A framework for user centred privacy and security in the cloud

Standardisation Requirements

Type (distribution level)	*Public*
Contractual date of Delivery	*30-06-2015*
Actual date of delivery	*29-07-2015*
Deliverable number	***D2.5***
Deliverable name	***Standardisation Requirements***
Version	*V1.0*
Number of pages	*33*
WP/Task related to the deliverable	*WP2/T2.4*
WP/Task responsible	*Trust-IT*
Author(s)	*Roberto Cascella (Trust-IT)*
Partner(s) Contributing	
Document ID	*CLARUS-D2.5-StandardisationRequirements-v1.0*
Abstract	This deliverable reports on the necessary requirements results and implementation roadmap, effort and analysis, considering the EU and global landscape.

Disclaimer

Revision History

Version	Date	Author	Description
V0.1	13/05/2015	Roberto Cascella (Trust-IT)	ToC
V0.2	28/05/2015	Roberto Cascella (Trust-IT)	Content in various sections. Inputs from partners about standards using the questionnaire
V0.3	23/07/2015	Roberto Cascella (Trust-IT)	Deliverable ready for internal review
V0.4	26/07/2015	Roberto Cascella (Trust-IT)	Integrated comments from Sven Rosinger
V1.0	27/07/2015	Roberto Cascella (Trust-IT)	Integrated comments from Antonio M. Ortiz and finalised the mapping of technical requirements from D2.2.

Reviewers:

Sven Rosinger (OFFIS)

Antonio M. Ortiz (Montimage)

Executive summary

Trust is essential to wider uptake of cloud services but it can only come from solid mechanisms that ensure greater control over the security and privacy of the user's data. Confidentiality and privacy are still major concerns when it comes to moving to the cloud. Many organisations are also reluctant to outsource sensitive data due to lack of control over its storage and management. What's more, it is also increasingly important to protect business assets from vulnerabilities or attacks and ensure that applications continue to operate and provide a good level of service even during an attack.

CLARUS is all about improving trust in cloud computing and securely unlocking sensitive data to enable new and better cloud services. CLARUS is developing a secure framework for storing and processing data outsourced to the cloud so end-users can monitor, audit and control their stored data while gaining the cost-saving benefits and capacity that cloud services bring.

CLARUS is paving the way towards more transparent, standardised, auditable and controllable cloud services, benefitting both consumers and providers of cloud services.

This deliverable reviews the current standard landscape for cloud computing, security, and data format relevant for CLARUS. The objective of this document is to define the standardisation roadmap built on a throughout analysis of the standards and their potential adoption in the design and implementation of the CLARUS proxy solution. By considering relevant standards in the early stage of the project will ease the integration of standardised solutions and follow an *interoperability by design* approach.

In the context of CLARUS, several aspects related to interoperability are considered: the Application Programming Interfaces (APIs) of the components and tools, the security solutions and the format of data. The objective is that the CLARUS solution will be able to interoperate and work without requiring further development with existing or future systems.

This document first recaps the cloud computing ecosystem and the objective of CLARUS to highlight the need for standards. Then, it discusses the CLARUS approach to standards and the methodology used to identify relevant standards; all partners have contributed to the selection of the most suitable standard working groups for CLARUS architecture design and implementation. Special attention has been given to the standards that could drive the design and implementation of the CLARUS architecture and of the two business use cases.

After reviewing the major standards in cloud computing, security & privacy, and data format, this document focuses on additional requirements stemming from the best practices and recommendations of Standard Developing Organizations (SDOs). Finally, this deliverable maps the technical requirements identified in Deliverable D2.2 "Requirements specification V1" to the identified standards.

The standardization effort in CLARUS is a joint activity with other work packages meant to analyse the standard landscape, orchestrate the co-operation with international organization (D7.1 "Dissemination and standards report V1"), and to derive new case studies and recommendations about standardization issues from the design of the CLARUS architecture (D4.3 "Standardisation and interoperability"), comprising the assessment of the standardization guidelines and requirements (D6.4 "Standardisation assessment") identified in this document.

The requirements identified at the early stage of the project will be further analysed and refined in D2.3 "Requirements Specification V2".

Table of Contents

List of Figures

1 Introduction

The objective of CLARUS is to enhance trust in cloud computing services by developing a secure framework for the storage and processing of data outsourced to the cloud. To have a wider acceptance of the technical solutions, CLARUS will implement a set of standardized cloud services, and therefore, transparent with regard to data management, privacy and security. This model change will give control back to data owners and will also encourage cloud services based on standards that can be certified as compliant with security and privacy.

The foreseen solution is a proxy installed in the end-user trusted domain, able to communicate and to interoperate with public cloud providers via the use of standardized communication interfaces and protocols:

- Between the end user and the CLARUS proxy;
- Among different CLARUS instances running at different organisations;
- Between the CLARUS proxy and the cloud service providers.

By means of standardisation, protocols and functions can be made homogenous for cloud providers and CLARUS proxies, so that:

- Interoperability can be achieved among otherwise heterogeneous cloud providers;
- Collaborative services (e.g., edition of documents outsourced to the cloud by several users) can be implemented through several CLARUS proxies;
- APIs based on standards can be made available to programmers for a seamless development of end-user cloud-based applications.

Standardization is of major interest for CLARUS. In order to guarantee that the use of standards or de-facto standards is enforced for the architectural design of the proxy solution and privacy-protecting tools, it is necessary to monitor the selected standards working groups and create synergies.

The two principal reasons for using standards in the development phases of CLARUS (apart from the obvious desire to produce an open source piece of software) are the leverage that CLARUS can create out of the project and receive into the project. Firstly, using well-known leading standards will facilitate the reuse of existing software in the project that is robust and well supported. Secondly, it also encourages others from outside the consortium to use the tools developed within the project and integrate them in their systems. This considerably increases the impact created by the project.

1.1 Scope of the document

This deliverable defines the set of requirements stemming from the best practices and recommendation of Standard Developing Organizations (SDOs) and then maps the requirements identified in D2.2 “Requirements specification V1” [3] to available standards. The objective is to guide the design and use of standardized interfaces for the development of CLARUS services.

The standardization effort in CLARUS is a joint activity with other work packages meant to analyse the standard landscape, orchestrate the co-operation with international organization (D7.1 “Dissemination and standards report V1” [4]), and to derive new case studies and recommendations about standardization issues from the design of the CLARUS architecture

(D4.3 “Standardisation and interoperability”), comprising the assessment of the standardization guidelines and requirements (D6.4 “Standardisation assessment”) identified in this document.

This deliverable is linked with D2.2 “Requirements specification V1” [3], which lists the case studies originated from CLARUS and its two business cases, and with D2.4 “Legal and ethical requirements” [54] that derives requirements from the analysis of the privacy legislation framework applicable to cloud services.

1.2 Acronyms

API	Application Programming Interface
CAMP	Cloud Application Management for Platforms
CDMI	Cloud Data Management Interface
CIMI	Cloud Infrastructure Management Interface
CPIP	Cloud Portability and Interoperability Profiles
CSA	Cloud Security Alliance
CSCC	Cloud Standards Customer Council
CSP	Cloud Service Provider
DMTF	Distributed Management Task Force
CSW	Catalog Service for the Web
DPCO	SNIA Data Protection and Capacity Optimization
EGI	European Grid Infrastructure
EHR	Electronic Health Record
EMR	Electronic Medical Record
HL7	Health Level Seven
IaaS	Infrastructure-as-a-Service
ICT	Information and Communication Technology
IEC	International Electrotechnical Commission
IETF	Internet Engineering Task Force
ISO	International Organization for Standardization
JSON	JavaScript Object Notation
LDAP	Lightweight Directory Access Protocol
NIST	National Institute of Standards and Technology
OASIS	Organization for the Advancement of Structured Information Standards
OCCI	Open Cloud Computing Interface
OGC	Open Geospatial Consortium
OGF	Open Grid Forum
OVF	Open Virtualization Format
PaaS	Platform-as-a-Service
PbD-SE	Privacy by Design Documentation for Software Engineers
PII	Personally Identifiable Information
PKI	Public Key Infrastructure
PMI	Privileged Management Infrastructure
PMRM	Privacy Management Reference Model
SaaS	Software-as-a-Service
SAML	Security Authorization Markup Language
SDO	Standard Developing Organizations
SIIF	Standard for Intercloud Interoperability and Federation
SNIA	Storage Networking Industry Association

SOA	Service Oriented Architecture
SOAP	Simple Object Access protocol
SSO	Single Sign-On
TOSCA	Topology and Orchestration Specification for Cloud Applications
UDDI	Universal Description, Discovery and Integration
W3C	World Wide Web Consortium
WFS	Web Feature Service
WMS	Web Map Service
WSDL	Web Services Description Language
XACML	eXtensible Access Control Markup Language

2 CLARUS and the Cloud computing ecosystem

2.1 Cloud computing ecosystem

Cloud computing has gained its momentum leveraging the pay-as-you-go concept and the idea of using commodity public cloud providers based on customers needs without requiring an in-house dedicated infrastructure. The National Institute of Standards and Technology (NIST) has classified the business models for cloud computing, based on the type of services offered to the customers, and identified the type of stakeholders of the cloud models.

In this section we briefly recap the concepts and discuss cloud computing as an interdisciplinary approach leveraging technologies that span from virtualization, networking, data, programming models, web services, and security. Hence, open standards related to these technologies are all under consideration for cloud computing at a different level of granularity.

NIST [5] identifies three different business models for cloud computing, each with its own characteristics and impact on the implementation of security and privacy enabling functionalities.

Infrastructure-as-a-Service (IaaS): the provider offers a pool of resources in the form of virtual machines running on a hypervisor, file systems that can be loaded and mounted, and networking capabilities to interconnect the virtual machines. Customers can run any type of operating system and software on the virtual machines to deploy their applications and have full control of the processes running, thus have the responsibility to maintain their software and applications. Privacy and security features are implemented by the users, who must manage the keys for authentication, integrity and confidentiality, and ensure that the communications between virtual machines are end-to-end encrypted.

Platform-as-a-Service (PaaS): the provider offers computing platforms and programming frameworks for the deployment of the applications. Customers can upload their software solutions and manage their applications without the need of managing the underlying infrastructure. The only responsibility of the customers is the code running in the cloud, which performs a given task. Security and privacy features are implemented by the provider and customer. Customers are in charge of securing their applications and can rely on tools offered by the cloud infrastructure.

Software-as-a-Service (SaaS): the provider offers services accessible via simple interfaces, and manages the cloud infrastructure running the applications. Customers access directly the applications without the need to deploy the code to implement a given task, and do not have to install any specialized software to run the application. Security and privacy are all implemented by the service provider, who must provide adequate authentication mechanisms to users who access the services.

The NIST cloud computing reference architecture defines five major actors: cloud consumer, cloud provider, cloud carrier, cloud auditor, and cloud broker. Relevant stakeholders of the CLARUS solution are essentially the *cloud provider* and the *cloud consumer*. The provider offers services to store, manage, and process data of CLARUS end-users in the cloud. The cloud consumer can have different roles based on the use case considered. We refer to the following section for a characterization of the actors in CLARUS and to deliverable D2.2 for a detailed analysis of the interactions between the different actors in CLARUS.

2.2 CLARUS motivations and background

Cloud computing offers many benefits to the users with a vast selection of services ranging from bare infrastructure resources, advanced programming models to process data, to services accessible via simple browser interfaces. Despite the potential of cloud computing, trust is essential to wider uptake of cloud services but it can only come from solid mechanisms that ensure greater control over the security of this multi tenant infrastructure and the privacy and confidentiality of data.

Customers rely on cloud providers for their daily and business activities, assuming that their data are safe. Current security mechanisms are commonly located within the cloud platform, hence compelling customers trust cloud providers for the way they manage their data. This leaves the cloud as an impractical solution for those customers that value the sensitivity of their data as critical or should comply with specific regulations that force them to treat data with special precautions, reaching higher importance when dealing with business and user sensitive data. Thus, to reach its full potential, cloud computing needs solid security mechanisms that enhance trust in cloud computing by allowing cloud customers greater control on the security and privacy of their data.

This is the main objective of the CLARUS project: enhance trust in cloud computing services by developing a secure framework for the storage and processing of data outsourced to the cloud.

CLARUS ultimate goal is to allow end users to retain control over their personal data: monitor, audit, and control the stored data without impairing the functionality and cost-saving benefits of cloud services. This model change will give control back to data owners and will also encourage cloud services based on standards that can be certified as compliant with security and privacy. The foreseen solution is a proxy installed in the end-user trusted domain, able to communicate and to interoperate with public cloud providers.

In the context of CLARUS, the relevant cloud computing business models are PaaS and SaaS. CLARUS will develop a set of security tools and services that will be integrated in the cloud provider infrastructure to enable storing encrypted data processing user's data in a secure way. The aim is to improve the security capabilities of the cloud providers in a transparent way for the user who will benefit from an enhanced framework able to preserve the privacy. Customers can rely their data on multiple providers at the same time with potentially different level of confidence, e.g. storing encrypted sensitive data or splitting anonymised data across providers in such a way that the privacy of the user is still preserved. Hence, interoperability for data formats and interfaces of cloud services is the key to ensure compatibility between independent systems.

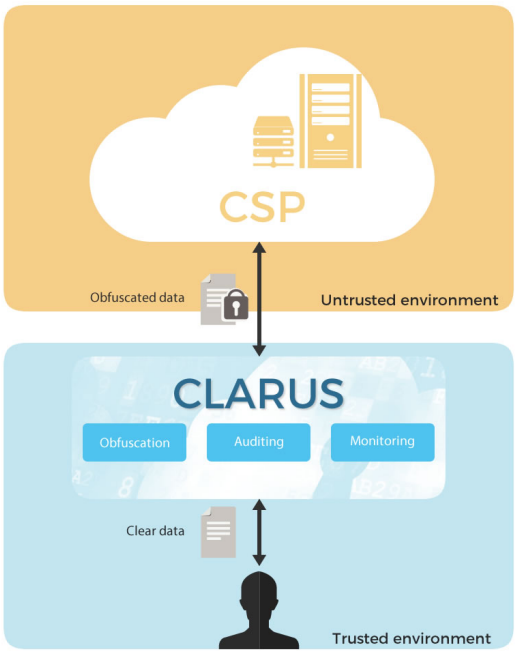

Figure 1 - CLARUS proxy solution

Figure 1 depicts the CLARUS proxy solution. The cloud service provider (CSP) is considered untrusted, honest but curious. The provider

may access the stored information for monitoring purposes or simply to provide the required services to its customers, such as processing data or returning results of queries to data stored in a database. This is the main justification of the need of the CLARUS solution. Customers reside in a trusted domain and access the CSP services via the CLARUS proxy, which encrypts or anonymise the data as required before accessing CSP.

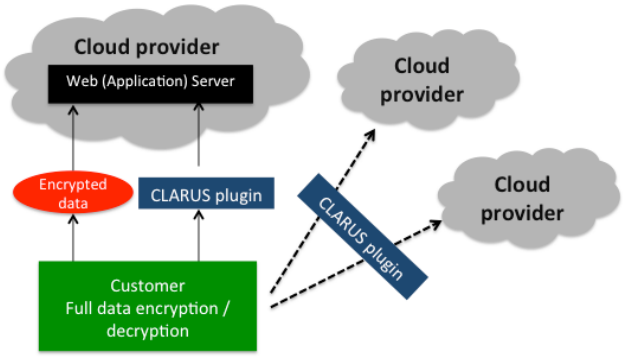

Figure 2 - CLARUS multi cloud scenario

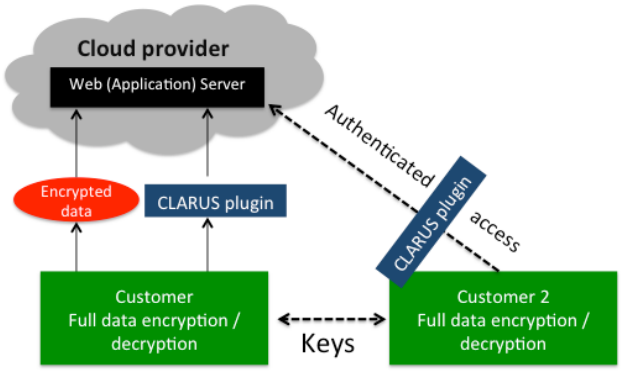

Figure 3 - CLARUS multi user scenario

Figure 2 depicts a multi cloud scenario, with the client accessing services offered by other providers by means of the CLARUS proxy. Figure 3 depicts a user (customer 2) who accesses the data of another customers via the CLARUS proxy after requesting authorization and retrieving the keys to decrypt the data.

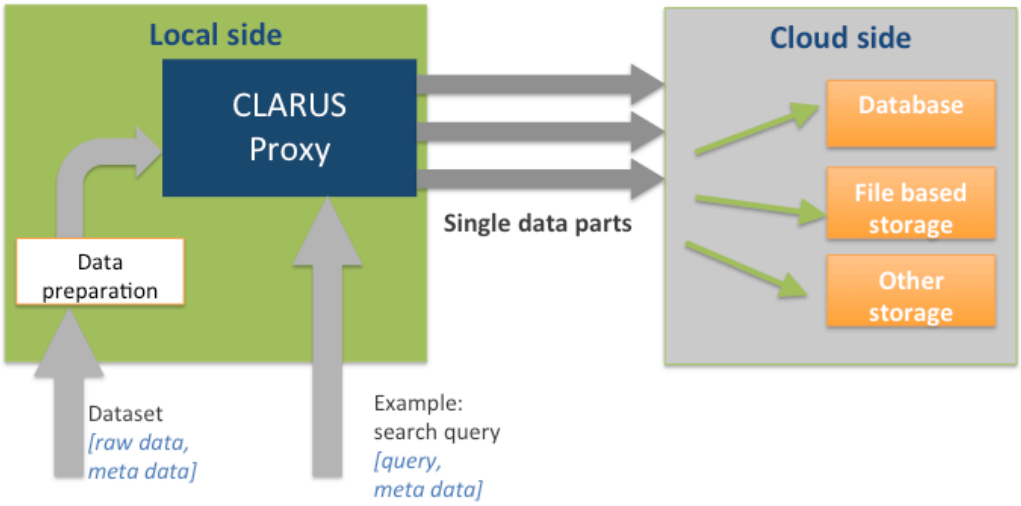

Figure 4 - Data splitting across different cloud services

Data can be split and stored on different components, as depicted in Figure 4 - Data splitting across different cloud services. This operation is performed since the users' data might require different levels of protection (encrypted, anonymised, or plaintext) as well as the meta-data that is used to enable a search in encrypted data and to reconstruct the original information.

These scenarios justify the need of a standardized solution for the service interfaces and data format since the CLARUS solution should be applicable across different domains, i.e. cloud service providers, and serve multiple customers, or support different storage type.

CLARUS will be demonstrated in two use cases: the publication of geo-reference data on the Internet and the eHealth scenario, described in deliverable D2.1. In use case 1, actors in the domain own and manage sensible information that is available to the public domain or it's

confidential, as such must be kept private. The functionalities will be provided via a set of security tools that need to work on anonymised and encrypted data that are exploited commercially by private companies, analysts or by other institutions in the public section. Geo-data information in the environmental domain possesses interesting characteristics like the size of the available data and the existing metadata descriptions (mostly according to the European Directive INSPIRE [6]) that require to be considered while applying the CLARUS solution. The second case study is related to eHealth and concerns a distributed e-health scenario that requires immediate access to medical data outsourced to cloud providers. The main actor in this use case is the hospital that is responsible to treat the Electronic Medical Records (EMRs) of the patients, which contain information that is highly identifying or confidential. A series of functionalities are needed, like creating, managing and updating medical histories, including results of clinical visits, searching for specific patients/histories and also a shared and cooperative access to these data according to access policies. The CLARUS solution will be applied to obfuscated data to provide security-aware access to functionalities, which should be backed up by a series of auditing tools.

In the remaining sections of this document, the scenarios herein described will be analysed to first determine the relevant open standards. Then, the requirements will be derived from best practices and recommendations of Standards Developing Organizations (SDOs) to guarantee the design and implementation of an interoperable solution thanks to the early adoption and support of standards. The threat model and a detailed analysis of the CLARUS use cases to derive the requirements to protect the system and guarantee security and privacy are part of Deliverable D2.2 [3].

3 The CLARUS approach to standards

3.1 Relevance of standards

Standardisation is of major interest for CLARUS, being the project committed to design and develop a standardised solution and to contribute to the European standard landscape with use cases and guidelines to fill in the current gap in standards for security and privacy in cloud services. Hence, standards have been monitored to identify those that are relevant to CLARUS to guide the architectural design and the development of CLARUS services implementing standardised interfaces and solutions.

One of the objectives of CLARUS is ***interoperability***, which ensures that the actors and systems providing cloud computing services and implementing the CLARUS solution exchange and make use of the information. Moreover, interoperability will reduce the risk of vendor lock-in for customers, which is one of the major obstacles to the adoption of cloud computing along with legal issues related to data protection and security, all aspects addressed by the CLARUS solution.

In the context of CLARUS, several aspects related to interoperability are considered so that the CLARUS solution is able to interoperate and work without requiring further development with existing or future systems. Interoperability is considered in different domains.

Application Programming Interfaces (APIs) of the components and tools developed as part of the CLARUS solution. These components need to be integrated in existing cloud infrastructures and easily exploitable by external components via the same APIs. Interoperability achieved via standardised APIs refers to the cloud computing domain and components of the CLARUS solution that need to interoperate among each other and with the existing cloud infrastructure to facilitate the integration with and the reuse of existing tools, and ease the penetration of the technology to a large set of applications and cloud infrastructures. In the CLARUS context, interoperability will be considered for the design of the architecture, contextualized to the cloud business models identified in Section 2.1, and for the implementation of the software solutions.

Security solutions to secure the communications, authenticate the users, and provide access to CLARUS cloud services. The intent is to provide compatibility with well-known protocols and best practices to ensure that the CLARUS solution can leverage existing authentication frameworks to facilitate the authentication of the users while at the same time ensure that the secure exchange of information between users, services, and components is done in a standardised way. Interoperability for security is needed to enable cooperative access to data from different CLARUS platforms, and according to policies set by the data owner.

Format of data processed by web services. The data and metadata stored and processed by CLARUS need to follow predetermined standards to allow the application of the CLARUS solution to a large set of applications. In the context of the two CLARUS use cases, the data format must follow well-defined standards, as documented in Deliverable D2.1: the Open Geospatial Consortium standards for the publication of geo-referenced data on the Internet, and the Health Level-7 standard for the transfer and format of healthcare information. Other standards can be considered for data format.

Interoperability demands common technical APIs, protocols and data/message format, which can be achieved by following best practices and common guidelines or in its more general form,

i.e., by design, adopting open or de-facto standards. Since the beginning of the project, CLARUS follows an ***interoperability by design*** approach by investigating the use of open standards in the architecture design and later in the implementation of the CLARUS components. The objective is to implement standards supported by a wide range of Cloud Service Providers (CSPs) and end users, thereby ensuring interoperability in collaborative, standardised and transparent cloud environments. By means of standardisation, the function calls implemented in the interfaces can be made homogeneous for cloud providers that provide similar services (e.g. data storage), so that data interoperability can be achieved among otherwise heterogeneous cloud providers. On the other hand, standards will allow the support of data splitting (for security enhancement, like to meet privacy constraints), merging and replication (for improved data integrity in front of potential CSPs' failures), thus facilitating the adoption of already available distributed backup solutions as such the integration of the CLARUS solution into the existing cloud infrastructure. Interoperability brings several benefits to CLARUS such as the possibility of implementing more robust security mechanisms and improving reliability and dependability, while increasing transparency and trust in cloud services.

3.2 Methodology in choosing standards

The landscape of standards in cloud computing is quite vast since there is not a single standard organization as reference due to the large number of communities that are involved in the design of cloud computing services. Cloud computing is quite a complex domain with institutions looking and contributing at different standard organizations for both business and technical requirements. In addition, many different areas contribute to make cloud computing successful such as networking, web services, security, and resource management.

The need to support and mainly use standards in the project is coming from promoting a new technology, like CLARUS, which leverages the solutions widely accepted by the research and industrial community. On the one hand, this could improve the CLARUS technology itself, either by reusing code available or developing code following clear specifications, and on the other hand, ameliorate the exploitation opportunities since some parts or a component could be taken as reference implementation for a particular standard. Moreover, the projects aims at supporting interoperability, which could be done by adopting widely accepted standards.

The CLARUS project proposes a proxy solution that will implement specific security functionalities offered as a set of PaaS and SaaS services. For completeness we will analyse also the standards relevant for the management of the resources at the infrastructure level, security and data format.

The process to determine to the most appropriate standards for the CLARUS technology has started by asking all partners of the project to identify the relevant standards for the project. Special attention has been taken with respect to the design and implementation of the CLARUS architecture and to the two business use cases. The responses to a questionnaire have been further elaborated to come up with a list of standards that covers at large all aspects of CLARUS. After a feasibility analysis, some of them might not be used for the final implementation, but it is still interesting to monitor their evolution.

4 Standard organizations landscape

The standards landscape for cloud computing is quite vast and standards, which enable interoperability and are applicable to cloud applications, data, or security methods, already exists. In this section we analyse the relevant standards that enable interoperability at different levels: management of cloud resources and applications, data format, and security. We first identify the major Standard Development Organizations (SDOs) and then we briefly discuss the standards that could be relevant for the design of the CLARUS architecture and implementation of the proxy solution.

4.1 Cloud computing standards for interoperability

The CLARUS project proposes a proxy solution that will implement specific security functionalities offered as a set of PaaS and SaaS services. For completeness we will also analyse the standards relevant for the management of the resources at the infrastructure level. The standardisation effort for model or applications at the SaaS level is still in its infancy due to the complexity of the large variety of software service solutions. In this deliverable, we consider the standards for security and data format, discussed in Section 4.2 and Section 4.3 that can be used to achieve a certain degree of interoperability for software applications.

4.1.1 Infrastructure-as-a-Service

In this section, we describe the open standards that help achieve user applications and cloud providers' interoperability for pivotal elements such as storage, infrastructure management, and application description formats. They are part of the larger effort of the standardisation community that addresses the definition of APIs for the management of cloud resources at the IaaS layer of the cloud reference model. Many organisations are involved in various standardisation activities on the common theme of clouds. Notable among them are the working groups operating within the Open Grid Forum (OGF) umbrella. Other prominent industry consortiums active in standardisation for cloud computing services are Distributed Management Task Force, Inc. (DMTF), and the Storage Networking Industry Association (SNIA).

The ***Open Grid Forum (OGF)*** [17] brings together industry and academia to devise a set of standards and specifications in several areas such as security, resource management, transfer protocols, scientific applications and so on. OGF is structured in groups (working, community, and research). Born as grid computing initiative, lately OGF develops standards and specifications for the cloud computing community, with flagship standards such as: Open Cloud Computing Interface (OCCI) for the management of cloud resources, and WS-Agreement for the negotiation of Service Level Agreements. In the context of CLARUS we only discuss OCCI below. The negotiation of SLAs is not part of the objectives of CLARUS.

- **Open Cloud Computing Interface (OCCI)** [46] is a major standard, widely supported by open source products (e.g., OpenNebula [55] and Openstack [56]) and international initiatives (such as the European Grid Infrastructure – EGI [57]). It attempts to standardise the RESTful protocol and API for tasks management and it is intended for the management of IaaS resources. The standard is quite extensible and can be used for PaaS and SaaS services as well. In the current release (version 1.1), it supports HTTP rendering and provides infrastructure extensions to deal with IaaS clouds. A new release (1.2, now available in public comment) is planned in 2015 with additional functionalities

and extensions like the support for Service Level Agreement negotiation and resource monitoring.

The ***Distributed Management Task Force (DMTF)*** [7] is an industrial organization that creates standards for the management of the infrastructure and enables interoperability among different vendor IT solutions for the IaaS systems. DMTF standards cover the definition of interfaces of systems and network management including cloud, virtualization, and storage. In the context of cloud computing and CLARUS, two noteworthy standards are the Cloud Infrastructure Management Interface (CIMI) standard for the administration of cloud resources and the Open Virtualization Format (OVF) for the packaging and deployment of virtual appliances, i.e. virtual images running on the hypervisor.

- **Cloud Infrastructure Management Interface (CIMI)** [44] is a standard from the DMTF consortium that targets management of resources within the IaaS domain. It implements a REST interface over HTTP and defines the REST APIs for both XML as well as JSON rendering. CIMI attempts to provide first-class support to Open Virtualization Format standard.
- **Open Virtualization Format (OVF)** [45] is a standard from DMTF, which aims to completely describe an application comprised of any number of virtual machines in a standard and portable format, written as an XML file.

Storage Networking Industry Association (SNIA) [8] proposes standards for storage elements and information management under the guidance of the Technical Council, composed by technical experts from the industry members. The SNIA flagship standard is the Cloud Data Management Interface (CDMI). The SNIA Data Protection and Capacity Optimization (DPCO) Committee has recently developed a product selection guide that lists products addressing backup and recovery, capacity optimized storage, and replication.

- **Cloud Data Management Interface (CDMI)** [43] defines a RESTful interface that allows cloud applications and users to retrieve and perform operations on the data from the cloud. The interface allows capability discovery of storage elements of the cloud. It also allows administrators to manage the containers, i.e., metadata, and user accounts and credentials pertaining to the cloud storage.

IEEE has established two working groups to work on cloud profile (P2301) and on Intercloud Interoperability and Federation (P2302). The P2301 working group proposes the Guide for Cloud Portability and Interoperability Profiles (CPIP) [9] to advise cloud computing practitioners of standards based choices for interfaces and file formats to enable application portability and interoperability of components from different vendors. The P2302 working group proposes the Standard for Intercloud Interoperability and Federation (SIIF) [16] to address the description of topology elements, functions, and governance for cloud-to-cloud interoperability, thus enabling the creation of a federation of cloud resources.

De-facto standards are widely used by the open source community, which has a major role in the implementation of cloud services. The major open source effort focuses on two initiatives, OpenStack and OpenNebula for the management of IaaS resources. In addition to their native interfaces, they also implement specific extensions implemented by the open source community to support OCCI standard.

4.1.2 Platform-as-a-Service

The effort to achieve interoperability at the Platform-as-a-Service (PaaS) layer is still in progress with few standards available at the moment. Interoperability at the PaaS level means having multiple components that are able to interact using well-defined messages and protocols. Compared to the IaaS model, where developers work with the bare resources, thus they can provide all the software needed to run their applications, PaaS services pose many problems linked to the portability of programming models across different platforms. Indeed, the lack of a consistent platform definition among PaaS providers is the main risk of deploying interoperable components that can be combined independently of the vendors implementing them.

The **Organization for the Advancement of Structured Information Standards** (OASIS) [10] works on standards to provide application portability and interoperability between cloud platforms. The available standards for PaaS are TOSCA and CAMP, which aim at reducing the complexity of deploying and managing cloud applications. OASIS also contributes to standards for security, privacy, access and identity policy, which will be discussed in Section 4.2.

- The **Cloud Application Management for Platforms (CAMP)** [35] standard aims to standardise the cloud PaaS management by defining APIs using REST and JSON that help packaging and deploying applications, and controlling PaaS workflows. CAMP also defines mechanisms for the monitoring and control of applications.
- The **Topology and Orchestration Specification for Cloud Applications (TOSCA)** [33] proposes an XML-based standard language to enable the interoperable description of cloud applications and infrastructure cloud services. Each cloud application is formalised as a typed topology graph and the management tasks as plans. TOSCA enhances the portability and management of software components, and facilitates the migration to any compliant cloud, but it does not execute the actual portability of the application.

De-facto standards for PaaS are emerging thanks to the effort of the open source community. Noteworthy open source projects with industrial support are Cloud Foundry [58], Heroku [59], OpenShift [60] and Docker [61].

4.2 Security & Privacy related standards

Security and privacy are the two pillars of the CLARUS proxy solution. On the one hand, the objective is to provide a secure framework for the management of the users and data, and on the other hand, it is to give the control to the users over their data. In this section, we review major standards of interest for CLARUS at the application (web services and SaaS solutions), transport, and networking layers. First, we discuss the security standards in the cloud, and then, we discuss other standards that can be applied to cloud services to address authentication, authorisation and confidentiality.

The **International Organization for Standardization (ISO)** and the **International Electrotechnical Commission (IEC)** under the joint ISO and IEC subcommittee, ISO/IEC JTC 1/SC 27, have published the 270xx ISO/IEC family standard, which is a set of internationally recognised good practices for information security management risk. These standards apply to the operations of IT systems and define the management of information security, the identification of potential risks and the implementation of security best practices to reduce these risks.

The foundation standard is the **ISO/IEC 27001** [36]. It contains the requirements for the design, maintenance, and implementation of an information security system, and it identifies the interested parties and their needs, the security risks and actions. On the other hand the best practices, or security controls, to maintain and implement the security system are part of the **ISO/IEC 27002** [37] standard. The recommendations and best practices are not specifically for cloud computing but they can be applied to cloud services; providers can certify the compliance to the standards for their cloud services. ISO/IEC is currently developing a new standard (**ISO/IEC 27017** [39]) that deals with cloud computing specifically by mapping the ISO/IEC 27002 to cloud services.

The **ISO/IEC 27018** [38] standard has been recently completed. It extends the established ISO/IEC 27002 standard to deal with the protection of Personally Identifiable Information (PII) in public clouds, which act as processors of personal data.

The ITU Telecommunication Standardization Sector (ITU-T) defines standards for telecommunications. ITU-T has produced the X.1600 [41] family standard that addresses the design security framework for cloud computing. The standard analyses the security threats and challenges in cloud computing, and provides recommendations to mitigate security risks. From these analysis and recommendations, a cloud computing customer should be able to do a risk assessment of adopting cloud computing. Major threats analysed in the standard for cloud customers are: data loss and leakage, insecure service access, and insider threats. The main security challenges for customers are: ambiguity in responsibility due to legal requirements, loss of governance and privacy by outsourcing services to cloud providers, and cloud service provider lock-in. Finally the standard ITU-T X1600 identifies recommendations to address security threats and challenges.

The **National Institute of Standards and Technology (NIST)** has defined a set of recommendation for cloud computing security in the form of report.

4.2.1 Identities management: authentication and authorisation

There exist several standards in security for authentication, authorisation, and management of user credentials.

The sensitive nature of the data managed by the CLARUS proxy requires suitable authentication mechanisms, separation of roles between data owners and data consumers, and fine grained access control when using cloud services. Some of the capabilities that CLARUS needs to support are federated IDs and single sign-on (SSO) to reduce the burden of multiple authentication steps and facilitate the management of users across trusted domains.

Federated identities (IDs) allow users to use credentials and IDs already set with a trusted identity provider to access cloud services. Thus, a federation identity consists of linking IDs of trusted providers. The IDs could be held directly by the client or by the trusted ID provider.

Single sign-on (SSO) enables a federation of identities. Users can access multiple services without being prompted for additional authentication. SSO requires trust relationships between service providers as the authentication token, issued by a provider, need to be trusted across multiple systems and domains.

Several open and de-facto standards provide Federated IDs and SSO capabilities. The ***Lightweight Directory Access Protocol (LDAP)*** [62] from IETF provides authentication and authorisation services. OASIS proposes the XML-based standard ***Security Authorization Markup Language (SAML)*** [11] for exchanging authentication and authorisation information between

identity provider and service provider. Based on SAML, the Internet2 initiative has defined the **Shibboleth** [63] web-based technology that implements identity management and federated identity-based authentication and authorization. **OpenID** [64] is an open standard of the OpenID foundation for the authentication of the users relying on identities verified and trusted by trusted third party providers. **OAuth2** is an IETF standard (RFC 6749 [12]) for authorisation; it enables the delegation of rights and permissions by creating dynamic credentials to provide a trustworthy communicating infrastructure. On top of OAuth, **OpenID Connect** provides user authentication via a simple API to verify the identity of the user and obtain basic profile information. **WS-Federation** [40] is a standard proposed by OASIS for implementing federation identity; it enables a security domain to broker for identities, identity attributes and authentication.

A relevant standard for authorisation is the **eXtensible Access Control Markup Language (XACML)** [14] defined by the OASIS consortium. It provides an XML-based language to express and evaluate authorisation policies to protect resources in a distributed computing environment.

4.2.2 Other relevant standards for security

Cloud computing does not add any further complexity for the protection of data transmitted between services or parties. Known security standards for authenticating one or both parties over the Internet and encrypt the data in transit can be used to provide data confidentiality at the application (HTTP over TLS), transport (TLS – IETF) or networking layer (IPSec - IETF).

These protocols can also use **X.509** certificates for establishing the identity of a party and authenticate. This is ITU-T standard for a Public Key Infrastructure (PKI) and Privileged Management Infrastructure (PMI) that specifies the formats of public key and attribute certificates [42].

In the context of web service security, the World Wide Web Consortium (W3C) and OASIS have defined a set of specification to help formatting messages with the intent to enable interoperability among services of different vendors and to specify the minimum information required to establish a secure exchange of messages. **WS-Security** [32] is an OASIS security standard that extends the Simple Object Access protocol (SOAP) to define how integrity and confidentiality can be enforced on messages by using Extensible Markup Language (XML) Signature and XML Encryption. **WS-Policy** [13] is a W3C standard that defines how web services can use XML to express their constraints in terms of policies on security and QoS and how end users can specify their policy requirements (expressed as policy assertions).

4.2.3 Privacy relevant standards

The loss of control over data represents one of the major threats for cloud service consumers that CLARUS addresses by designing solutions that guarantee the security and privacy of sensitive information. Few standards exists that target specifically privacy and in particular the protection of personal data in information and communication technology (ICT) systems.

The **ISO/IEC 29100** [18] specifies a privacy framework for the protection of Personally Identifiable Information (PII). It is a general framework that targets organisations and supports them for the definition of privacy requirements that should be considered complementary to legal ones, whenever personal information is processed. In addition, the standard specifies a common privacy terminology, defines the actors and their roles in processing personally identifiable information, and includes a set of eleven privacy principles for ICT systems. The

ISO/IEC 29101 [19] defines a privacy architecture framework. It specifies important concerns that should be considered for the design of ICT systems that process personal identifiable information, lists the components of the system for the implementation of such systems, and provides architectural views contextualizing these components.

The OASIS organisation has established two technical committees that have produced two standards working on related privacy aspects. The **Privacy Management Reference Model (PMRM)** [20] provides a guideline for developing operational solutions to privacy issues. The purpose of the standards is to define a methodology for the analysis of privacy policies and to serve as an evaluation framework of the different privacy management solutions, but it does not provide a specific implementation.

The other OASIS technical committee is the **Privacy by Design Documentation for Software Engineers (PbD-SE)** [21] and provides guidelines that enable software organizations to embed privacy into the design and architecture of IT systems, without diminishing system functionality. The guidelines follow the foundational principles of *privacy-by-design* [1].

4.3 Standards for web services and data format

The last set of standards analysed concerns web services and data format. These standards are not specific for cloud computing, but they are widely used to provide cloud services to end-users. In the context of CLARUS, the former is used for the definition of the interfaces and protocols to transmit the information at the Software-as-a-Service (SaaS) layer. Indeed, cloud computing leverages, among other technologies, Service Oriented Architecture (SOA) that uses web services to facilitate the implementation of the system. The latter is considered in the framework of the two CLARUS business cases, the eHealth and the Geo-reference data use cases that require the management of data in a standardised way.

4.3.1 Standards for Web services

The World Wide Web Consortium (W3C) and OASIS have developed the majority of the standards for web services. In this section, we focus on the messaging and metadata specifications, as standards for web services security have been already discussed in Section 4.2. In CLARUS, these standards could be relevant for specific use cases that require the publication of services.

The **Simple Object Access protocol** (SOAP), specified by W3C, is a standard for the definition of the XML-based format of information exchanged in the implementation of web services. It is an extensible protocol (e.g. WS-Security for security specifications) and uses HTTP as transport protocol. The **Universal Description, Discovery and Integration** (UDDI) [22] is an OASIS specification for web service directory; it defines a XML-based registry, used to list and locate web service applications, that can be interrogated by SOAP messages. The **Web Services Description Language (WSDL)** [23] is a W3C specification that describes the functionalities of the web service and defines the message formats required to access it.

Another relevant standard, not specific to cloud computing, is **JavaScript Object Notation (JSON)** [65], which uses human-readable text to transmit data objects between a server and a web application. JSON is a widely used standard for the development of web services.

4.3.2 Data format

The application of CLARUS to a number of different application scenarios requires the support of standard data formats. The two use cases of CLARUS are: the Publication of geo-reference data and the eHealth scenarios. In the following sections we discuss relevant standards for each use case.

4.3.2.1 Publication of geo-reference data

The Open Geospatial Consortium (OGC) [24] issues the standards relevant for the publication of geo-reference data use case. It is an international member-based consortium of organisations from the public and sector, academia and research with the intent of developing standards exclusively on geospatial data and related topics, including the geospatial law. OGC standards address specific interoperability challenges for data, processing systems, and geospatial technology.

The **OGC Catalogue Service for the Web (CSW)** provides guidelines for the implementation of catalogue services that support the ability to publish and search collections of metadata for data, services and related information objects. It defines the interfaces to enable a client to query the catalogues and discover of resources. The catalogue manages a metadata repository that can store multiple descriptions of the same resource, based on metadata scheme. Thus, it enables the use of multiple query languages to the same catalogue.

The **OGC Web Map Service (WMS)** specifies the standard interface that a user uses to view maps from a diverse range of data providers (each one storing a layer of the map, or different portions of the map) through a single application. WMS enables two or more maps with equivalent geographic parameters and output size to be accurately overlaid to produce a composite map. The standard also defines the rendering (visualization and querying) of the image delivered based on the location of interest selected by the user. The OGC WMS Interface Standard has been adopted by ISO under standard ISO 19128.

The OGC **Web Feature Service (WFS)** specifies the interface to download geographical information via a web service. It defines the operations of the service to provide access to the geographical features: discovery, query, transaction (features can be created, changed, deleted) locking, and stored query.

4.3.2.2 eHealth

International Standard Organisation (ISO)

The ISO/TC 215 technical committee addresses the standardisation of health informatics to facilitate the exchange of health related data, thus, addressing compatibility and interoperability between independent systems. Below we discuss some of the guidelines proposed by ISO for health informatics. We do not cover ISO/TR 22221, that deals with principles and practices for the clinical data warehouse since CLARUS stores encrypted medical records on the cloud, without requiring specific data management.

The **ISO/TS 21547** [25] discusses the security requirements for archiving of electronic health records in any format for the long term. This standard specification discusses the document management and privacy protection, rather than specific messages and protocols, and applies the same care for the management of Electronic Health Records (HERs) as in the paper form. Document management is intended as the practise to archive documents, which can be implemented as a separate independent archive or a federated one. HERs management includes maintenance, retention, disclosure and destruction. The standard also focuses on

security requirements (integrity, confidentiality, availability and accountability) and privacy requirements to protect the patient records for their long-term digital preservation in digital archives. The ISO/TS 21547 covers the functional requirements for security and archiving but it does not discuss the technology and the archiving models. The **ISO/TR 21548** [26] complementary Technical Report provides additional guidance for implementation of requirements defined in ISO/TS 21547. It discusses practical methods and tools for the development and management of digital archives that satisfy the security requirements.

The **ISO 22600** [34] standard defines principles and specifies services needed for managing privileges and access control to data distributed across policy domain boundaries. It proposes a template for policy agreement for the different stakeholders of the healthcare information system, including patients and staff members, and defines how the communication should be managed. The policy agreement must include all the differences in the security systems of the stakeholders in different domain boundaries and the agreed solutions on how to overcome the differences.

The **ISO 22857** [27] provides guidance on data protection requirements to facilitate the transfer of personal health data across national or jurisdictional borders. The standard does not require the harmonisation of the national legislations in terms of data protection and national guidelines to prevent threats to the privacy of the individual, i.e. ensure that medical data of a patient is adequately protected when transmitted and processed by another organisation. The goal is to ensure compliance to security policy principles of an organisation in the trans-national transfer of personal data. If a legal agreement exists between two nations, that agreement has the precedence over the standard.

The **ISO/TS 25237** [28] contains principles and requirements for privacy protection using pseudonymization of health records. The specification defines organisational and technical aspects for pseudonymization (reversible and irreversible) and gives a guide to risk assessment in case of re-identification. Furthermore, it specifies a policy framework and minimal requirements for pseudonymization.

The **ISO 27799** [29] is a standard (the updated version is in draft) that provides guidance for the application and implementation of ISO/IEC 2700 for the health sector. The target is organisations holding or processing personal health information and the standard describes how these organisations should protect the information and maintain the confidentiality, integrity and availability of personal health information.

Health Level Seven International (HL7)

The Health Level Seven International defines a set of international standards for the exchange, integration, sharing and retrieval of clinical and administrative health information between information systems used by various healthcare providers. The standards focus on the application layer (7^{th}) of the OSI model, thus giving the name to the family Health-Level 7 (HL7). HL7 standards are meant to facilitate information transfer between healthcare organisations by defining a set of rules that allow information to be shared and processed in a uniform and consistent manner. The HL7 standards have also influenced ISO standards.

The **HL7 Health Level Seven Version 3 (V3)** standard focuses on interoperability of the health and medical transactions. It specifies how the information should be presented in a clinical context to ensure that the two parties of a transaction share the semantics of the data exchanged. The messaging standard defines a set of interactions, i.e. XML-based messages, to support all healthcare workflow. The Reference Information Model (ISO/HL7 21731 [30])

expresses the data content needed in a specific clinical or administrative context. The HL7 Development Framework (ISO/HL7 27931 [31]) specifies messaging, processes, tools, actors, rules, and artefacts relevant to development of all HL7 standard specifications for the development of an interoperable healthcare framework.

The security technical committee of HL7 has produced a set of guidelines for the security and privacy policy management, privilege management, access control and auditing. Some of these standards are: HL7 Healthcare Privacy and Security Classification System (HCS), Role-based Access Control Healthcare Permission Catalog (RBAC), HL7 Version 3 Standard: Privacy, Access and Security Services; Security Labeling Service (SLS), and the Privacy, Access and Security Services (PASS).

5 Requirements: foundation, best practices and mapping

This section describes the methodology adopted in the requirement gathering process. The approach is common to deliverable D2.2, and it is an extension of the requirement classification method proposed in WP2.

CLARUS adopts the following pre-requisites for the definition of requirements. They must be *clear* to understand and unambiguous; *verifiable* to check whether a requirement is met or not; *traceable* to refer always to the source and need for a requirement; *numbered* to be easily referred to in the design and implementation phase; *prioritised* to understand which are the requirements critical for the project execution.

5.1 Requirement gathering process

Requirements are gathered via an interactive process inside the CLARUS consortium, by categorizing the source of the requirements. The main classification is proposed in D2.2 and herein we focus on the process aimed at deriving requirements from reports and best practices that consider standards or de-facto standards. In addition to the technical standards relevant for CLARUS and discussed in Section 4, best practices from SDOs have been analysed as the main sources for requirements. The gathering process consisted in several steps; first, they have been identified from the sources below described, refined based on the requirements identified in D2.2, and then they have been assigned a different priority and a description.

The **ETSI Cloud Standards Coordination Final Report** [47] document is a collection of 90 use cases that have been further categorized in high-level use cases targeting the different phases of cloud service provision: (i) Set-Up Cloud Service, (ii) Prepare & Procure service, (iii) Operate the service, (iv) Use Service, and (v) Assure Quality. We consider those that are relevant for CLARUS including Data Security and Privacy, Interoperability, Data Portability, Support EU Policies, and Based on Real life situations. The standards considered in this report and relevant for CLARUS are: OVF (DMTF), CIMI (DMTF, v1 ISO/IEC 17203), CDMI (SNIA, ISO/IEC 17826:2012), TOSCA (OASIS), OCCI (OGF).

ENISA report on "Standardisation in the field of Electronic Identities and Trust Service Providers" [48] analyses the standard landscape in cyber security in Europe and provides a reference to the standards. The report does not include use cases, but a set of requirements is derived from recommendations and best practices in cyber security to limit the security threats.

ENISA report on "Privacy and Data Protection by Design – from policy to engineering" [49] gives recommendations for design of systems following a *private-by-design* approach. The report focuses on the legal framework, but it gives hints for technical requirements that must be satisfied to ensure data protection. It does not include a list of use cases.

ENISA report on "Standards and tools for exchange and processing of actionable information" [50] focuses on information sharing standards. CLARUS must support some of those to ensure compatibility with external components and facilitate the processing of data. Relevant for CLARUS are the standards implementing data transport and serialization such as REST and RESTful architecture style, XML, JSON.

NIST report 7956 on "Cryptographic Key Management Issues & Challenges in Cloud Services" [51] focuses on guidelines to follow for the design of security capabilities for cloud services. The report identifies the objectives and the security capabilities that enable the operations to fulfil the objectives. We consider the capabilities for PaaS and SaaS solutions to derive requirements.

Cloud Security Alliance (CSA) document on "Security Guidance for Critical Areas of Focus in Cloud Computing v3.0" [52] discusses best practices for secure cloud computing operations. It is structured in 3 different sections: cloud architecture, governing the cloud, and operating in the cloud.

Cloud Standards Customer Council (CSCC) paper on "Cloud Security Standards: What to Expect & What to Negotiate" [53] focuses on information security requirements for public cloud deployment. It discusses the 10 steps that customers should take to evaluate and manage the security of the cloud. In the context of CLARUS we map the recommendations for cloud customers to the requirements that CLARUS must satisfy as cloud solution.

5.2 Naming scheme and priorities

WP2 has agreed on a common naming scheme to facilitate the tracking of the requirements, thus reducing the risk of possible duplication of requirements and identifying conflicting ones coming from different use cases (project-level and best practices). For completeness, the conflicting requirements will be reported and then refined in the updated version of D2.3 once the first architectural design will be available. Each requirement is identified with an ID, which is unique in CLARUS. In this document, the, requirements are named according to a common template:

- Source derived from the best practices or SDOs reports.
- Category of the requirement, which indicates requirements sharing common purpose or background and are likely to be listed and analysed together. The categories of interest are security (SEC), monitoring (MON), architecture (ARCH), privacy (PRI)
- Progressive Number in the category section.

As an example, ETSI.SEC.03 is a valid name indicating a requirement derived from the ETSI use case addressing Security.

Priorities for the requirements have been determined following the MoSCoW scale [2][15].

- Must have [MH]: the project will be seriously impacted if this requirement is not met.
- Should have [SH]: a requirement that, if it is not met, seriously impacts the project, but it can be delivered on a different timescale than the must-haves.
- Could have [CH]: optional requirements that can improve the project outcome and the value of its results, but these requirements are not essential to the main delivery (not too many Could-have requirements should be dropped in the design).
- Would like [WL]: truly optional features which can be delivered if there is sufficient effort available within a task, or if they can be found elsewhere and integrated easily.

5.3 Requirements from best practices

Requirement ID	Description	Priority	Comment
ETSI.ARCH.01	Enabling Data Portability: Agreement on common formats of the Data [48].	SH	The CLARUS proxy should support standardised data format.

Requirement ID	Description	Priority	Comment
ETSI. ARCH.02	Enabling Interoperability: Agreement on common interfaces between the provider and the customer, including management and administration interfaces [48].	SH	The CLARUS solution should be interoperable with different cloud service providers
NIST.ARCH.03	PaaS security capability: The ability to set up secure interaction with deployed applications and/or development tool instances [51].	MH	The CLARUS architecture must support tools to enable secure data transmission.
NIST.ARCH.04	PaaS security capability: The ability to securely store static data (data not directly processed by applications) [51].	SH	The CLARUS architecture could have the ability to store data securely.
NIST.ARCH.05	PaaS security capability: The ability to securely store application data in a structured form (e.g., relational form) using a Database Management System (DBMS) [51].	SH	The CLARUS architecture should allow secure storage of data in structured databases.
NIST.ARCH.06	PaaS security capability: The ability to securely store application data that is unstructured [51].	CH	The CLARUS architecture should consider the use of unstructured data storage.
NIST.ARCH.07	SaaS security capability: The ability to set up secure interaction with an application [51].	MH	The CLARUS architecture must consider secure communication with SaaS services.
NIST.ARCH.08	SaaS security capability: The ability to store application data (structured or unstructured) in an encrypted form [51].	MH	The CLARUS architecture must consider secure storage of data provided via SaaS services.
CSA.ARCH.09	When possible, use platform components with a standard syntax, open API's, and open standards [52].	MH	The implementation of the CLARUS solution should consider the use of standard interfaces.
CSA.ARCH.10	Develop application architecture with layers of abstraction to minimize direct access to proprietary modules [52].	MH	CLARUS PaaS solution must be implemented with an abstraction layer to ease the support for interoperability.
CSA.ARCH.11	Cloud customers should not depend on a singular provider of services and should have a disaster recovery plan in place that facilitates migration or failover should a supplier fail [52].	CH	CLARUS should enforce mechanisms to avoid vendor lock-in and migration of data and services to other providers.

Requirement ID	Description	Priority	Comment
CSA.ARCH.12	Cloud computing architecture patterns that explicitly mitigate threats should be used [52].	WL	The design of the architecture could account for patterns that mitigate threats.
ETSI.SEC.01	The customer must analyse its Data Privacy obligations with respect to the PII (if any) that will be processed by the cloud services, and build a set of security and privacy requirements that must be fulfilled by the cloud service provider [48].	MH	The CLARUS proxy must support the security and privacy requirements defined by the customer
ETSI.SEC.02	Administration of users, identities and authorizations [48].	MH	The CLARUS solution must implement mechanisms for managing users, identities and authorisations
CSA.SEC.03	Detecting and Preventing Data Migrations to the Cloud [52].	WL	The CLARUS solution would like to have mechanisms that control data migration to cloud and prevent data loss.
CSA.SEC.04	Protecting Data Moving To (And Within) the Cloud: Encrypt all sensitive data moving to or within the cloud [52].	MH	CLARUS must protect the transmission of data to the cloud via encryption.
CSA.SEC.05	When using data protection, pay particular attention to key management [52].	MH	The CLARUS proxy must implement key management schemes.
CSA.SEC.06	Encrypt sensitive data in PaaS applications and storage [52].	MH	Sensitive data managed by the CLARUS PaaS solution must be protected.
CSA.SEC.07	When using application encryption, keys should be stored external to the application whenever possible [52].	MH	Keys must not be stored within the application.
CSA.SEC.08	Use best practice key management practices when using any form of encryption/decryption [52].	MH	CLARUS must implement key management for the proxy solution.
CSA.SEC.09	Use best practice key management practices and obtain technology and products for encryption, decryption, signing, and verifying from credible sources [52].	MH	Verify the source of the tools used for cryptographic operations.

Requirement ID	Description	Priority	Comment
CSA.SEC.10	Use standard algorithms. Do not invent/use proprietary scrambling techniques. Proprietary encryption algorithms are unproven and easily broken. Avoid old insecure encryption standards such as Data Encryption Standard (DES) [52].	MH	Use standard and secure encryption algorithms.
CSA.SEC.11	Implementers should, where possible, use federation based on open standards such as SAML and OAuth [52].	SH	Use standards for access rights delegation and id management.
CSA.SEC.12	Implementers should ensure all sources of Identity/Attributes provide organizational Identity [52].	MH	CLARUS should verify the level of assurance of the Id provider.
CSA.SEC.13	Implementers should ensure that Attributes are validated at master/source whenever possible, or as close as possible [52].	MH	Attributes for authorisation need to be checked as often as needed.
CSCC.SEC.14	Ensure CSP supports federated IDs and single sign-on using one or more of the following standards: LDAP, SAML 2.0, OAuth 2.0, WS-Federation, OpenID Connect, SCIM [53].	MH	CLARUS must implement open standards to provide federated IDs and single-sign-on scheme.
CSCC.SEC.15	Ensure CSP provides access control and security policy decisions leveraging a standard such as XACML [53].	MH	CLARUS must implement open standards for access control.
CSCC.SEC.16	Ensure CSP supports one or more of the following standards for security certificates: PKCS, X.509, OpenPGP [53].	MH	CLARUS must implement a security certification scheme.
CSCC.SEC.17	Ensure CSP supports one or more of the following standards for data in motion: HTTPS, SFTP, VPN using IPSec or SSL [53].	MH	CLARUS uses secure transfer protocols for data transmission between components and with the client.
ENISA.PRI.01	Providers of software development tools need to provide tools that enable the intuitive implementation of privacy properties [29].	MH	CLARUS must have development tools that make the design and implementation of privacy properties intuitive.

Requirement ID	Description	Priority	Comment
ENISA.PRI.02	Privacy supporting components, such as key servers, anonymising relays, should be included [49].	MH	The CLARUS proxy should include anonymising relays and encryption libraries.
ETSI.MON.01	Monitoring Service Levels: Availability, Incident Management, Storage performance, Processing performance, Access security event information, Uptime [48].	MH	The CLARUS solutions must implement monitoring and report to the user.
ETSI.MON.02	Keep a record of past transactions, under data retention obligations [48].	MH	On termination of a cloud service, the CLARUS proxy must keep the record of past transactions

5.4 CLARUS requirements mapping with standards

In the following table we map the requirements reported in Deliverable D2.2 "Requirements specification V1"[3] to the standards that we have identified in Section 4.

Requirement IDs	Relevant standards	Comment
REQ-NF_IMP-1.2, REQ-F_FEA-2.1, REQ-F_FEA-2.2, REQ-F_FEA-5.41, REQ-F_FEA-5.42, REQ-F_FEA-5.45	Storage standard: CDMI Authorisation: OAuth2, XACML	Authorisation via delegation and policies evaluation on stored data.
REQ-NF_OPE-1.1, REQ-F_BRU-1.3, REQ-F_BRU-1.9, REQ-F_USI-1.1	Storage: CDMI PaaS proxy solution: TOSCA, CAMP. Data communication: JSON	Interoperability of the CLARUS solution to allow data distribution and replication across different providers, Integration of existing applications for data management and compute.
REQ-NF_SEC-1.1; REQ-NF_SEC-1.2, REQ-F_FEA-1.11, REQ-F_FEA-3.1, REQ-F_FEA-4.1, REQ-F_FEA-4.2, REQ-F_FEA-5.15	Federated Id management: Shibboleth, OpenID, OAuth2, WS-Federation Authorisation: XACML	Management of the users and authentication to perform operations on data.
REQ-NF_TRSP-1.1, REQ-NF_TRSP-1.2	Transport security: HTTP over TLS, TLS, IPSec Web service security: WS-Security	Confidentiality and integrity of data transmitted.
REQ-F_FEA-5.0.2, REQ-F_FEA-5.0.3	LDAP, SAML	Remote user information for authentication and authorisation

Requirement IDs	Relevant standards	Comment
REQ-F_FEA-5.10	PaaS services: TOSCA Federated Id management: Shibboleth, OpenID, OAuth2, WS-Federation	Interoperability of CLARUS proxy solution
REQ-F_FEA-6.10, REQ-F_FEA-6.10.1, REQ-NF_ FEA-8.2, REQ-F_USI-1.1	CDMI	Standard for data storage elements
REQ-F_FEA-6.10.1	OVF	Specify storage end-point for the description of an application
REQ-F_FEA-8.1	Authentication and user management: Shibboleth, OpenID, OAuth2, WS-Federation	Single sign-on service and Federated Id management

6 Conclusion

This deliverable reviews the current standard landscape for cloud computing, security, and data format relevant for CLARUS. The objective of this document is to define the standardisation roadmap built on a throughout analysis of the standards and their potential adoption in the design and implementation of the CLARUS proxy solution.

The documents analyses all relevant standards for CLARUS that could be potentially adopted for the design of the architecture and implementation of the CLARUS solution. It focuses on identifying the standards for the definition of the Application Programming Interfaces (APIs) of the components and tools, the protocols for supporting security (authentication and authentication), and the format of data.

An analysis of the best practices for implementing cloud services is also discussed to derive additional requirements stemming from these recommendations. Finally, this deliverable maps the technical requirements identified in Deliverable D2.2 “Requirements specification V1” [3] to current standards.

The requirements identified at the early stage of the project will be further analysed and refined in D2.3 “Requirements Specification V2”.

References

[1] Foundational Principles. [Online]. https://www.privacybydesign.ca/index.php/about-pbd/7-foundational-principles/

[2] Dai Clegg and Richard Barker, *Case Method Fast-Track: A RAD Approach*, Addison-Wesley, Ed., 1994.

[3] CLARUS Consortium, "D2.2 Requirements specification V1 ," Deliverable 2015.

[4] CLARUS Consortium, "D7.1 Dissemination and standards report V1 ," Deliverable 2015.

[5] National Institute of Standards and Technology (NIST), "The NIST Definition of Cloud Computing," NIST Special Publication 800-145, 2011.

[6] (2007) Directive 2007/2/EC of The European Parliament and of the Council. [Online]. http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32007L0002

[7] Distributed Management Task Force. [Online]. http://www.dmtf.org

[8] Storage Networking Industry Association (SNIA). [Online]. http://www.snia.org

[9] IEEE. P2301- Guide for Cloud Portability and Interoperability Profiles (CPIP). [Online]. https://standards.ieee.org/develop/project/2301.html

[10] Organization for the Advancement of Structured Information Standards (OASIS). [Online]. https://www.oasis-open.org/

[11] OASIS, "Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0," Standard 2005.

[12] IETF, "The OAuth 2.0 Authorization Framework," Standard RFC 6749, 2012.

[13] W3C, "Web Services Policy 1.5 - Framework," Standard 2007.

[14] OASIS, "eXtensible Access Control Markup Language (XACML) v3.0," Specifications 2013.

[15] MoSCoW Prioritisation. [Online]. http://www.dsdm.org/content/10-moscow-prioritisation

[16] IEEE, "Standard for Intercloud Interoperability and Federation (SIIF)," Intercloud WG (ICWG) Working Group, P2302,.

[17] Open Grid Forum. [Online]. https://www.ogf.org/

[18] ISO/IEC, "ISO/IEC 29100 - Information technology - Security techniques - Privacy framework," Standard 2011.

[19] ISO/IEC, "ISO/IEC 29101 - Information technology - Security techniques - Privacy architecture framework," Standard 2013.

[20] OASIS, "Privacy Management Reference Model (PMRM) 1.0," Standard 2013.

[21] OASIS, "Privacy by Design Documentation for Software Engineers (PbD-SE)," Standard.

[22] OASIS, "Universal Description, Discovery and Integration (UDDI) v3.0," Standard 2005.

[23] W3C, "Web Services Description Language (WSDL) 2.0," Standard 2007.

[24] Open Geospatial Consortium (OGC). [Online]. http://www.opengeospatial.org

[25] ISO, "ISO/TS 21547 - Health informatics - Security requirements for archiving of electronic health records - Principles," Standard 2010.

[26] ISO, "ISO/TR 21548 - Health informatics - Security requirements for archiving of electronic health records - Guidelines," Standard 2010.

[27] ISO, "ISO 22857 - Health informatics - Guidelines on data protection to facilitate trans-

border flows of personal health data," Standard 2013.

[28] ISO, "ISO/TS 25237 - Health informatics - Pseudonymization," Standard 2008.

[29] ISO, "ISO 27799 - Health informatics - Information security management in health using ISO/IEC 27002," Standard.

[30] ISO, "ISO/HL7 21731 - Health informatics - HL7 version 3 - Reference information model," Standard 2014.

[31] ISO/HL7, "ISO/HL7 27931 - Data Exchange Standards - Health Level Seven Version 2.5 - An application protocol for electronic data exchange in healthcare environments," Standard 2009.

[32] OASIS, "Web Services Security v1.1," Standard 2004.

[33] OASIS, "Topology and Orchestration Specification for Cloud Applications (TOSCA) v1.0," Standard 2013.

[34] ISO, "ISO 22600 - Health informatics - Privilege management and access control," Standard 2014.

[35] OASIS, "Cloud Application Management for Platforms (CAMP) v1.1," Standard 2014.

[36] ISO/IEC, "ISO/IEC 27001 - Information technology - Security techniques - Information security management systems - Requirements," Standard 2013.

[37] ISO/IEC, "ISO/IEC 27002 - Information technology - Security techniques - Code of practice for information security controls," Standard 2013.

[38] ISO/IEC, "ISO/IEC 27018 - Information technology - Security techniques - Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors," Standard 2014.

[39] ISO/IEC, "ISO/IEC FDIS 27017 - Information technology - Security techniques - Code of practice for information security controls based on ISO/IEC 27002 for cloud services," Standard.

[40] OASIS, "Web Services Federation Language (WS-Federation) Version 1.2," Standard 2009.

[41] ITU-T, "X.1600 - Security framework for cloud computing," Standard 2014.

[42] ITU, "X.509 - The Directory: Public-key and attribute certificate frameworks," Standard 2012.

[43] SNIA, "Cloud Data Management Interface (CDMI) v1.1.1," Standard 2015.

[44] DMTF, "Cloud Infrastructure Management Interface (CIMI) v2.0," Cloud Management Working Group, Standard DSP0263, 2015.

[45] DMTF, "Open Virtualization Format (OVF) Specification," Open Virtualization Format Working Group, Standard DSP0243, 2013.

[46] OGF, "Open Cloud Computing Interface (OCCI) v1.1," Standard GFD.183-5, 2011.

[47] ETSI, "Cloud Standards Coordination - Final Report," Report 2013.

[48] ENISA, "Standardisation in the field of Electronic Identities and Trust Service Providers v1.0," Report 2014.

[49] ENISA, "Privacy and Data Protection by Design - from policy to engineering," Report 2014.

[50] ENISA, "Standards and tools for exchange and processing of actionable information," Report 2014.

[51] NIST, "Cryptographic Key Management Issues & Challenges in Cloud Services," Report NISTIR 7956, 2013.

[52] Cloud Security Alliance (CSA), "Security Guidance for Critical Areas of Focus in Cloud Computing v3.0”," Report 2011.

[53] Cloud Standards Customer Council (CSCC), "Cloud Security Standards: What to Expect & What to Negotiate," Report 2013.

[54] CLARUS Consortium, "Legal and ethical requirements," Deliverable 2015.

[55] OpenNebula. [Online]. http://opennebula.org

[56] OpenStack. [Online]. https://www.openstack.org

[57] European Grid Initiative (EGI). [Online]. http://www.egi.eu

[58] Cloud Foundry. [Online]. https://www.cloudfoundry.org

[59] Heroku Cloud Application Platform. [Online]. https://www.heroku.com

[60] OpenShift by RedHat. [Online]. https://www.openshift.com

[61] Docker. [Online]. https://www.docker.com

[62] IETF, "Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map," Standard RFC 4510, 2006.

[63] Internet2. Shibboleth. [Online]. http://www.internet2.edu/products-services/trust-identity-middleware/shibboleth/

[64] OpenID Foundation. OpenID. [Online]. http://openid.net

[65] JavaScript Open Notation (JSON). [Online]. http://json.org