/Hacker, Terrorists and Industrial espionage. Other possible attackers are more or less close to the three groups defined here. And the slight differences in motivation and budget do not provide more details to the discussion.

The first group, the Kiddie/Hacker has a limited budget. The amount of money they can invest in the attack ranges from no investment up to several thousands of Euros. Additionally, there are free tools and scripts available in the Internet that can be used by this group of attackers, what further decreases the needed investments. The attackers from this group are interested in either showing themselves to the public (being famous) or in gathering some private data of individuals (neighbours espionage). Thus, their attacks are rather limited in the area they cover and they are usually not causing large damage. If the goal is to become famous then a big damage may be undesirable (to avoid legal consequences), but it depends on the audience. Similar, if the goal is to gather some private data of another stakeholder, then the attack is mainly based on eavesdropping. In both cases, if large damage occurs, then it was rather unintentional by the attacker and, on the other hand, exposes huge security breaches in the system. However, an organized and distributed attack by a reasonable number of Kiddie/Hacker attackers can already be considered as an attack belonging to one of the following groups and can have also corresponding goals.

The second group, the Terrorists has a much greater budget, ranging from thousands to millions of Euros. The aim of a terrorist attack is to maximize damage and to maximize the publicity after the fact. This means that some leaking of the attack sources may be acceptable. It does not mean that all the information about the attack sources and other details on the attack approach have to be exposed, but especially in case of a suicide-like attack, maximizing the damage has priority over hiding the attack source.

The third group, the Industrial espionage also has a great budget, going up to millions of Euros. However, compared to the Terrorists group, the goal of the Industrial espionage group is rather the gathering of data, high quality and specific information done in a way that the attack cannot be detected neither while it is performed nor after the fact. The disclosure of the fact that an attack was performed and proof of it could lead to lawsuit actions. Further, in case the attack indeed has the goal to cause damage, then hiding the attack source has an even higher priority.

Attacking Security Goals

The following paragraphs name the envisioned attack ways, name the attacker groups and discuss the potential loss and drivers for the attacks. The attacks may be targeting a specific goal, requiring knowledge about the system or may be executed as unspecific attack just to generate some imbalance in the grid.

Eavesdropping allows executing espionage targeting the individual customers, groups of them as well as other stakeholders. The gathered information can be further used by the attacker or a third party for business purposes or even for blackmailing. Eavesdropping includes listening to the medium (powerline, wireless communication). However the devices of the e-balance system can be physically exposed to potential attackers. This allows of more sophisticated attacks from the area of side channel attack, like measuring the power consumption or measuring the electromagnetic fields or temperature around and within the devices.

To avoid these attacks it is very important to provide proper physical protection for the devices and procedures, i.e., limiting physical access to authorized persons only. Further, for such an attack to be successful it is necessary to have the knowledge on how to interpret these gained measurements. Both these aspects are often related to human factors, i.e., the physical access can be gained due to employees not following defined security procedures, while hardware documentation leaks and human skills misuse can also be caused by corrupted employees.

The modification of the data as well the direct modification of the system processes (running malicious code or influencing valid code by changing data in memory) allows influencing the system state and behaviour. Depending on the specific target of the attack and its dimension the attack may be aiming at disabling parts of the grid, causing a blackout or even at physical destruction of the grid assets or customer devices.

The attack can be realized by a legal participant of the system that interacts with the other system participants and components according to the defined protocols, but misbehaves. It may also be realized by a third and external party that has to find a way to interact with the system first, e.g., via a legal participant of the system. These two options differ in the initial access to the system details, i.e., the knowledge on the used protocols and possession of credentials recognized by the system.

Further, the attack can be realized with hiding the identity of the attacker, e.g., by masquerading, requiring knowledge of the system or may be realized as a suicide attack – without taking care of the consequences.

2.3 Standard security approaches

The smart grid is also a complex system with respect to the diversity and the size of the different areas it covers. These areas can have different security requirements due to different level of exposure to the attacks. Our architecture supports an approach that aims at addressing all security gaps properly and to define the proper mechanisms to close these gaps, where the specific target strength of the applied cryptographic mechanisms is finally adapted for the target deployment. Thus, the actual protection level and its costs for a specific implementation/deployment are adapted according to the individual requirements defined by the individual system provider.

Standard security approaches use symmetric and asymmetric cryptographic approaches that are used for different purposes. While the asymmetric cryptography, like Elliptic Curve Cryptography (ECC) or RSA (*Rivest-Shamir-Adleman*) are mainly used for authentication, non-repudiation and sometimes integrity of messages, symmetric approaches like the Advanced Encryption Standard (AES) are used for confidentiality, i.e., encrypting large amounts of data. There are also other mechanisms, like Message Authentication Codes (MAC) and Hash-based Message Authentication Codes (HMAC) that are also based on symmetric approaches and help to provide data integrity and authentication. Other approaches like protocols for secret agreement, like Diffie-Hellman (DH) help to exchange credentials in a secure way. Public Key Infrastructure (PKI) provides means to share and verify certificates for authentication purposes, or solutions for protecting the communication, like the Transport Layer Security (TLS) providing means for exchange of security credentials that are further used to secure the communication.

Table 1: The security zones

Zone	What is included
L5	Households, Customers, Internet
L4	IT, Billing services, historian data systems, GIS, Office environment
L3	SCADA EMS/DMS
L2	Secondary grid assets, like measurement equipment, RTUs, etc.
L1	Primary grid assets, like switchgear, transformers, cables, lines, etc.

In order to help coping with the security issues, the smart grid can be split into zones. This approach is widely used in practice as it helps to define more specific security requirements for each zone and rules for the interactions between the zones. One of the examples is the structuring of the system as proposed by the IEC 62351 standard. Here five zones are proposed from which zone L1 has the highest security requirements

and zone L5 the lowest. One of the rules controlling the communication between the zones can be that communication requests are only allowed to be issued from a zone with a lower number, i.e., request from zone L4 targeting zone L3 is not allowed. In this approach the view is usually more from the Distribution System Operator (DSO) perspective and the security requirements become more restrictive with the lowering the zone number. The coverage of a zonal approach applicable for the e-balance is described in Table 1 and is further depicted in Figure 3.

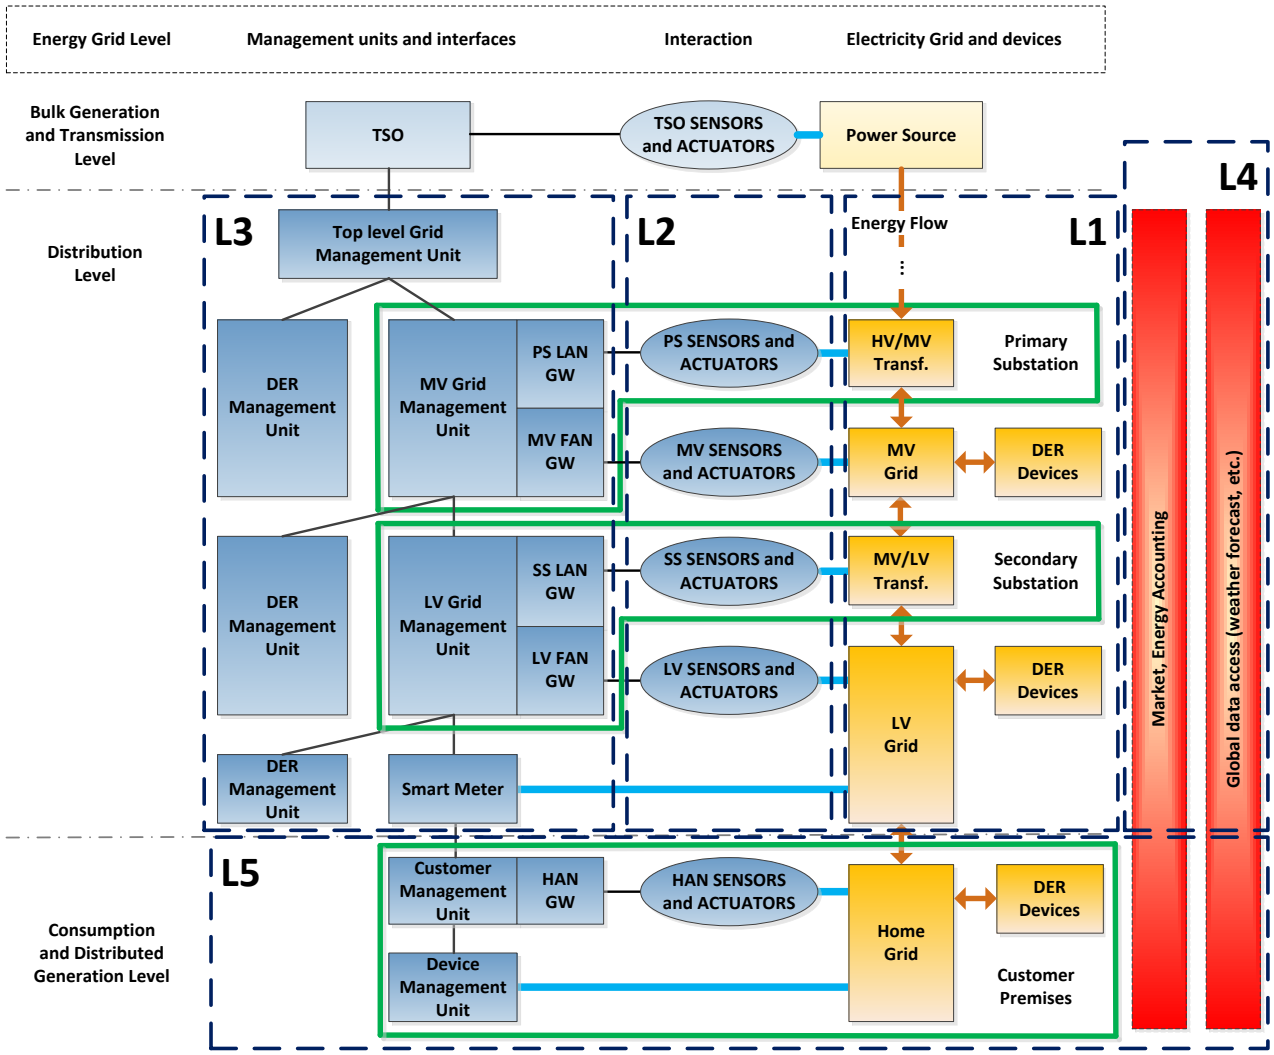

Figure 3: The security zones mapped onto the e-balance system

In order to reduce costs standard security solutions from the office IT can be applied in the Smart Grid, but they have to be adapted for the slightly different requirements and this adaptation has to be done while keeping in mind the critical character of the underlying infrastructure – the energy grid.

Supporting of the zonal approach simplifies the definition of security requirements, but it may also slightly increase the complexity of the system architecture, especially in the networking part. There is a need to provide firewalls and proxies between the zones that monitor the traffic and check if the rules defined for data exchange between these zones are fulfilled.

3 Implementation

This section provides the details of the implementation of the individual security related modules on the different management units in the e-balance system. In all the subsections we describe the general implementation of the respective module that is relevant to all the management units (CMU, LVGMU, MVGMU and TLGMU). If it is not the case, then the subsection describes the exceptions, i.e., names the specific devices that are using the specific implementation.

Table 2 lists the security features supported by the individual modules relates to privacy and security in the e-balance system.

Table 2: The security features provided by the e-balance modules

Security feature / Module	Node maintenance, protection, parameterisation and reset	Data access control	Secure storage	Secure group management	Network stack security
Data confidentiality		٧	٧		٧
Data integrity		٧	٧		٧
Data availability	٧				٧
Non-repudiation		٧		٧	
Authentication		٧		٧	٧
Execution process protection	٧				
Privacy		٧	٧		٧

3.1 Node maintenance, protection, parameterisation and reset

This security related module is in fact a set of measures that provide the required features for monitoring, maintaining and protecting the individual devices in the system. Its focus is rather local to the specific machine and networking is limited to remote performing of actions that can also be taken locally – remote status monitoring, management and update.

The management units (CMU, LVGMU, MVGMU and TLGMU) work under the control of Linux system, e.g., our test implementation uses Ubuntu 14.04.3 LTS for all the management units. Thus, the main part of the node maintenance, protection, parameterisation and reset is realized using the Linux tools.

The low level maintenance and management is mainly realised based on the system users rights defined in the operating system and using the available tools, e.g., the bash console. In our test implementation the access to the maintenance functions is ruled by the Linux access control and management based on the system user definitions.

Different system users are defined for different purposes, the e-balance processes (middleware and services, respectively) run using the proper system user settings to provide proper system access rights. This approach provides the protection against malicious code execution. But it also relies on proper system configuration and maintenance – careful choice of installed software to avoid exploits to gain other user rights and execute code on their behalf, e.g., root. Further protection may be provided by the virtual machines related to programming languages, e.g., our implementation uses the Java virtual machine and the Mono (C#) virtual machine. These allow providing even more detailed configurations of the sandboxes the processes are working within. This approach allows protecting the system against potentially malicious service code, but it also helps to protect the services against each other, each running in a separate sandbox. The level of isolation between the individual processes can be different, starting with the basic protection of context

(process code and allocated memory) and going further to protecting individual processes to communicate with each other and the outside world.

Different levels of virtualisation are possible. The management units above the CMU level can also be realized as systems running on servers in virtual machines. This approach allows configuring the parameters of the sandboxes on an even higher level. In case of malfunction or failure due to malicious code execution such a system can easily be recovered. This approach allows also distributing the system components onto several virtual machines connected via local networking means and processing requests. Separating data storage (the communication platform middleware) and processed code (the energy management services) and letting them to communicate over the data interface enables even finer protection and recovery schemes.

The operating system tools for managing system packages are used for updating and parameterisation of the installed management unit system and software modules. For the Debian Linux based systems (Debian, Ubuntu) this task is done using the Advanced Packaging Tool (APT). The system components are provided in form of deb files, what helps to manage and update the software of the management units. The parameterisation and reset of the management units is also done this way.

3.2 Data access control

The data access control module is part of the request processor in the communication platform middleware. This specific module is responsible for allowing or denying the accesses to the data stored in the communication platform middleware. The requests for the request processor can be issued from external sources (from the middleware point of view) using the data interface, but also internally within the middleware (e.g., data accesses between different management units). The data access control takes care that the source of the access request is authorized to access the data. The data access control module uses the access rules stored local on the management unit in its private database tables to decide if the accesses are authorized. These access rules are defined by the security and privacy policies of the data owners (see deliverable D5.4), as well as by the process dependent access rules for the middleware internal accesses.

The authorization is based on a verification of a stakeholder certificate. This certificate is installed together with the service software and is accessible to that software only. Compared to the hardware certificates, the stakeholders' certificates prove the identity of the specific stakeholder. The services running on a management unit can work on behalf of different stakeholders, thus it is necessary that they can also authenticate themselves, while accessing the data stored in the e-balance middleware. The only exception here is the CMU, where the hardware certificate is also the stakeholder (customer) certificate, since on that particular management unit only this one stakeholder can run services. The other stakeholders have to ask the certification authority to generate certificates for them to be included in their privacy and security modules in the energy management platform.

The details of the implementation of this module are provided in deliverable D5.4. This description was provided there as a counterpart of the privacy and security module for the energy management platform. Beyond the details provided there, the module has also been further ported to C# to be integrated within the communication platform middleware.

3.3 Secure storage

The database is very often the prime target for the application level attacks. These attacks can lead to exposure of the database structure and, even more critically, the database content. They are mainly based on exploiting vulnerabilities due to exposed data storage structure (e.g., data stored in unencrypted and freely available files), due to vulnerable data access code (e.g., possible data leaks at the database interface) or due to application level vulnerabilities. We proposed means to eliminate the latter by providing the data access control module (as described in previous section) that allows only a given set of operations on the database – the data interface operations. But, since the database is an external software module provided by a third party, proper database features and configuration are necessary in order to protect the database against the first two kinds of attack, targeting directly at the database.

The database shall be configured to provide the following features:

- The storage protection – content of the tables related to the e-balance middleware shall be stored in the database (on the hard disk) in an encrypted form.

- The access protection – the access to the tables related to the e-balance middleware is possible only for the middleware – the access is protected by a password.
- The interface protection – the communication between the e-balance middleware and the database shall be encrypted to protect the data exchange – Structured Query Language (SQL) strings.

In our implementation we are using SQLite [4] as the database. The code of the database that is accessing the permanently stored data is compiled as a library and interacts with the e-balance middleware based on function calls. Thus, the eavesdropping of the communication between the middleware and the database would require monitoring of the memory (stack, heap) of the management unit working on Linux.

The permanent storage of the SQLite database can be encrypted using the SQLite Encryption Extension [5]. In this approach the database files are encrypted using the provided encryption key, supporting AES encryption – Output Feedback (OFB) mode with AES-128 or AES-256 and AES-128 in CCM mode (Counter with CBC-MAC). Thus, the possession of the database files will not reveal the content or the structure of the database. The plain data is only present in the memory of the management unit, but again accessing this temporal data would require monitoring the memory of the device running Linux operating system.

Depending on the implementation details of the specific deployments the secure database solutions can be different. But they should provide the above mentioned features. For our C# middleware implementation we have chosen the above mentioned software blocks for the prototypes, but other combinations of state-of-the-art solutions are possible as well.

3.4 Secure group management

The secure group management module takes care about creating and managing the group of trusted devices the management unit communicates with. The approach proposed in e-balance is a three-step approach, where the first two are performed in the initialisation phase and the third one is performed at run-time. These steps are sketched in Figure 4, i.e., to increase the readability of the figure only the most important data exchange by the messages is named at the arrows representing the communication. The secure group management module is implemented in C#.

Each management unit is identified by a unique identifier that is further confirmed by a certificate, signed by a trusted party – the certification authority. Our implementation of the certification authority is realized in Java and uses the X.509 certificates with ECC-256 signature using the Elliptic Curve Digital Signature Algorithm (ECDSA) with Secure Hash Algorithm SHA-1 hash function.

In order to obtain its certificate, in the first step, during the installation of the management unit, the unit connects to the certification authority and the certificate of this respective management unit is generated and signed. This step is done either by authorized personnel or using a single use credential (e.g., token), generated for this specific customer – both these authorisation methods are supported. As a result of this step the management unit obtains the certificate that confirms its hardware identity in the e-balance system and it also obtains the certificate of the trusted party, to be able to verify the certificates of its future communication partners.

In the second step, also in the configuration phase, the management unit is registered at its parent management unit in the e-balance hierarchy. Thus, a CMU is being registered at the respective LVGMU, a LVGMU at the respective MVGMU, etc. During this step, the two management units exchange their certificates, verify them (check the certification authority signature) and store them in their local secure certificate storage. In case there are redundant parent management units in the hierarchy, the registration is done at all of them and their certificates and identities are stored together with the priority number. The identity of the corresponding parent management unit is either provided by authorized person during this configuration step or it is generated automatically, based on a single use credential, generated for the specific customer. Both these authorization methods are supported by the implementation.

The certification authority can also be involved in the second step, to provide additional level of protection against attacks, e.g., certificate forging and to check if the used certificate is not present on a black list. The trusted lists of each management unit are also periodically checked at the certification authority, to identify potential treats. Additionally, during each communication session the identity of the communicating parties is verified.

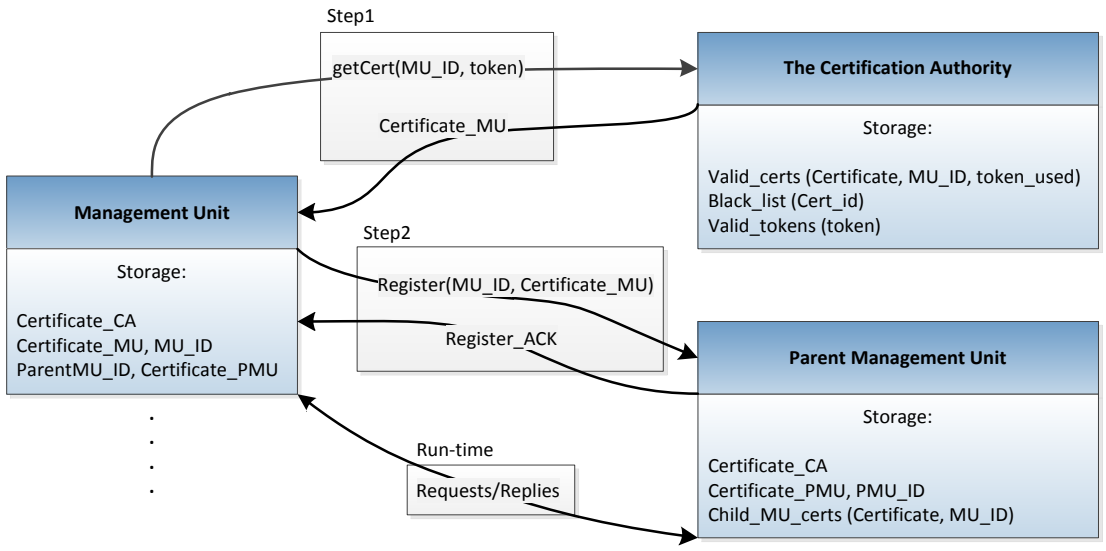

Figure 4: The secure group management details

In case the management units – the current one (the child) and its parent, are located in different zones in the zonal approach, then the sketched group management approach is slightly modified to incorporate an additional networking component (buffering proxy) that is buffering the requests, in order to satisfy the defined communication rules, e.g., that the requests can be only issued from a zone with a lower number. Figure 5 presents the first two steps, i.e., the initialisation phase in this case. Please note that, compared to Figure 4 Step2 is now a decoupled set of messages, because depending on the polling interval of the parent management unit, the request and the acknowledgement can be realized in two distinct communication sessions.

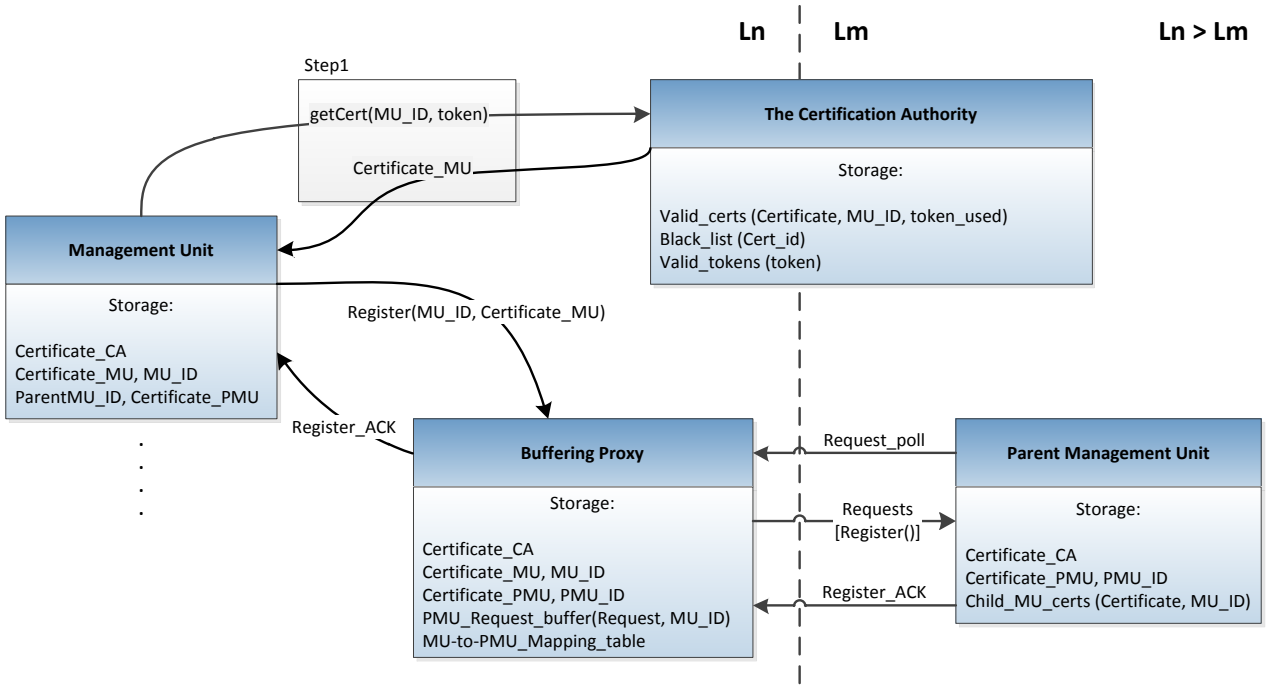

Figure 5: The secure group management approach with a buffering proxy

In this case, where the two communicating management units are located in different zones, the buffering proxy is located in the zone with the higher number (where the child management units are located) and the child management unit connect to this server, instead of connecting directly to the parent management unit. The parent management unit is located in the zone with lower number and it polls the buffering proxy for requests originating from the higher zone. Thus, all requests from the child management unit are stored at the buffering proxy. Further, the flexible implementation of the proxy supports also bidirectional buffering, i.e., the buffering proxy stores the requests targeting both interacting kinds of management units and supports

polling from both sides, providing an universal implementation of a flexible connector between the zones in the zonal approach. It may then also be located in between the two zones, i.e., providing the buffering of communication requests corresponding to the defined access rules.

In the modified approach with zones, the certification authority has to be present in each zone, allowing the registration and certificate list verification (black list check). Thus, the first step remains unchanged.

In the zonal approach the second step is realized via the buffering proxy. The proxy stores the registration request and waits until the parent management unit polls for its request buffer. The registration process is stalled until the parent replies to the buffering proxy. The buffering proxy stores the mapping between the parent and child management units, e.g., if connecting CMUs and LVGMUs, the buffering proxy stores the list of CMUs for each LVGMU, together with their identities and certificates, in order to be able to verify their identity. For each LVGMU it also stores the identity and certificate to be able to verify the identity of the LVGMUs as well. The management units store also the certificate of the buffering proxy.

The requests during the runtime are buffered in the respective buffer for the target management unit that can poll for its content. The communication is further protected by the networking means, supported by the use of the certificates of the communicating management units (and the buffering proxy).

The implementation of the buffering proxy is realized in Java programming language.

3.5 Network stack security

In the networking stack several mechanisms are used to protect the exchanged data against eavesdropping and modifications.

The communication between the management units (and the buffering proxies) is protected by means of Transport Layer Security (TLS). The certificates of the individual management units (and buffering proxies) are used to generate session keys to encrypt the communication.

Communication redundancy is supported and can be used whenever possible. In order to do so the communication manager module in the middleware can use different networking stacks (and communication technologies) from the available ones to increase the resistance to denial of service attacks targeting the specific communication channel. The CMU per default uses the WiFi Local Area Network connection to connect to the LVGMU (via Internet) , but it can use the Power-line Communication (PLC) via the smart meter (SM), or it can be equipped with a 3G USB dongle and can use the alternative communication technologies as backup, if the primary channel fails. The implementation details of the communication manager module are provided in Deliverable D4.3 [2].Firewalls are used on the management units to filter the communication to limit potentially malicious requests. In our approach we are using filtering of messages based on the iptables [6] utility for configuring the firewall, available in the Linux operating system. Depending on the system provider requirements the firewall can be configured to accept only known messages from known sources. From the security point of view, the optimum configuration is to allow only outgoing messages and replies to these. This configuration is especially practical with the combination with the buffering proxies, configured in a way that the requests are buffered for all management units and all these management units have to poll for the requests targeting them – initiate each communication with a defined structure of the content.

The remote access to the maintenance and management tools is provided using the Secure Shell (SSH). We use the implementation provided by the OpenSSH [7] suite.

The above mentioned network stack security mechanisms apply to all the management units. In the area of sensor networks the security is based on the specific applied protocols. In our implementations we use ZigBee for sensor networks. The ZigBee Pro protocol supports data protection on two levels – application level security allows protecting the data exchanged between two distinct nodes using a key shared only by these two nodes; while the network level security allows for a network wide data protection involving all the nodes in the network, i.e., the data is protected using a key shared by all the nodes in the network. The protection in this case provides data confidentiality and integrity – the data is encrypted using AES-128.

The sensor network is connected to the management unit by means of a gateway that is a node in the sensor network, but also provides an interface for the management unit to interact with the sensor network, e.g., to read/write data or send commands.

4 Conclusions

This document provides the implementation details of the security and privacy solution for the e-balance smart grid system. It provides a holistic overview of the system that follows our considerations related to the security requirements.

The implementation uses the available solutions and libraries whenever possible. We have also implemented components that did not have available commercial or open source implementations available or required special functions not provided by available solutions. This includes the homomorphic encryption, the certification authority, the buffering proxy as well as the security and privacy related modules of the communication platform.

The proposed solution as a combination is based on some standard cryptography approaches. We have named the crypto base, what allows defining the current strength of the applied protection. However, the applied cryptography base can be considered as a library that can be exchanged or updated in future deployments to keep the system security at a level that is up-to-date and conform to the current security level requirements.

Additionally, the chosen solutions from the set of available ones are also not obligatory. Alternative choices are possible, but as long as the alternatives provide the needed features and work based on the same or similar interfaces. The latter makes the use of standard solutions an advantage.

References

[1] K. Piotrowski, et al., "Deliverable D3.2 – Detailed System Architecture Specification", Public deliverable of e-balance project, FP7-Smartcities-2013, Project number 609132, 2015.

[2] D. Garrido, et al., "Deliverable D4.3 – Detailed middleware specification and implementation", Public deliverable of e-balance project, FP7-Smartcities-2013, Project number 609132, 2015.

[3] K. Piotrowski, et al., "Deliverable D5.4 – Detailed specification, implementation and evaluation of security and privacy means", Public deliverable of e-balance project, FP7-Smartcities-2013, Project number 609132, 2015.

[4] SQLite Home Page, https://www.sqlite.org/index.html, last viewed: 07.02.2016.

[5] SQLite Encryption Extension Home Page, http://www.sqlite.org/see, last viewed: 07.02.2016.

[6] IPTABLES manual page, http://ipset.netfilter.org/iptables.man.html, last viewed: 07.02.2016.

[7] OpenSSH Home Page, http://www.openssh.com, last viewed: 07.02.2016.

