Jostein Jensen

Federated Identity Management in the Norwegian Oil and Gas Industry

Thesis for the degree of Philosophiae Doctor

Trondheim, March 2014

Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics
and Electrical Engineering
Department of Computer and Information Science

NTNU – Trondheim
Norwegian University of
Science and Technology

NTNU
Norwegian University of Science and Technology

Thesis for the degree of Philosophiae Doctor

Faculty of Information Technology, Mathematics and Electrical Engineering
Department of Computer and Information Science

ISBN 978-82-326-0156-1 (printed ver.)
ISBN 978-82-326-0157-8 (electronic ver.)
ISSN 1503-8181

Doctoral theses at NTNU, 2014:116

Printed by NTNU-trykk

Smerte er godt, smerte er midlertidig
(Aspirant Heiskel, Jørstadmoen 1999)

Abstract

The Norwegian oil and gas industry has a highly collaborative, but at the same highly competitive nature. Most of the daily oil and gas production takes place on the Norwegian continental shelf in the North Sea. The production facilities are expensive to develop and maintain, and it is therefore necessary to take advantage of new and innovative solutions; both above and beneath the sea surface. Close collaboration between the operators and the contractors[1] is needed. At the same time these companies can be strict competitors in other projects. Information security is thus essential; only information relevant for a given collaboration should be available to the involved parties.

Federated identity management (FIM) is a concept that allows cooperation on technologies, processes and policies for identity management, as well as sharing of identity data across organizational boundaries and across security domains. Many current information security challenges within the industry are related to access control and identity management. The goal of this PhD project has been to analyze companies involved in the Norwegian oil and gas production in order to explore their perceived benefits, challenges and other security risks related to adoption of FIM.

In order to meet our research goal we have based our research on three research methods: a design science approach, systematic literature reviews, and a case study. Empirical evidence related to the oil and gas industry and its perception of FIM is mainly collected through the case study, using semi-structured interviews to collect data.

First, our research shows that a focus on security is needed throughout the whole software development lifecycle when developing identity management solutions. It is especially important to protect the identity assertion. Federated identity management is more than just technology. Collaborators within the federation must agree on common rules and security policies for all phases of the identity management lifecycle.

Second, the federated identity research community should spend more effort on empirical research. Great initiatives exist to move the technology into academic perfection, however, little empirical evidence exists to document real world expectations and needs.

Third, this research has listed the benefits and challenges of FIM from an academic perspective and from an industry perspective. We have also documented many of the challenges the industry is faced with today related to access control.

Our interviews with the industry practitioners show that some of the benefits of FIM are offset by their challenges. However, we believe that some forms of federated identity management will be implemented in some form sooner or later. This research can be used as input to tailor new identity management solutions to the Norwegian oil and gas industry's needs, it can be used to highlight the need for security in the software development process, and it can be used to understand the strengths, weaknesses, opportunities and threats related to adoption of federated identity management in the industry.

[1] including e.g. equipment vendors, oil service companies, engineering companies

Preface

This thesis is submitted to the Norwegian University of Science and Technology (NTNU) for partial fulfilment of the requirements for the degree of philosophiae doctor.

This doctoral work has been performed at the Department of Computer and Information Science, NTNU, Trondheim, with Professor Torbjørn Skramstad (NTNU) as the main supervisor and with senior research scientist Martin Gilje Jaatun (SINTEF ICT) as co-supervisor.

The work has been funded by the Norwegian Research Council through the GoICT project, grant 183235/S10.

Acknowledgements

I would like to thank my supervisor Torbjørn Skramstad for giving me the opportunity to pursue my research interests through this PhD project. I must also thank my co-supervisor, colleague and friend Martin Gilje Jaatun for his support from start till end. His rich knowledge, analytical skills and passion for information security in general have been invaluable for my motivation to pursue the PhD degree. My fellow PhD candidate on the GoICT project, Åsmund Ahlmann Nyre also deserves attention: Thank you for the collaborative writing on two papers, for sharing joy and frustration related to the PhD work, and for discussing work-unrelated topics, such as house building projects, which has relieved the mind from the narrow focus on a specialized research topic.

I have also had a part time position at SINTEF in parallel with the PhD work at the NTNU. I am grateful for this opportunity given by my research director at SINTEF ICT, Department for Software Engineering, Safety and Security, Eldfrid Øfsti Øfstedal. Thanks to all my colleagues at SINTEF, and especially the information security research group for your encouragement to grow as security researcher and your catching enthusiasm for information security.

The most important support during the work with this PhD thesis has come from my family. My wife, Tuva, I am forever thankful for your unreserved support, care and love. My children: Christina, Markus and Morten, you give me inspiration and motivation.

Contents

List of Tables

List of Figures

Abbreviations

ABAC	Attribute Based Access Control
API	Application Programmer Interface
CIA	Confidentiality, Integrity, Availability
COTS	Commersial Off The Shelf
CVE	Common Vulnerability Enumeration
FIM	Federated Identity Management
HTTP	Hyper Text Transfer Protocol
ICT	Information and Communication Technology
ID	Identity / identifier
IdM	Identity Management
IEC	International Electrotechnical Commission
IO	Integrated Operations
IP	Internet Protocol
ISO	International Standards Organization
IT	Information Technology
ITU	International Telecommunication Union
LOA	Level of Assurance
MDD	Model Driven Development
NSM	Norwegian National Security Authority (norwegian: Nasjonal Sikkerhetsmyndighet)
NTNU	Norwegian University of Science and Technology
PST	Norwegian Police Security Service (norwegian: Norsk Sikkerhetsmyndighet)
RBAC	Role Based Access Control
RQn	Research Question 'n'
SAML	Security Assertion Markup Language
SCADA	Supervisory Control and Data Acquisition Systems
SDK	Software Development Kit
SOA	Service Oriented Architecture
SOAP	Simple Object Access Protocol
SPn	Secondary Paper 'n'
SSO	Single-Sign-On
SWOT	Strengths, Weaknesses, Opportunities and Threats
US	United States
WS	Web Service

Chapter 1

Introduction

1.1 Motivation

The oil and gas industry on the Norwegian continental shelf has migrated from the use of closed, autonomous offshore production facilities to production facilities where offshore (including subsea) and onshore facilities are highly interconnected. The use of modern information and communication technology has facilitated this migration, and the concept has been called integrated operations (IO). The new work and information flows improve information sharing, and facilitate collaboration among personnel e.g. from the operator, the service companies and vendors [2]. Figure 1.1 gives an indication of the collaborating actors in the IO environment. The new opportunities, however, also lead to new challenges. Jaatun and his colleagues [40] explain that a migration from stand-alone proprietary systems to commercial-off-the-shelf (COTS) systems, combined with increased connectivity between office networks and supervisory control and data acquisition (SCADA) systems, increase the likelihood of security incidents.

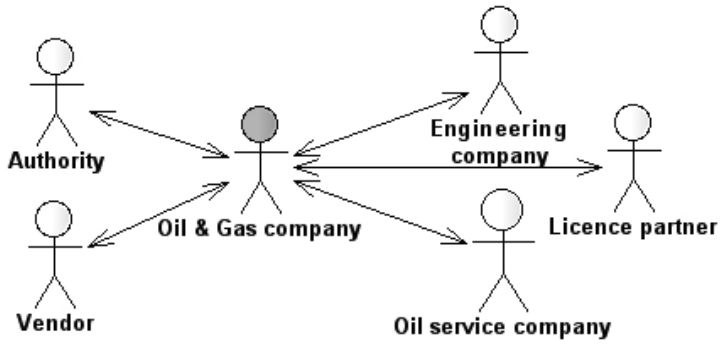

Figure 1.1: Actors in the IO domain. Adapted from Tungland et al. [78]

Information security is an important success factor for integrated operations, and new means to achieve secure access to IT- and process control systems must be devised [2]. A solid identity management foundation is essential to ensure that information is disclosed only to *authorized* personnel (confidentiality), that only *authorized* entities can modify or delete information (integrity), and that information is available to *authorized* individuals upon request (availability). Federated Identity Management (FIM) is a concept specifically targeted at collaborative environments, and is a concept that allows cooperation on processes, policies and technologies for identity management across company borders to facilitate secure and seamless information sharing.The following goal was stated for the this thesis:

> *This PhD project will analyse companies involved in close inter-organisational collaboration in Integrated Operations for the purpose of exploring the perceived benefits, challenges and security risks related to adoption of federated identity management in the Norwegian oil and gas industry.*

1.2 Research Questions

FIM has gained popularity over the last few years. Governments around the world (also in Norway) have deployed federated identity management services that can be used by citizens to access digital government services. Online banks have taken advantage of the FIM technology so that they can share, and benefit from, a common IdM infrastructure. Finally, there are several examples on the Internet, where the credentials issued by your favourite social media site can be used to gain access to other digital online services. FIM adoption, however, has been slower than expected [74] [55] [37], and it's close to non-existent within the Norwegian oil and gas industry.In order to gain a better understanding of the industry's perception of FIM we stated the following research questions:

RQ1 What are the main security challenges faced during development of secure identity management solutions for a distributed service platform?

RQ2 What empirical evidence exists on the development and adoption of FIM in industry?

RQ3 What are the benefits and challenges related to FIM from the perspective of the academic community and the industry, respectively?

1.3 Research Context and Path

This thesis is financed by the GoICT project[1] funded by the Norwegian Research Council. The early results, however, also draw upon the European Commission funded research project MPOWER[2].

The goal of the GoICT project was to create guidelines for development of high quality, dependable software within the energy sector, and more specifically in the Norwegian oil and gas industry. Critical components within this industry are increasingly more reliant upon software components, which affect the risk associated with on-going operations. The software dependability guidelines resulting from the project should allow development of software with qualities related to ensuring business continuity, safety and security, to keep risk at an acceptable level. This thesis contributes to the security aspect of the GoICT project.

The aim of the MPOWER project was to develop a middleware platform to support implementation of distributed services in an assisted living scenario. Smart home and sensor technology were to be integrated with profession and institution-specific systems to allow remote follow up of, and interaction with, elderly and people with cognitive disabilities. Health scenarios are automatically subject to strict security requirements in regard to confidentiality, integrity and availability of sensitive personal data [49]. Identity management, including user authentication, is a fundamental element to ensure fulfilment of the security requirements in such a setting. Consequently, the project included development of security services to handle the identity management. Due to usability requirements, a solution to achieve single-sign-on was selected.

Although the goals of the two projects seem unrelated, many elements from a distributed health service scenario and the oil and gas industry's initiatives for integrated operations are similar. The most important being:

- A number of actors are involved
- Collaboration and information sharing among the different actors are key to success
- Both rely on distributed services
- Both scenarios imply strict security requirements to protect sensitive information

The MPOWER project was nearing its end while the PhD work started. Being responsible for the security design of the healthcare platform, and given the similarities between the two scenarios, the idea of developing a common IdM solution for the oil and gas industry's IO environment seemed compelling. The original plan of the PhD work was therefore to reuse knowledge about IdM solutions as prerequisite for secure data exchange and single-sign-on from the MPOWER project, and adapt it to a new industrial environment with a

[1] GoICT, grant 183235/S10 from the Norwegian Research Council

[2] MPOWER, Contract No. 034707, Specific Targeted Research or Innovation Project (STREP) within the 6th Framework Programme

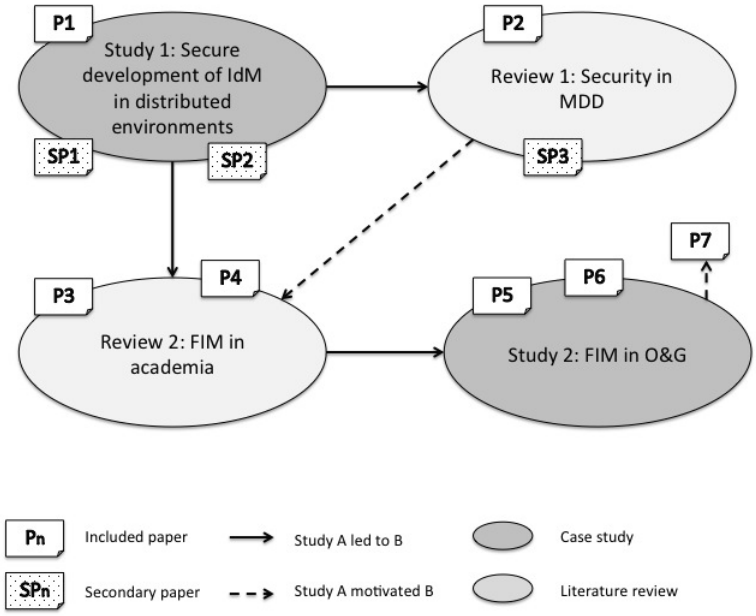

Figure 1.2: Studies, papers and contributions

scientific approach to tailor the solution and its functionality to a new environment. Instead of protecting sensitive personal health data, the solution could be used to facilitate secure seamless sharing of sensitive business information among the industry collaborators on the Norwegian continental shelf. The idea of federated identity management fit very well into this approach, and was pursued during the PhD project.

Figure 1.2 shows the relations between the studies that have been carried out within the scope of this thesis. It illustrates the progress that took place from the initial idea of having a technical design science approach based on the artifact developed in the MPOWER project (input to Study 1 in figure 1.2) to the end where a research approach based on an empirical case study focused on the Norwegian oil and gas industry's needs and expectations was followed (Study 2).

In Study 1, the design and the implementation of the IdM solution for the healthcare platform was available for testing and use. Before starting the adaptation towards usage in a new domain, however, we chose to do a security test of the solution. We wanted to learn from "past mistakes" - if any. As we will see later in section 4.1 this study identified several challenges with the existing solution; The implementation did not meet the requirements. This led to Review 1.

Model driven development (MDD) is an approach where design models are automatically translated to software code. Several initiatives have used an MDD approach to integrate

security in the design models, and as such claim to convert the design models to code with inherent security properties. In Review 1 we therefore carried out a literature review to learn from academics' past experiences, and see if this approach could be used to avoid some of the errors that had happened between the design and implementation that we found in Study 1. Unfortunately, we did not find empirical evidence to support claims that secure MDD practices would lead to more secure code, and in our case a more secure IdM solution. This approach was therefore not pursued further. The lack of empirical research on the topic of MDD, however, motivated Review 2. Would there be empirical evidence within another area of security research, namely research related to FIM adoption?

Close to the end of Review 1 we initiated Review 2 to learn more about the concepts of federated identity management. The ideas of FIM are very much in line with the SSO-inspired IdM approach taken in MPOWER. Academic literature was therefore studied to build upon existing research during design of a new identity management solution for the oil and gas industry. An important part of this study was to identify which benefits and challenges academics had identified during implementation of FIM solutions - again to build on the strengths of past experiences, and to avoid already identified pitfalls. Through this study we confirmed that the idea of FIM would be suited in an IO scenario in the oil and gas industry. However, again we found little empirical evidence related to actual implementation and use of FIM solutions. This led to a change of direction in the work towards this thesis and resulted in Study 2:

Instead of building a technical system that would have been built on good ideas and intentions and probably tested and evaluated in a small scale academic laboratory environment, we went out to the industry to catch real world empirical data as to what their expectations towards a FIM solution in a collaborative environment are. We wanted to learn more about the current challenges the industry is faced with regarding secure information sharing and identity management in a collaborative environment, and how they perceive the benefits and challenges of adopting FIM solutions. Is there a need for FIM at all, and is it realistic that the industry at large will adopt a common solution to facilitate collaboration? With lack of existing empirical evidence in academic FIM literature, this contribution represents something new to build upon when future FIM solutions are to be developed and tailored to industry needs.

1.4 Selected Papers

Figure 1.2 shows the papers that resulted from the different studies. The list below sums up the papers that were selected for inclusion in this thesis. The numbering follows the ID given in the figure. Additionally, the figure refers to three secondary papers that are related to Study 1 and Review 1. These provide background results and complementary information to the primary papers. Their abstracts are found in Appendix B.

1.4.1 Primary Papers

P1 Jostein Jensen, Åsmund Ahlmann Nyre: *"SOA Security - an experience report"*, Proc. The Norwegian Information Security Conference (NISK), 185-196, 2009

Reference ID in this thesis: [47]

P2 Jostein Jensen, Martin Gilje Jaatun: *"Not Ready for Prime Time: A Survey on Security in Model Driven Development"*, International Journal of Secure Software Engineering, 49-61, 2011

Reference ID in this thesis: [45]

P3 Jostein Jensen: *"Benefits of Federated Identity Management - A Survey from an Integrated Operations Viewpoint"*, Availability, Reliability and Security for Business, Enterprise and Health Information Systems, volume 6908 of Lecture Notes in Computer Science, 1-12, 2011

Reference ID in this thesis: [42]

P4 Jostein Jensen: *"Federated Identity Management Challenges"*, Proc. Seventh International Conference on Availability, Reliability and Security (ARES '12), 230-235, 2012

Reference ID in this thesis: [43]

P5 Jostein Jensen, Martin Gilje Jaatun: *"Federated Identity Management - We Built It; Why Won't they Come?"*, IEEE Security & Privacy, 34-41, 2013

Reference ID in this thesis: [44]

P6 Jostein Jensen, Åsmund Ahlmann Nyre: *"Federated Identity Management and Usage Control - Obstacles to Industry Adoption"*, Eigth International Conference on Availability, Reliability and Security (ARES '13), 2013

Reference ID in this thesis: [46]

P7 Jostein Jensen, *"Identity Management Lifecycle - Exemplifying the need for Holistic Identity Assurance Frameworks"*. Information and Communication Technology, volume 7804 of Lecture Notes in Computer Science, pages 343 - 352. Springer, 2013

Reference ID in this thesis: [48]

1.4.2 Secondary Papers

SP1 Jostein Jensen, Inger Anne Tøndel, Martin Gilje Jaatun, Per Håkon Meland, and Herbjørn Andresen, *"Reusable Security Requirements for Healthcare Applications"*.

Fourth International Conference on Availability, Reliability and Security (ARES '2009), 2009

Reference ID in this thesis: [49]

SP2 Richard Sassoon, Martin Gilje Jaatun, Jostein Jensen, *"The Road to Hell is Paved with Good Intentions: A Story of (In)secure Software Development"* Fifth International Conference on Availability, Reliability, and Security (ARES '10), 2010.

Reference ID in this thesis: [66]

1.5 Contributions

Table 1.1 shows the key findings of this thesis, their relation to the research questions and the papers where more details are found.

Table 1.1: Relations between key findings, research questions and papers

No	Key finding	RQ	Paper
1	IdM security specifications alone do not guarantee secure applications.	1	P1, SP1, SP2
2	Identity assertions must be properly protected.	1	P1
3	Secure identity management is more than secure technology.	1	P7
4	Little empirical evidence exists on the development and adoption of federated identity management.	2	P2, P3, P4
5	Academics expect increased privacy, security and usability for end users.	3	P3
6	Academics expect businesses to benefit from reduced administrative cost and complexity, improved data quality and security, and easier co-operation.	3	P3
7	Academics expect technical challenges related to interoperability, attribute synchronization and consistency, revocation and identity provider discovery.	3	P4
8	Academics expect organizational challenges related to investment cost, liability issues, identity assurance, security, knowledge and trust.	3	P4
9	Practitioners perceive that FIM will improve user administration and usability, make collaboration more efficient, reduce cost, facilitate audit and lead to better protection.	3	P5
10	Practitioners expect trust issues, technological challenges, investment cost and security challenges to be obstacles to adoption of FIM in industry, and there is a risk that they confuse identity management with access management.	3	P5
11	Practitioners question whether there is sufficient organizational maturity to adopt new identity management solutions.	3	P5, P6

The PhD thesis provides a body of knowledge that gives insight into federated identity management in general, and with a special focus on the Norwegian oil and gas industry. Table 1.2 provides an overview of the target audience of this work, and how they can take advantage of the results.

Table 1.2: Who this PhD will be useful for, and how they can benefit

Actors	**Benefit**
IT strategy planners and solution architects	Understand strengths, weaknesses, opportunities and threats of adopting identity federation technology in an inter-organizational industry environment.
Researchers	1) Understand federated identity management in an industrial domain, and focus research activities towards industry needs. 2) Motivate further empirical studies on federated identity management
Developers/product manufacturers	Understand company needs and develop solutions that meet requirements in a inter-organisational collaboration context in industry.

1.6 Structure of the Thesis

The remainder of this thesis is organized as follows: Chapter 2 presents background information to the field of digital identities and federated identity management. Chapter 3 provides a brief overview of the context of this research, i.e. the Norwegian oil and gas industry, as well as an overview of the selected research approaches and the motivation for their use. Chapter 4 presents the results sorted by the topics of each research question. The content of this chapter is copied from, and sometimes a narrative of the selected papers to improve the readability of the thesis. The most relevant results are included, however, more details and complimentary information can be found in each paper. Chapter 5 discuss the implications of the work, the limitations of the thesis results, and provides some recommendations for future research on the topic. Chapter 6 concludes the thesis. In Appendix A we have enclosed all the selected papers, while Appendix B provides references and abstracts of each supporting paper.

Both the research community and industry practitioners are identified as target groups of this thesis. Consequently, we provide details in the background section so that both groups can gain sufficient insight to understand our results and their implications.

Chapter 2

Background

Identification and authentication processes have always been important to hinder leakage of resources to unauthorized entities, also prior to the IT-era; Bank employees should verify their customers' ID-cards, such as driver's license or passport, before money can be withdrawn over the counter, pharmacy staff should verify their customers' ID before they deliver prescribed medications, and case workers should control the identity information of the person in front of them before they continue discussing personal sensitive information. These are standard procedures to make sure that the person asking for some information or resource really is who he/she claims to be, in order to determine whether they are authorized to get whatever they are requesting. There are many dangers to consider when migrating from human-in-the-loop processes to modernized, fully digital online services. One of the more important is related to the identification and authentication process. How can you trust that the person requesting access to some protected resources online really is who he claims to be? What kind evidence do you need to be able to say: "I believe you - here's your sensitive data" to a person that you cannot see, do not know, and who can sit anywhere in the world? This is where topics on digital identities (section 2.1) and identity management (section 2.2) play a central role.

The essence of information security is to protect assets according to given requirements for confidentiality, integrity and availability (CIA). Confidentiality implies that information should be disclosed only to authorized subjects. Integrity is ensured if only authorized subjects are allowed to modify assets, and availability means that the asset should be accessible or usable to authorized subjects upon request. Since authorization decisions in IT-environments are based on claimed digital identities, we argue that secure management of digital identities is a fundamental prerequisite to ensure information security in every computer environment.

2.1 Digital Identity

Identity is a word, which is interpreted differently by different people, and depending on the context in which it is used. The identity concept can for instance refer to 1) a collective set of characteristics by which a thing is definitely recognizable, 2) the set of behavioral or personal characteristics by which an individual is recognizable as a member of a group, 3) the distinct personality of an individual, or 4) as an expression in mathematics[1]. The origin of the word stems from the late 16th century with the original meaning 'quality of being identical'[2]. With the last few decades' digital evolution, the identity concept has also been adopted by the information and communication technology community (now in the sense digital identity), and also in this community with slight different definitions of the term. In their ISO/IEC 24760-1 standard, the International Standardizations Organization (ISO) defines identity as *"information used to represent an entity in an ICT system"* [11], while the International Telecommunication Union (ITU) defines identity as *"Information about an entity that is sufficient to identify that entity in a particular context"* [8]. In the strictest sense, the latter definition limits the identity to the information that identifies a particular entity, whereas the ISO standard allows more descriptive information about the entity to be included in the concept.

Hansen et al. [34] present a comprehensive view on digital identities, where the entity term from the ISO and ITU definitions translates to human beings. They state that digital identities can consist of one or more attributes, and that each attribute refers to personal data or personal information about an individual. Personal information will be collected and stored in digital registries in various situations throughout a person's life; governments create and store birth certificates, health care systems store medical data, and educational institutions and companies will store personal attributes. Hansen et al. argue that the total sum of all these attributes of personal information constitutes a person's digital identity (consistent with definition 1) above). The subset of identity attributes needed to represent an individual in a specific situation is context dependent. They call these subsets for *partial identities*, and say that an *"individual typically appears under different partial identities for work, other for leisure activities* [...] *or dealing with companies"* (definition 2) above).

This research is focused on an industrial context, and on how professionals can access resources in their working environment. In the context of this thesis we refer to identities as digital information that represent a physical person in an ICT system, and with the clarification of Bertino and Takahashi [23] on the building blocks of the digital identity; A digital identity consists of three key elements: 1) an *identifier* used to identify the owner of the identity 2) *attributes*, which describe different characteristics of, or related to, the identity owner 3) *credentials* which are evidence/data that is used by the identity owner to establish confidence that the person using the identity in the digital world corresponds

[1] http://www.thefreedictionary.com/identity
[2] http://oxforddictionaries.com/definition/english/identity

to the claimed person[3]. This implies that a digital identity can be used identify the person that requests a resource, that the identity claim can be verified through an authentication process involving the credentials, and that authorization decisions can be based on an individuals' identity attributes, such as the role attribute in an role-based access control (RBAC) regime, or the location and nationality in an attribute-based access control (ABAC) regime.

2.2 Identity Management

One of the first pieces of information a new employee is presented with when joining a company is related to how the corporate computers, networks and services can be accessed. A username (identifier) and temporary password (credential) is often delivered in a sealed envelope. This information is then used to log into and access company resources. Prior to the issuance of the credentials, the human resource or IT-department has created a user account for this person, including identity attributes such as, name, address, phone number, organizational role, and of course the mentioned username and password. Most of the personally identifiable information that constitutes a digital identity (or user account) is presented to the organization by the employee itself. Price [61] makes a point of this and says that if the initial verification of the user's identity [and his attributes] is flawed, then the service provision is undermined. Since we argued that digital identities are fundamental prerequisites to ensure information security in computer environments we agree with Price that security must be considered in all activities associated with management of identities. Identity Management comprises the processes an organization must have in place to create, use, update, and revoke digital identities, and the policies that exist to govern each of these activities. The IdM lifecycle is illustrated in Figure 2.1, and also presented by Bertino and Takahashi [23]. The rigor and quality of all steps of the IdM process can vary substantially between different organizations, and this affects the level of trust that can be associated with a digital identity. Dishonest individuals can exploit weaknesses in any of the identity management lifecycle steps to gain unauthorized access to resources, and as such threaten confidentiality, integrity and availability of assets [26] [44]. Insufficient focus on identity management within an organization can also lead to security incidents caused unintentionally by a company's own employees [32].

2.2.1 Digital Identity Creation

The first step of the identity management lifecycle is to create digital identities. In a business case, identity attributes will be *collected* and *registered*, credentials will be *defined*, and finally *delivered* to the user during this process. The rigor in this phase can vary from company to company. The creation phase may include identity proofing practices involving screening and vetting of users [11]. In contrast, many of the popular services on

[3]Thus a hybrid between the ISO and ITU definitions.

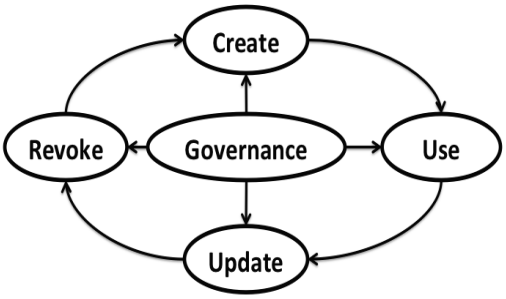

Figure 2.1: Identity Management Lifecycle. Adapted from Bertino and Takahashi [23]

the Internet, such as the social media sites allow users to enter all necessary identity information without further verification. Rigor in the identity creation process is of utmost importance for systems that require a high to moderate security level, i.e., for systems where there is a need to know that the physical person trying to access a system resource correlates with the digital identity provided to the system. The creation process is the foundation for all subsequent use of digital identities. If false or unverified data is entered into the identity management system in this phase, then the system creates a situation in which security is reduced [27].

2.2.2 Use of Digital Identities

Once a digital identity is created and issued, it is time to start using it in electronic transactions. The most common use of digital identities is to use it to authenticate the users by means of the identifier and credentials (ref. section 2.1). The authentication process is central in this respect, and includes two steps [3]:

- Identification, which means that attributes, e.g. identity attributes, are presented to the authentication service.
- Verification, which means that evidence is created to be able to trust the binding between the attribute and the entity it represents.

"Hi, my name is Jostein Jensen" is an everyday example of identification, where I claim to be this person. An example of verification would be to continue the previous conversation by saying: "and here is my passport that proves I am Jostein Jensen", whereby the receiver takes a look at the presented passport and checks to see if there is a match between the photo and the physical person, the claimed name with the name written in the passport, and validity of the passport. The most common authentication example for IT-systems today is the use of usernames and passwords. Users identify by saying: "My username is: jostein.jensen" and "here is my secret password: Password123". Since the user and the service are the only entities that know the password, the service can be confident that the

person asking for access is the same person that the identity was created for.

In addition to user authentication (by means of the identifier and credentials), the other identity attributes in the digital identity can be used for various purposes. RBAC and ABAC schemes can use attributes as input to the authorization decisions in access control modules, and in other scenarios the identity attributes can be used to provide tailored content to the visiting user. In an e-commerce scenario, relevant identity attributes could be a person's name, shipping address, phone number and payment card data. Other attributes would be more relevant in other contexts, such as gender and age on social media networks or organizational role (to determine authorization level) and social security number in a professional context. Digital identities can be created for the purpose of use in one system only, or on multiple services including options to provide single-sign-on.

2.2.3 Update Identity Attributes

Camp [27] divides identity attributes into three different categories [27]: *Persistent attributes* (such as date of birth and eye color), *temporary attributes* (such as employer, organizational role, age), and *long-lived attributes* (such as passport numbers). Long-lived and temporary attributes can and will change over time: employees' role in a company can change, people can move and change address, and credit cards and digital certificates will expire and new ones be obtained and so on. The identity management process must therefore include good procedures to keep identity attributes up to date to ensure their correctness. Identity adjustment, reactivation, maintenance, archive and restore are all activities that are part of the identity update process [11].

2.2.4 Identity Revocation

Identities, including credentials should be revoked if they become obsolete and/or invalid [23]. The ISO/IEC standard 24760-1 [11] separates revocation into identity attribute suspension and identity deletion. The former means that some or all identity attributes are made unavailable so that access rights associated with these attributes are made temporarily unavailable to the user. An example of this can be that the association between a user and a certain group membership is removed to reduce a user's access rights. Another example is the deactivation of all access rights associated with a user. Identity deletion means the complete removal of registered identity information. Information about revocation should be distributed to all relevant stakeholders to ensure that access is not given based on invalid credentials.

2.2.5 Identity Management Governance

There is a need to have policies in place and govern all steps of the identity management lifecycle. Regarding creation of identities, for instance, there should be policies in place

that regulate e.g. who can create identities, how they are created, how the quality of attributes can be assured, how credentials are issued and so on. Flawed and inconsistent procedures throughout the identity lifecycle can make foundation for severe vulnerabilities that can be exploited by attackers to impersonate users, elevate privileges, perform denial of service attacks, and so on [26].

2.3 Identity Management Models

The previous sections on identity management and the identity management lifecycle is primarily concerned with the processes and governance of identity management. The IdM concept on the other hand must also be backed up by architectures and technical solutions. There must be systems in place to register identity information, and there must be mechanisms in place to authenticate users and otherwise take advantage of identity attributes. At an architectural level, one of the decisions that must be made is concerned with how the IdM components are distributed. Ahn et al. [14] and Jøsang and Pope [50] describe that there are three predominant models for IdM:

In *isolated IdM* Ahn et al. explain that businesses form their own identity domains, and have their own way of maintaining identities of users, including employees, customers and partners. That is, each company establishes, uses and maintains a local user repository where credentials are stored and used for authentication purposes to access company internal resources. This example does not exclude single-sign-on to services within a single company. Jøsang and Pope explain the the isolated model from an Internet-based viewpoint, where each service provider online keeps a local user repository, and perform user authentication to authorize access to the own service.

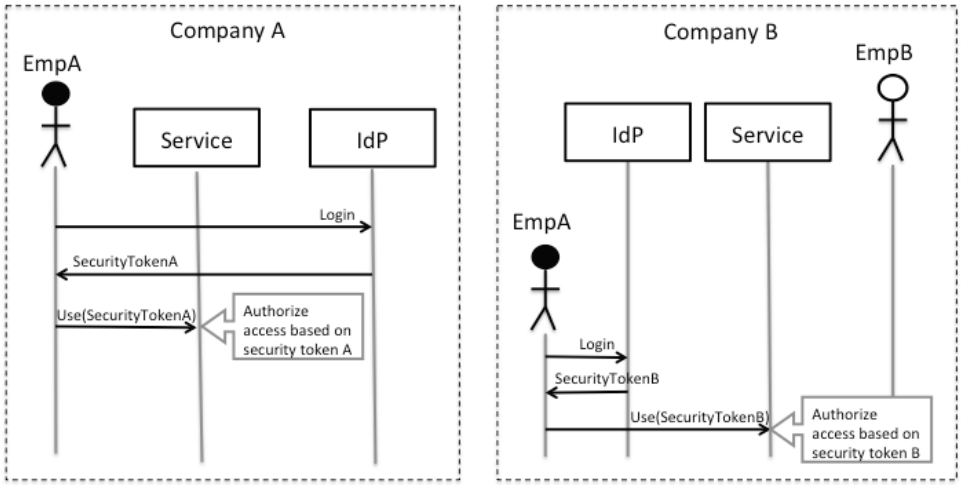

Figure 2.2: Example illustration of isolated IdM

Figure 2.2 shows an example of isolated IdM, which is the most commonly used model today. This example illustrates two collaborating companies where an employee from Company A (EmpA) needs access to resources hosted by Company B. A user profile (digital identity) is created in both Company A and Company B for EmpA. This information is stored and maintained in a company-internal Identity Provider (IdP) service, which is also used to authenticate the users and issue security tokens/security assertions used to prove identity in subsequent service requests. In this model each company is responsible for the identity management lifecycle both for internal employees, and external collaborators that need access.

The *centralized IdM* model involves a single IdP that can be used by several other service providers [50]. The central IdP is responsible for all identity management tasks, such as registering users, issuing credentials, and authenticating users [14]. Identity assertions are issued upon a successful authentication process, and these assertions can then be used to access services outside company borders, or access services distributed across the Internet. IdM lifecycle activities are outsourced to a trusted IdP.

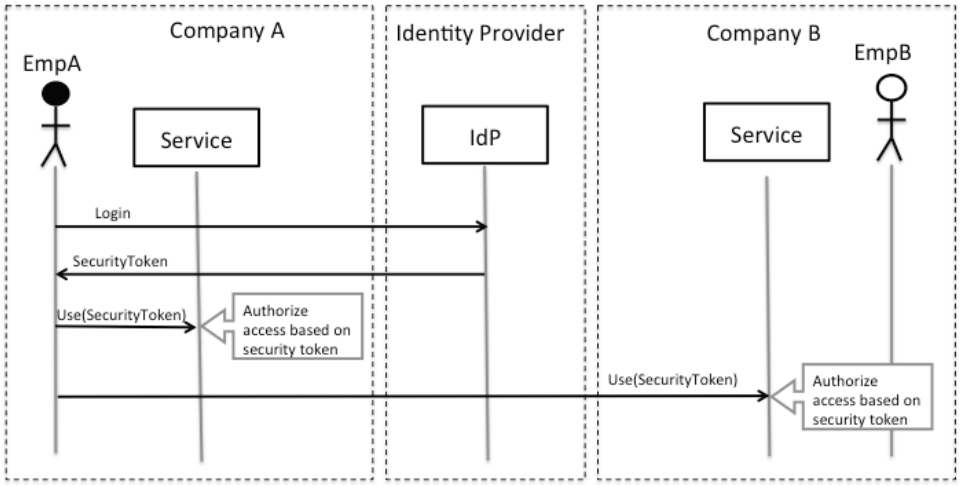

Figure 2.3: Example illustration of centralized IdM

In figure 2.3 we illustrate how a centralized IdM model could be realized. A trusted IdP could host all digital identities within a collaborative environment. Successful authentication towards the IdP would result in a security token that could be used to prove identity towards both internal and external services. In this illustration, all authorization decisions are still being made internally, but based upon the centrally issued security token.

The last model that is described is a *distributed IdM* model. Each company, or service provider on the Internet can keep a local user repository. Users authenticate towards the company or service where they are registered. Successful authentication results in an assertion that can be used to prove identity towards collaborating companies, or online

services. In this case distributed services do not have to manage external users and their digital identities.

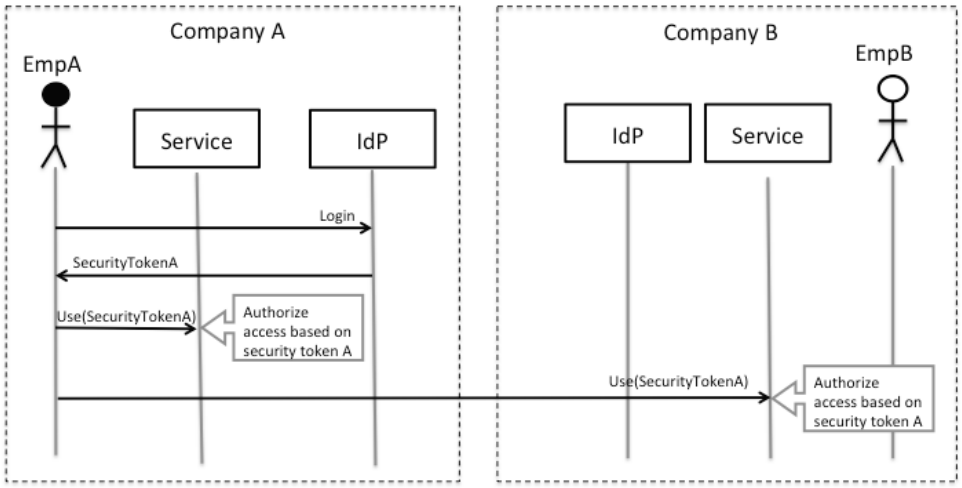

Figure 2.4: Example illustration of distributed IdM

In figure 2.4 a distributed example of IdM is illustrated. Employees are enrolled in an internal IdP within each collaborating company, but the collaborators' services are configured to allow service access based on security tokens issued by external collaborators. All IdM lifecycle activities would be related to a company's own employees alone.

These three models are conceptual models that illustrate different IdM architectures. Hybrid solutions are possible, such as a mix between centralized and distributed models; Each company can maintain an internal user data base and be responsible for most IdM lifecycle activities related to own employees, but where a centralized IdP is responsible user authentication and issuance of security tokens. Such hybrid solutions would require tight integration between local user databases and the centralized IdP.

2.4 Federated Identity Management

Federated Identity Management (FIM) is a concept that allows cooperation on Identity Management across corporate boundaries. Within a federation, partners collaborate on IdM processes, agree on IdM policies, and use interoperable IdM technology. Both the centralized and distributed IdM models presented in the previous section (section 2.3) are examples of technical FIM architectures [14]. The most important contrast to the isolated model, in addition to the technical aspects, is that FIM facilitate identity tasks across security domains [57], i.e. the identity management lifecycle activities are entrusted to an entity outside the local company. Smith emphasizes this and states that *the fundamental*

concept underlying federations is trust." [74], which is also stressed by Chadwick and Inman [29]. Collaborators must rely on each other to only *create* identities for trustworthy employees, *update* attributes whenever changes occur, and *revoke* identities when employees leave the company. Further, the level of trust between federation partners must be sufficient to be willing to exchange identity messages between themselves, including authentication credentials [28] (IdM lifecycle *use* phase).

FIM is considered a promising approach to facilitate secure resource sharing among collaborating partners in heterogeneous IT environments. FIM is about inter-organisation and inter-dependent management of identity information rather than identity solutions for internal use, and that it has emerged with the recognition that individuals frequently move between corporate boundaries [24]. The federation model enables users of one domain to securely access resources of another domain seamlessly, and without the need for redundant user login processes [17].

2.5 Identity Federation Building Blocks

2.5.1 Technological Building Blocks

There are currently two predominant protocols used for building federated identity management systems, which are OpenID and the security assertion markup language (SAML) [57][4]. These two protocols are based on similar ideas with respect to federation of identities and single-sign-on functionality. However, the security requirements driving their designs are quite different.

The original intention behind the OpenID framework was to create a lightweight, decentralized authentication mechanism to avoid blog comment spam [57]. This protocol would save users the effort of going through a new registration process for each social media/blog site where they wanted to participate in discussions [30]. Since its origin in 2005, the OpenID protocol has evolved to become a widely used user-centric online identity solution. The technology is administered by the OpenID Foundation[5]. Internet users can obtain a digital identity by creating a user account at their favourite website that takes advantage of this technology, and the digital identity can then be used to sign into other OpenID enabled web sites, in addition to the site where the user registered. OpenID follows the principles of a distributed IdM model.

Figure 2.5 shows the login screen belonging to the Sourceforge[6] software repository. Users can either create a dedicated user account for this site (isolated IdM model) or log in by means of the OpenID identity provider services offered by Google, Yahoo or AOL.

[4]This paper includes a discussion on a third protocol, the Identity Selector Interoperability Profile underlying Windows Cardspace, which is no longer supported.

[5]http://openid.net

[6]http://sourceforge.net/

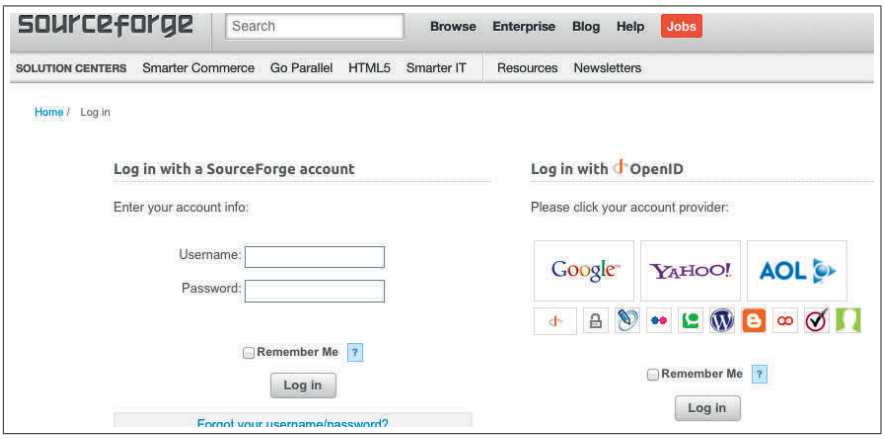

Figure 2.5: OpenID enabled website

The OpenID Foundation provides open source libraries so that the general software developer can realize their own interoperable identity provider solutions, and it provides plugins to popular content management systems, such as Drupal, Wordpress and phpBB.

SAML is a standard defined by OASIS[7], and is an XML-based framework for relaying asserted identity attributes across organizational boundaries [6]. Maler and Reed explain that it is an identity solution that serves several purposes, and that its *"design is driven by strong requirements for true, high-value transactions, and privacy"* [57]. Unlike the user-centric, community-driven development of OpenID, SAML is drafted by industrial companies to meet the goals of professional businesses [30]. OASIS' claimed benefits of taking advantage of the SAML technology include:

- Single-sign-on - authenticate once, receive a security/identity token and reuse it to access resources until its expiry
- Federated identity - use a locally issued security/identity token to access services from collaborating partners
- Attribute based authorization - a SAML assertion may contain attribute information that can be used for authorization decisions.

The SAML standard includes a set of basic concepts. *Assertions* carry information about a user, such as identity attributes, that an asserting entity claims to be true. The *protocols* defined in the standard describes request/response messages used to obtain and relay assertions. *Bindings* describe how SAML protocols can be mapped to existing messaging protocols, such as HTTP or the simple object access protocol (SOAP). *Profiles* describe combinations of assertions, protocols and bindings that can be used to achieve specific business oriented use cases.

[7] http://www.oasis-open.org/

The SAML assertion specification can be used in concert with other technologies, which can substitute the protocol, bindings and profiles defined in SAML. In particular, there are two more standards that are relevant for building business-oriented federated identity management infrastructures, namely WS-Trust and WS-Federation. WS-Trust [4] defines a security token service framework, which describes how security tokens should be requested and issued. The framework is general in that it is not bound to one specific security token type. However, SAML assertions are one of the token candidates. WS-Federation [9] complements the WS-Trust framework, and defines mechanisms to allow federation between different security domains, so that access to resources managed within one security domain can be given based on a security token, e.g., a SAML assertion, issued and managed in a different security domain.

Delft and Oostdijk [30] conclude that OpenID should only be used in situations where the consequences of erroneous authentication are close to zero, which was also the original intention; to provide a good enough solution that fulfils online users' needs with regards to usability, and service providers' needs to verify that there is a correlation between the claimed identity and the person who originally created an identity. SAML, and potentially WS-Trust and WS-Federation should be the preferred choice among the protocols listed here to be used as federated identity management building blocks in business environments where the security requirements are stricter, and where the consequences of erroneous authentications potentially result in breaches on confidentiality, integrity or availability of valuable assets.

2.5.2 Assurance Frameworks

Trust among collaborators is a fundamental prerequisite for the establishment of identity federations [74] [29]. Building a secure and trustworthy IdM infrastructure, e.g., by means of the technologies presented in the previous section (section 2.5.1), may help in trust establishment. However, the presentation of the identity management lifecycle in section 2.2 clearly shows that identity management is more than technology; Identity management is also about processes and policies. The quality and rigor concerning staff vetting and identity issuance during an employment process can vary substantially between different companies. Variations in the quality of authentication mechanisms and strength of authentication credentials will also be found [28]. The same goes for the quality of identity attribute update and revocation procedures, in addition to governance mechanisms in general. A company's level of assurance (LOA) that an external user really is who he/she claims to be in a FIM environment depends on the collaborators' registration processes, and the strength of the authentication process [28]. The rigor in the other IdM lifecycle phases will also affect the assurance level. Common requirements for carrying out each IdM lifecycle phase among federation partners will contribute to increased trust. Identity assurance frameworks are documents that include requirements for each IdM lifecycle phase, and in existing assurance frameworks these requirements are bundled to form identity assurance levels; the higher the assurance level, the stricter

requirements. Information about the assurance level of a digital identity can be used by service providers to determine whether they trust the identity presented to them or not.

The US government defines four identity assurance levels in their framework [10]:

- Level 1: Little or no confidence in the asserted identity's validity
- Level 2: Some confidence in the asserted identity's validity
- Level 3: High confidence in the asserted identity's validity
- Level 4: Very high confidence in the asserted identity's validity

Identities that fulfil requirements at level 1 can be used to access content that has limited concerns regarding confidentiality, integrity and availability, while identities fulfilling level 4 requirements can be used to access assets at the highest classification level. This will balance needs for usability and security. Identity assurance contributes to ensure *confidence in the vetting process used to establish the identity of the individual to whom the credential was issued, and confidence that the individual who uses the credential is the individual to whom the credential was issued* [26]. Consequently, identity assurance is important for the risk management associated with identity management [19].

Examples of identity assurance frameworks can be found within the government sector, such as the Australian National e-Authentication Framework [7], the US government Electronic Authentication Guidelines [26], and the Norwegian Framework for Authentication and Non-Repudiation in electronic Communication with and within the Public Sector [5]. Also the industry supported Kantara initiative has developed an identity assurance framework [75].

2.6 Adoption of Federated Identity Management

There are many examples of federated identity management systems in operation today. Several of these examples can be seen on the Internet, where large service providers, such as Facebook and Google act as identity providers, and allow users to access third party services by means of their Facebook or Google account. These large service providers offer software development kits (SDK) and application programmer interfaces (API) so that application developers (for web and mobile apps) can take advantage of existing authentication mechanisms, and attract new users to their services without new registration processes. Figure 2.5 illustrated one existing example, where the OpenID technology is used to build an identity federation between an online software repository and three other service providers that also act as identity providers. Figure 2.6 illustrates a potential benefit of FIM from a user perspective, through existing federation solutions.

Identity federations are also seen within sectors, such as the government and educational sectors, further illustrated by two Norwegian examples: The Norwegian government defines a roadmap to a more digital society in its report "Digital Agenda for Norway" [1],

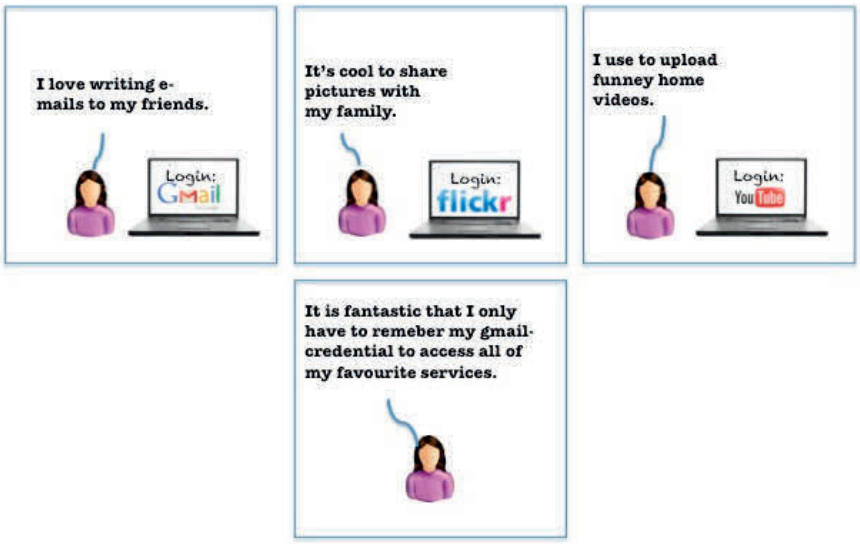

Figure 2.6: Example of online identity federation

where it is recognized that Internet-based communication between the government and its citizens will become increasingly important. Most government agencies provide internet-based services today, many of which allow access to sensitive and personal content, such as personal and professional tax data or personal medical information. Secure use of digital services is of utmost importance. The government has realized that digital identities enable secure use of digital services, and consequently it has developed a FIM solution that can and should be used by all online services offered by the government. The result is that users meet the same familiar login mechanism for all government services. This digital identity solution is called MinID / ID-porten[8]. The solution is built on a centralized IdM model (section 2.3), with SAML (section 2.5.1) as the underlying technology.

The Norwegian educational sector has also chosen to develop an identity federation solution for its users, called FEIDE[9]. System administrators at Norwegian campuses can develop services for their students with low overhead, since authentication is "outsourced" to FEIDE. Students have the benefit of meeting the same login mechanism for the resources they need access to both at their university and at collaborating universities. The FEIDE solution is a hybrid of the centralized and distributed IdM models. User administration, including issuance of user credentials, is performed at the local educational institution. These local user databases are then integrated with the FEIDE solution, which is used as a central authentication point. SAML is used as underlying technology.

Academia has been quite optimistic about FIM technology, and through surveys of academic FIM literature we have seen great effort to move the technology toward academic

[8]http://www.difi.no/digital-forvaltning/id-porten-minid
[9]https://www.feide.no/introducing-feide

perfection [42] [43]. Despite maturing technology and successful examples of identity federations, however, the FIM adoption rate has been slower than expected [74] [55], and focused on specific projects [37]. In the next sections we present our research, which helps gain a better understanding of FIM, and its potential adoption in the Norwegian oil and gas industry.

Chapter 3

Research Context and Design

3.1 Research Context

In section 1.1 we stated that the Norwegian oil and gas industry has migrated towards the concept of Integrated Operations (IO). Before the modernization of the industry 1) daily operational decisions were made offshore, 2) plans were made and changed fragmentally and at fixed times and 3) IT-solutions were specialized and silo-oriented [62]. Around mid 1990 there was an increasing awareness in the industry that only parts of the data collected during oil and gas production was being used, and that this amount of data increased drastically. The idea that these data could be used real-time as decision support and result in better decisions led to a transition to IO. Not only can IO allow decisions to be based on timelier and accurate data, but the integration of offshore and onshore facilities also allow virtual teams to emerge. Integrated operations improve the effectiveness of, and information flows within, the collaborative environment, and among the actors we illustrated in figure 1.1 [2]. Further the IO concept allows several work processes and decisions to be automated, and it allows vendors to deliver their services digitally [62]. The desired effect of IO is to allow better decisions to be made, which consequently should lead to increased profit for the oil companies. The Norwegian Oil and Gas Associations[1] has indeed valued the potential of the IO initiative to NOK 250 billion (net present value) [12].

Integrated Operations has allowed tighter collaboration, through real time data sharing among the actors in the oil and gas industry. The future vision of the industry, however, is not fully reached. Although the IT and process control solutions within a single company have been integrated, there is still work to do to achieve integration across company borders. Today, most IT-systems used in the IO scenario are designed to support intra-organisational collaboration. External collaborators e.g. employees external to the oil company must be given access to the oil company's internal resources. Thus, each company incorporates external collaborators in their identity management regimes. The

[1] previously named Oil Industry Association (OLF)

Norwegian Oil and Gas Association has a plan for the future; To take the industry a step further, and also enable inter-organisational collaboration through collaborative IT-solutions, where partners (Figure 1.1) will share information and knowledge seamlessly across company borders [78].

The migration towards IO is not without challenges, from an information security point of view [62] [40]. Jaatun et al. [40] have pointed to the connectivity between traditional office networks and SCADA systems as a challenge that increases the likelihood of security incidents. Further, in a workshop with industry participants from one of the IO installations they found evidence that *"information security was not satisfactorily integrated in projects and new installations"*, and that *"productivity goals sometimes were prioritized ahead of information security requirements"*. Jaatun et al.'s study conclude there is indication of weak emphasis on information security in the industry as a whole[2]. This is quite a contrast to the last years threat reports from the Norwegian National Security Authority (NSM), the Norwegian Police Security Service (PST) and the Norwegian Intelligence Service[3] where they specifically point to the oil and gas industry as high value targets for attackers, and where the threat agents span from international security and intelligence agencies, via competing industry to hacker groups and private persons.

Information security is key for the integrated operations to be a success in the new threat landscape, and new means to achieve secure access to IT- and process control systems must be devised [2]. Tungland and his colleagues have written a report to describe the reference architecture for next generation of Integrated Operations. Among their suggestions the state that authentication should be carried out at the local company, i.e. the assertion obtained after a successful authentication should be used to access services anywhere within the collaboration [78]. The FIM concept is specifically developed for purposes similar to Tungland's suggestion. It is interesting to learn more about whether a potential technology adoption in the Norwegian oil and gas industry can mitigate some of the new risks that are being introduced with the future generation of integrated operations.

3.2 Research Goal and Questions

Federated identity management as a concept has matured over the last few years, and real life deployments of FIM solutions show the potential of such systems. Consequently, we wanted to learn more about FIM and whether it is a concept suitable for the highly collaborative Norwegian oil & gas industry. As such, we stated the research goal presented in section 1.1

After years of research on secure software development [58] [77] [49] [41] we have learnt

[2]Note that this study was carried out before the Stuxnet, Flame and Duqu virus infections targeted specifically at process control systems in the energy sector, and that these incidents have opened the eyes of the industry to some degree.

[3]The joint 2013 NSM, PST and intelligence service Threat and Vulnerability Report is found at: `http://www.pst.no/media/59018/Trusler_og_sarbarheter_2013.pdf`

that it is hard to develop software where security is an inherent property. Development of secure identity management infrastructures is not different. Security vulnerabilities continue to show up, also in software whose main purpose is to deliver security services. Searches on the National Vulnerability Database[4] using search phrases such as "active directory", "OpenID" and "SAML" (which all are mature IdM related technologies) all return results where related vulnerabilities are rated with high and medium severity, i.e., attackers can potentially exploit these vulnerabilities to achieve their goals. A technical identity management solution with vulnerabilities will lead to a false feeling of being secure. In most cases this is worse than taking a calculated decision, knowing that your assets are at risk. In the early PhD work we intended to build a FIM prototype and test the technology within the oil & gas industry to gain empirical data through prototype testing. This motivated the first research question:

RQ1 What are the main security challenges faced during development of secure identity management solutions for a distributed service platform?

If the main challenges of building a FIM solution are known prior to development, the likelihood of avoiding known pitfalls should increase. The answer to this question is thus a prerequisite to build a secure system that will be accepted for use in an industrial environment.

RQ2 What empirical evidence exists on the development and adoption of FIM in industry?

Hevner and Ram [36] argue that technology being developed as part of a research strategy should be relevant. RQ2 was stated to learn from existing experience, and build a solution based on existing empirical knowledge.

RQ3 What are the benefits and challenges related to FIM from the perspective of the academic community and the industry, respectively?

RQ3 is motivated by our desire to learn from academic security professionals; What do they consider as the benefits of taking advantage of FIM, and which challenges have they identified related to the technology? These benefits and challenges should certainly be taken into consideration by industries investigating the possibilities of adoption FIM solutions. Further, it is highly interesting to know the industry's own perceptions of adopting the technology in order to tailor solutions to meet their expectations.

3.3 The Applied Research Approach

> *No plan of operations extends with certainty beyond the first encounter with the enemy's main strength.*
> (Helmut von Moltke)

[4] http://cve.mitre.org/cve

The results of this PhD project are based on three research methods: design science as presented by Hevner et al. [35] [36], systematic literature reviews as described by Barbara Kitchenham [51] and the case study approach described by Robert K. Yin [80].

3.3.1 Design Science: Study 1

The original plan for this PhD work was to follow Hevner and Ram's [36] framework and seven guidelines for design science throughout the project. *Guideline one* is to develop an artefact. For the purpose of this research the artefact would be an identity management solution that could be experimented with in a limited scenario within the Norwegian oil and gas industry. The *second guideline* is to develop a solution that is relevant to the selected scenario. The *third guideline* Hevner and Ram present is to evaluate the designed artefact. These evaluations can be based on a number of different strategies, such as simulations, analysis, experiments, case studies or mathematical proofs. The original idea was to implement a FIM solution and deploy it in a laboratory environment so that stakeholders in the oil and gas industry could test in a controlled environment. Results about the final FIM solution would be obtained by observation and interviews to catch their perception of the technology. In their *fourth guideline* Hevner and Ram say that design science must lead to clear research contributions, and that the contributions must be verifiable. Further, in *guideline five* they claim that research relies upon rigorous methods in both construction and evaluation. *Guideline six* describes design science as an iterative research approach, and that it is a search process to discover an effective solution to a problem. Hevner and Ram's last advice (*guideline seven*) is to communicate the research, which is a natural part of all PhD works.

The result for this plan would have been a technological FIM solution that had matured over a number of iterations, and which had been based on stakeholders' perception of the technology. The stakeholder perception of FIM concepts would also be an important contribution in this approach. As mentioned, this was the basis for the original plan, however, things changed after the first design science iteration. The artefact subject to first iteration of the design science approach was the Identity Management functionality developed as part of a healthcare platform within the MPOWER project. In this first iteration we focused on evaluating the security properties of the platform, to find its weaknesses before continuing to adapt it from a healthcare scenario to the industrial oil and gas case. We wanted to learn more about the challenges of implementing a secure identity management solution in a distributed environment, and as such provide answers to RQ1. Later iterations would contribute to also answering RQ3.

The evaluation (guideline three) of the original artefact (guideline one) showed that the selected IdM solution contained severe vulnerabilities. In our search to discover a more effective solution (guideline six) to improve the IdM artefact, we chose to carry out a systematic literature review (section 3.3.2) to see if a model driven approach to software development would be a more suitable way to convert the security design of the IdM solution into secure code. Before the second design science iteration we also wanted

to study academic literature to improve the relevance of the technical solution (guideline two). Existing empirical evidence could guide us in the design towards a solution targeted at the industry (RQ2), and insight into academic work on federated identity management (RQ3) would inform us how to improve the existing solution. Again, a systematic review was the selected approach. In both cases the approach was selected to fulfil the design science requirements for rigorous research (guideline five).

Hevner and March claim that *"purposeful artefacts are built to address unsolved problems"* [35]. After initial discussions with stakeholders in the oil and gas industry, we realized that there are many unsolved challenges with respect to Identity Management within the industry. FIM could as such be a purposeful artefact, however, the meeting with industry also revealed that the original planned academic approach was a bit naive. The oil and gas environment and its collaborative nature was far more complex than previously imagined. It would not be realistic to develop a FIM artefact that could be tested in an operative setting, at least not within the reach of a one person PhD project. The applicability of laboratory results would be hard to generalize to a real life setting. Thus, lacking an existing published academic knowledge base and understanding of the oil and gas environment with respect to information security and identity management, we elected to deviate from the original research plan. Instead of continuing the artefact development in a second iteration, we chose to focus effort on creating a knowledge base about the collaborative oil and gas environment through a case study (section 3.3.3). This knowledge should then be used as input by other design science researchers, or FIM developers, to improve new or existing FIM artefacts to meet industry expectations.

To sum up; Our design science approach includes the first iteration of artefact development and evaluation. The approach led to two structured literature reviews, and a case study to create a knowledge base that can be used by later iteration FIM design science.

Further information about the artefact, and the evaluation and analysis of results can be found in Paper 1 [47] and in two of the supporting papers [49] [66].

3.3.2 Systematic Literature Review: Review 1 and Review 2

Systematic literature reviews are forms of secondary research studies, where existing research about a topic is identified and synthesized in an objective manner. Barbara Kitchenham has developed guidelines for performing systematic literature reviews within the software engineering field [51]. There are several reasons for performing a systematic literature review. Kitchenham points to the most common:

- To summarize existing evidence within a research field
- To identify gaps and propose further research
- To provide background to position new research initiatives

Table 3.1: The systematic literature review process

Planning	Identification of the need for a review Commissioning a review Specifying the research question(s) Developing a review protocol Evaluating the review protocol
Conducting	Identification of research Selection of primary studies Study quality assessment Data extraction and monitoring Data synthesis
Reporting	Specifying dissemitainon mechanisms Formatting the main report Evaluating the report

Kitchenham argues that there are many advantages of systematic reviews; Rigorous methods leads to less bias, information about the effects of a phenomenon can be analyzed across a wide range of settings and empirical methods, and data from different studies can be combined using meta-analytic techniques. The major disadvantage, however, is that systematic reviews require considerable more effort than traditional literature reviews[5].

The systematic literature review involves three main phases in the review process. Table 3.1 is an overview of these phases and their sub-activities. Further descriptions of the guideline is found in Kitchenham's report [51].

Review 1 and review 2 followed the guidelines for systematic review. The need for reviews was motivated by the design science process as mentioned in the previous section. In review 1 we were looking for evidence that model driven development could be used as development strategy for building a more secure identity management artefact to provide evidence regarding RQ2. In review 2 the use of systematic review was triggered by our aim to learn more about federated identity management; the benefits and challenges of adopting such technology (RQ3), and industrial experiences of taking advantage of such technology (RQ2), which would be used to increase the relevance of the identity management artefact. All of the three bullet points above were thus drivers for this research approach.

Systematic mapping studies are another type of review that are complimentary to systematic literature reviews [51]. They have broader research questions, search terms will be less highly focused, the data extraction process is broader, and the analysis stage is about summarizing the data to answer the research question posed without going into in-depth analysis techniques such as meta-analysis and narrative synthesis. Kitchenham further claims that systematic mapping studies are appropriate in situations where very little evidence exists. In our two review studies we found little or no empirical evidence. The

[5]This statement is confirmed through review 1 and review 2 in this PhD thesis.

outcomes of the studies are papers giving overviews of the research fields. As such, one can argue that our two literature reviews can be categorized as systematic mapping studies rather than systematic literature reviews, but where a rigorous systematic review approach has been followed to allow replication of the studies.

The starting point for both reviews was a research protocol where the research questions and the search strategy were defined. To support the paper selection process, the protocols also specify inclusion and exclusion criteria. A rigorous and comprehensive search is key to identify all the relevant scientific literature. Both sources for scientific literature and search phrases were specified prior to the search. We used four online databases for scientific literature to search for studies, for the purpose of review 1:

- IEEE Xplore[6]
- ACM Digital Library[7]
- ISI Web of Knowledge[8]
- Compendex[9]

According to experiences made by Dybå et al. [31], the use of these databases should be sufficient to find all relevant literature within the information systems field. The use of other databases will lead to duplicate findings, and as such, lead to extra work. However, we kept this list and added SpringerLink[10] as source in review 2 as a special precaution to not miss important literature.

Our search strategy identified a large number of papers in both reviews. All references and abstracts were imported to the reference tool EndNote, which helped us eliminate duplicate findings from the reference database, and to filter papers based on title and abstracts. The remaining papers were read in full and sorted according to the inclusion and exclusion criteria.

Once we had our primary studies we continued to analyze the papers. In review 1 on model driven development, we grouped all papers treating the same research initiatives and wrote a narrative for each major initiative. The lack of empirical research eliminated the need to do a further analysis to answer the research questions. In review 2 we printed out all primary studies, and color coded academics' statements about benefits and challenges of FIM. Coded quotes were then copied into the tool MindManager, and grouped into similar concepts. This approach shares similarities with the constant comparison method [69], but was a result of a pragmatic approach to the analysis phase, rather than a deliberate act to follow the more rigorous constant comparison approach. This pragmatic and highly manual approach provided insight into qualitative data analysis, which was used as foundation to learn and improve practices in study 2 (section 3.3.3).

[6] http://ieeexplore.ieee.org/Xplore/dynhome.jsp
[7] http://portal.acm.org/dl.cfm
[8] http://apps.isiknowledge.com
[9] http://www.engineeringvillage2.org/
[10] http://www.springerlink.com

The model driven development review, including further details, is reported in Paper 2 [45], while the FIM review including details was reported in Paper 3 [42] and Paper 4 [43].

3.3.3 Case Study Research: Study 2

With the recognition that there was a lack of empirical evidence related to industrial adoption of FIM, and that there was a need to develop more knowledge concerning the industrial oil and gas environment, we selected to deviate from the original design science approach. Hevner and March [35] claim that a design science artefact that solves a non-existing problem is of equally low importance to the research community as a theory produced by behavioral-science research that is not useful for the environment. We did not want to fall into any of these traps, and consequently we elected to follow an empirical research approach to build new and needed knowledge about current identity management challenges, and the perception of federated identity management within the oil and gas industry. This decision, however, affected the original design science research plan, as it did not allow for more iterations beyond Study 1 within the timeframe of this PhD project.

Robson [64] gives advice on the selection of empirical research strategies. Our research were going to be exploratory, since our aim was to develop new knowledge about a topic that was not well documented, nor understood within the industrial research context. Consequently, a flexible (qualitative) design would be appropriate. Robson mentions case studies, ethnographic studies and grounded theory studies as possible flexible research strategies. FIM was not implemented in the context we were investigating, so ethnographic studies including participant observation were out of the question. Klein and Myers argue that interpretive research attempts to understand a phenomena through people's perception of them, and that interpretive research methods are designed to produce an understanding of an information system's context [52]. Both case studies and grounded theory could have been feasible approaches, but the choice fell on a case study strategy based on the work of Robert K. Yin [80], with inspiration from Klein and Myers' principles for conducting and evaluating interpretive field studies[11] [52].

Robson gives the following definition of case study research [64][12]: *"Case study is a strategy for doing research which involves an empirical investigation of a particular contemporary phenomenon within its real life context using multiple sources of evidence."* It is emphasized that case study is a strategy rather than a mere data collection method such as observation or interviews. Yin [80] presents a rigorous approach for carrying out case studies, which involves six phases: 1) plan, 2) design, 3) prepare, 4) collect, 5) analyze, 6) share. The approach is linear, but at the same time iterative.

In the *planning* activity we stated the following research questions: What are the industrial expectations regarding federated identity management? How can the industry benefit

[11] Klein and Myers state that field studies are either in-depth case studies or ethnographies.
[12] The definition is inspired by Yin [80].

from FIM and which challenges will the industry face during adoption of this technology? These questions contribute to answering this thesis' research question RQ3.

The study is based on an embedded single case *design*, where we look at FIM in relation to the industrial collaboration within the Norwegian oil and gas industry. Our embedded units of analysis include the most interesting stakeholders within this domain, including one oil and gas company, two vendors of production equipment and process control systems, one consultancy, and one service company. Informants from each of these stakeholders were included in the study, and a total of eleven informants were selected based on our industry contacts' recommendations.

The case study and the research protocol was *prepared* together with co-supervisor Martin G. Jaatun, as part of a PhD course on empirical software engineering, whose aim was to teach research methods applicable to the software engineering field. The protocol was presented in this course and improved based on feedback from the class. The protocol included the research questions, data collection plan including an interview guide and a data analysis plan. A pilot test of the interview guide was carried out to further improve the guide prior to data collection.

The data *collection* was based on interviews. Multiple sources of evidence should ideally be obtained and triangulated to strengthen the empirical evidence in case study research [80]. However, since our study was of an exploratory, interpretative nature and targeted at perception of technology that had not already been implemented, there would not be other types of relevant evidence available within this context. Each in-depth interview lasted between 45 minutes and one hour. We selected a semi-structured approach where an interview guide was used to obtain the same information from each participant, but where additional follow-up questions were asked to elaborate on interesting points. The interviews were held in person and via telephone, and were recorded to allow for full transcription.

The transcribed interviews were *analyzed* using the constant comparison method [69]. Data were first coded into three preformed categories: perceived benefits of FIM, perceived challenges related to FIM, and current security challenges. The two first categories come directly from the research questions of the study and RQ3 of the thesis. Data sorting into each of these high-level categories were then subject to new iterations of coding (postformed) to discover common concepts in the interview material within each category. The qualitative analysis tool Nvivo was used to support the coding and analysis process.

Results from the study was *shared* with the research community through Paper 5 [46] and Paper 6 [48].

Chapter 4

Results

This thesis consists of seven selected papers that highlight different aspects of federated identity management in collaborative industrial environments. The knowledge shared to through these papers is synthesized in this chapter. The following sections represent each of the three research questions stated in section 3.2

Table 4.1: Relations between key findings, research questions and papers

No	Key finding	RQ	Paper
1	IdM security specifications alone do not guarantee secure applications.	1	P1, SP1, SP2
2	Identity assertions must be properly protected.	1	P1
3	Secure identity management is more than secure technology.	1	P7
4	Little empirical evidence exists on the development and adoption of federated identity management.	2	P2, P3, P4
5	Academics expect increased privacy, security and usability for end users.	3	P3
6	Academics expect businesses to benefit from reduced administrative cost and complexity, improved data quality and security, and easier co-operation.	3	P3
7	Academics expect technical challenges related to interoperability, attribute synchronization and consistency, revocation and identity provider discovery.	3	P4
8	Academics expect organizational challenges related to investment cost, liability issues, identity assurance, security, knowledge and trust.	3	P4
9	Practitioners perceive that FIM will improve user administration and usability, make collaboration more efficient, reduce cost, facilitate audit and lead to better protection.	3	P5
10	Practitioners expect trust issues, technological challenges, investment cost and security challenges to be obstacles to adoption of FIM in industry, and there is a risk that they confuse identity management with access management.	3	P5
11	Practitioners question whether there is sufficient organizational maturity to adopt new identity management solutions.	3	P5, P6

4.1 Security Challenges Related to Identity Management Development (RQ1)

Key finding 1: IdM security specifications alone do not guarantee secure applications

Our first study (Study 1) was centered around a healthcare service platform. Its identity and access management services were of direct interest to this PhD project, as described in section 3.3.1. A thorough process to document security requirements and develop a secure design was carried out prior to implementation and coding of the services. In supporting paper 1 [49], we document the extensive work that was carried out in relation to reviewing existing laws and translating these to technical security requirements to ensure compatibility between the two. Several of the requirements are directly relevant to the identity and access management solution. In table 4.2 we list the most relevant security requirements that were elicited in this regard, and give a short description of how they are relevant for FIM solutions.

In paper 1 [47] we document security challenges related to distributed service environments. Further, we give a brief introduction to the design of the identity management modules in the healthcare platform[1], before we describe how we tested the security properties of the implemented solution.

Paper 1 clearly shows a mismatch between the security requirements, the security design and the implemented solution. This situation is also confirmed by a second and more extensive security test of the platform presented in Supporting paper 2 [66]. Both tests revealed severe security vulnerabilities. In supporting paper 2 we argue that there is a need to focus on security throughout the software development lifecycle. The security work is not done when the security specification is delivered. Further, we conclude that *"the main lesson learned is that it is necessary that every person involved in such a project* [development project] *is aware of the consequences of not thinking about, implementing and testing security from the beginning. Only then will it be possible to achieve more secure systems".* [66]

Key finding 2: Identity assertions must be properly protected

Federated identity management is about inter-organizational collaboration, and sharing of identity data across organizational boundaries and security domains. Internet is one of the communication means that can be used to relay identity information. It is difficult to control how digital information is being used behind firewalls, inside the trusted organizational premises. Once information has left the security domain, this control goes from difficult to impossible, unless proper security controls are applied to the information. Confidentiality and integrity controls become extremely important to prevent unauthorized access to, or modification, of data. Identity assertions are used to prove identity at an

[1] a design based on the security requirements specified in Supporting Paper 1.

Table 4.2: Security requirements from privacy legislation and their relevance to FIM

Security requirements from [49]	**Relevance to FIM solutions**
Services should identify and verify the identity of their human users before allowing them access to their resources.	Services that provide access to assets that need protection should utilize access control mechanisms. A FIM solution could be the choice to achieve identification and verification (i.e. authentication) in a distributed environment
Services should identify and verify the identity of corresponding services before they are allowed to communicate.	Identity management should not only be associated with human users, but also services and equipment. If not, attackers can spoof other services or equipment to obtain service access.
Services should verify the authorization level of users before access to sensitive data can be given.	The FIM identity assertions will likely be used in the authorization process in distributed FIM environments.
The platform should be able to detect unauthorized manipulation of data that is being transmitted.	FIM assertions contain information about the authentication status, and possibly identity attributes that can be used to take authorization decisions. As such, it is of utmost importance to hinder unauthorized manipulation of these to prevent opportunities for users to elevate their privileges.
The platform should be able to log security incidents, such as failed login attempts or unauthorized access attempts to services in order to discover and trace system abuse.	Information sent in the FIM identity assertion, along with the service requests, can facilitate logging and tracing of individuals' activities.
Data freshness should be controlled to prevent chances of replay attacks.	Authorization decisions are based on the FIM identity assertions. If the right precautions are not taken, a "man-in-the-middle" can copy these assertions and resend them at a later time to obtain service access.

external service, without the need to do a local authentication. Further, identity attributes may be used as input to role based or attribute based authorization decisions. Adversaries can craft their own identity assertions, or modify existing ones if identity assertions are not properly protected. This is why security requirements like the ones presented in table 4.2 are particularly important in federated environments. In Paper 1 [47] we demonstrated several challenges related to unprotected identity assertions:

- Lack of confidentiality protection allowed us to inspect all identity attributes. This may reveal sensitive identity data.
- Lack of integrity protection allowed us to modify identity data. Lack of integrity protection means that adversaries may be allowed to elevate or remove access privileges associated with an identity.
- Identity assertions were vulnerable to replay attacks. We could record the data flows between to entities, including identity data. The identity data could be resent at a later time, and successfully be used as authentication proof since there were no mechanisms to control its freshness.

In addition to this we verified whether entity authentication was implemented at a service level. Since this was not the case we were allowed to craft service requests directly towards the identity assertion service (security token service), and as such omit the whole authentication process. One can argue that the application we tested was not representative for real production solutions, however, it illustrates how bad things can go if the identity management solutions contain errors and security vulnerabilities. Implementation errors are made, also in operational systems: CVE-2008-3891 is one example of a high severity vulnerability that was discovered in relation to the SAML identity assertion issued by one of the larger service providers on the Internet. Before the vulnerability was corrected, it allowed *"remote service providers to impersonate users at arbitrary service providers"*[2].

The use of identity assertions in FIM environments make them as attractive and vulnerable to theft and exploitation as traditional user credentials, such as usernames and passwords. With a wider adoption of federated identity management solutions, we may in the future experience identity assertion phishing in the same scale as that of traditional phishing and spear phishing for usernames and passwords. Further, if identity assertions can be exploited, then the whole authentication process is circumvented, no matter how many authentication factors are used or what strength of authentication mechanisms is employed.

Key finding 3: Secure identity management is more than secure technology

The first two key findings highlight the importance of a solid, high quality development strategy for the technical identity management platform, and with a special focus on securing the identity assertions. Identity management, however, is more than technology, as we wrote in section 2.2. In Paper 7 [44] we carried out threat modeling of each identity management lifecycle phase. For each phase we listed the most likely threats. The conclusion is that if one activity carried out in one of the lifecycle phases is flawed, then the assets being protected by the identity management solution is at risk. Following is a summary of the associated threats, as reported in Paper 7. The figures represent threat models, using the misuse case diagram notation of Sindre and Opdahl [72].

[2] http://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2008-3891&cid=2

There are numerous motives for attackers to somehow manipulate the **identity creation** process, one of which is to assume the identity of another person during the establishment of a digital identity. This can, e.g., be done as shown in Figure 4.1 by presenting forged identity information (e.g. false passport) during the identity proofing process, or exploit the fact that identity proofing is not operationalized in the creation process. University enrollment under a fake alias, establishment of credit cards or establishment of phone subscriptions in another person's name are examples of this threat. The consequence of this is that the attacker obtains full access to resources by means of a valid user identity.

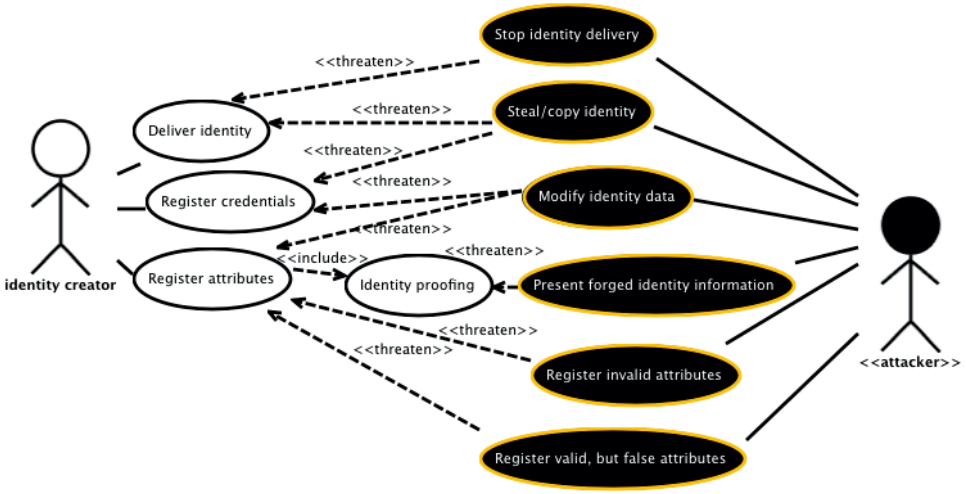

Figure 4.1: Threats to the identity creation phase

Invalid attributes can be inserted in the user database, attributes can be modified by unauthorized entities, or valid, but false attributes can be registered during the attribute registration if proper countermeasures against these threats are not in place. The effects of this must be seen in light of specific system implementations, and an analysis of how the attributes are being used. Still, we must acknowledge that these threats can have serious consequences knowing that attributes can be used to determine access level e.g. based on group memberships/roles in role based access control (RBAC) schemes or possibly any other attribute in attribute-based access control (ABAC) schemes. Attackers can obtain elevated privileges for their own valid user account by manipulating attributes. The credential registration process must also be protected so that attackers cannot steal or copy credentials, such as username password pairs. If attackers get access to valid credentials, they can impersonate valid users to obtain protected information. If the attackers modify credentials the result can be that they alone possess the valid credentials to access resources, or that they block access for another user. These challenges also exist during delivery: Attackers can obtain access to digital identities, which can be used in subsequent malicious activities by intercepting the communication channel used to deliver the credentials, such as mail or e-mail.

Most users of IT systems recognize that there is risk involved in the use of online services, especially related to **use of digital identities**. Figure 4.2 illustrates typical threats related to use of digital identities. Access credentials can be lost or stolen so that attackers can authenticate, and thereby impersonate valid users. There are many attack vectors used to obtain valid credentials. Communication lines can be intercepted to copy plaintext data, password files can be stolen and decrypted, social engineering can be used to trick users into giving away their credentials, and so on. The introduction of SSO and federated SSO has added to this complexity in that security assertions are issued based on a successful authentication. This security assertion is stored by the client and used as proof of identity in subsequent service requests. This means that an attacker can copy this assertion and add it to malicious service requests, or replay previously sent messages. If the receiving service trusts this assertion, it will provide information as requested. Since authentication data (assertions) are shared across services in SSO, and across services within different trust domains in federated SSO, the attack surface in weakly designed systems is greatly increased compared to having separate systems.

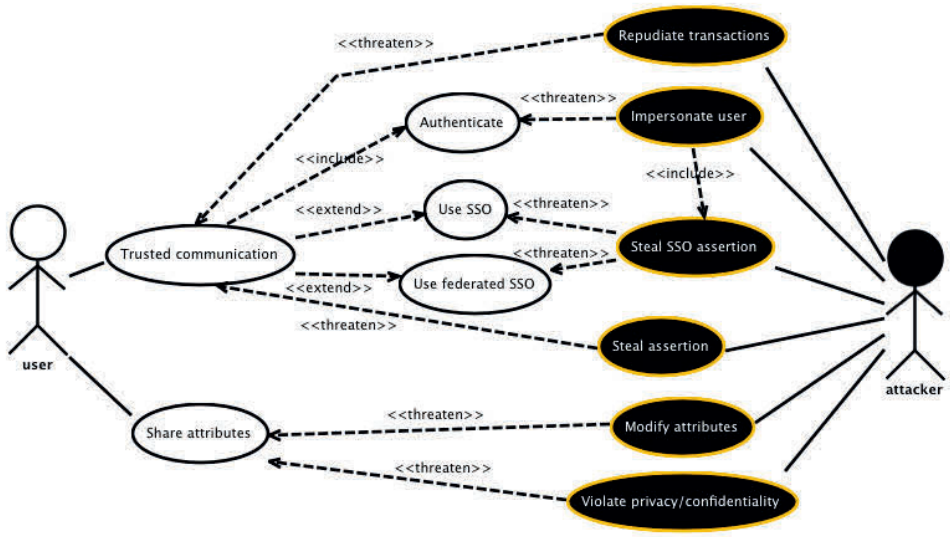

Figure 4.2: Threats to the identity use phase

In addition to these impersonation threats, there is always a probability that one of the entities in a trusted communication repudiates that a given transaction has taken place. In an e-commerce context, users can deny that they have placed an order, or deny that they have accepted terms of use if the service providers do not take precautions. Similarly, service providers can deny that they have received information or requests from users.

As already mentioned, RBAC and ABAC models allow taking access control decisions based on identity attributes. If attackers can modify attributes during transmission, they

can be allowed to elevate their privileges by manipulating attributes. Another scenario is that attackers modify, e.g., shipping address so that ordered goods is paid by one user, but is sent to the attacker's destination. The disclosure of identity attributes may also violate user's privacy preferences, or reveal company internal information.

Figure 4.3 illustrates typical threats to the **update phase**. Similar threats can be found in the create and update phase: credentials can be copied or stolen and false attributes can be provided. It is still interesting to treat them separately. In operative environments one can experience that the responsibility for identity creation and identity update are placed at different levels in the organization. While the human resource department may be responsible for creation of user identities, e.g., in connection with a new employment, the responsibility for updating user profiles may lie with the IT-support department. Consequently, attackers can approach different parts of an organization to achieve the same goals.

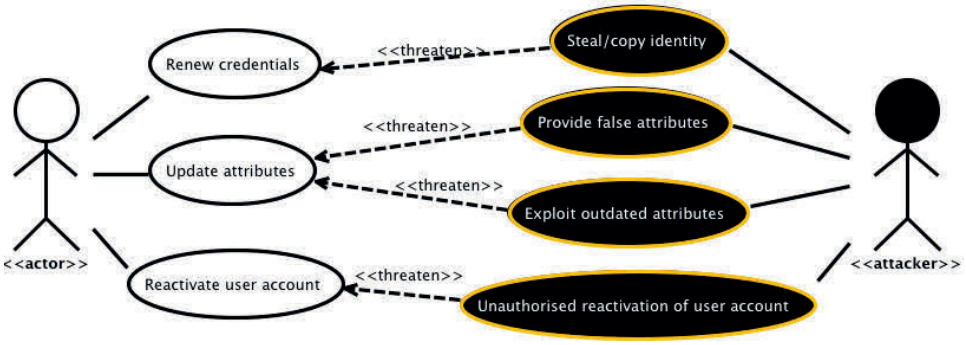

Figure 4.3: Threats to the update phase

It is also important to note that attackers can exploit weaknesses in the update procedures. Delays in the update procedure can allow users to access content based on old but still valid access credentials and attributes, and attacks towards update management interfaces can allow unauthorized reactivation of user accounts.

For the **revocation phase** it is hard to see that suspension and deletion of identity information can be misused otherwise than to block access to resources, i.e. to deny service to authorizes users. The confidentiality and integrity of assets will as such be maintained, but if resource availability is a critical success factor, then this kind of denial of service attack can have major consequences.

Figure 4.4 also indicates that insufficient distribution of revocation lists, e.g., in distributed systems can be exploited by attackers. Distributed services without an updated access revocation list can continue to grant access, even after the credentials have been revoked officially.

Figure 4.5 provides an overview of threats that are related to **identity management gov-**

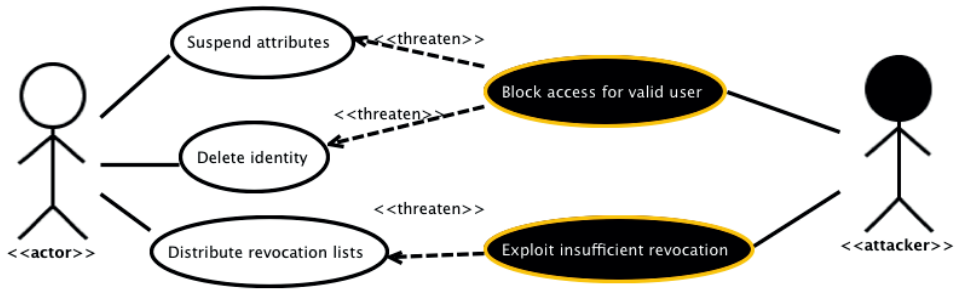

Figure 4.4: Threats to the revocation phase

ernance. Password policies are among the policies that affect all phases of the identity management lifecycle, so let's use this as an example. The password policy should include requirements for password length, complexity and validity period. It should also contain statements about the password being personal, and that it should not be shared with anyone.

As indicated in the figure, attackers can exploit weak or non-existent polices throughout the phases of the identity management lifecycle. If a password policy is non-existent or weak, then users can associate their digital identities with insecure passwords. This affects the use of the digital identity. Since the password is weak and can easily be hacked, e.g., through brute force attacks or guessing attacks, then there is a challenge to establish trust that the the entity requesting a service is who he claims to be. Non-existent or poor requirements for password change (update) and revocation also affect the trustworthiness of credentials. With infinite password lifetime, attackers can exploit compromised credentials as long as the user account is active.

Policy incompliance means that policies exist, but that they are not complied to, e.g., due to lack of policy enforcement. It does not help to have password length and complexity requirements if the technical platform still allows users to select shorter and weaker passwords. Further, many users will continue to reuse their passwords after expiry, despite a policy stating that passwords are valid for 90 days and that reuse is not allowed.

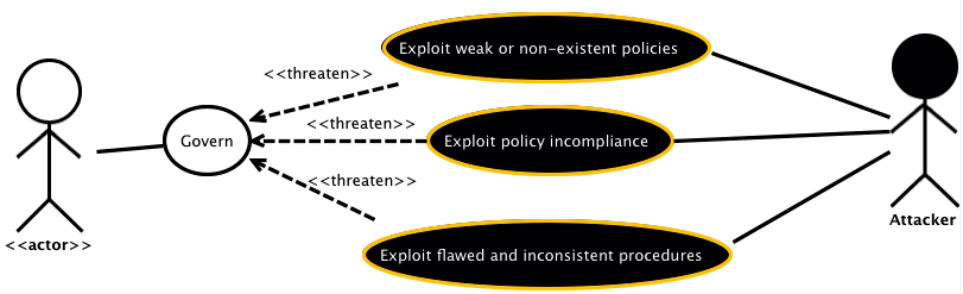

Figure 4.5: Threats to the identity management governance

Even if policies are in place and compliance is achieved throughout the organization, attackers can exploit flawed and inconsistent identity management procedures.

Source Papers

- Paper 1: Jostein Jensen, Åsmund Ahlmann Nyre: **"SOA Security - an experience report"**, Proc. The Norwegian Information Security Conference (NISK), 185-196, 2009
- Paper 7: Jostein Jensen, **"Identity Management Lifecycle - Exemplifying the need for Holistic Identity Assurance Frameworks"**. Information and Communicatiaon Technology, volume 7804 of Lecture Notes in Computer Science, pages 343-352. Springer, 2013
- Supporting paper 1: Jostein Jensen, Inger Anne Tøndel, Martin Gilje Jaatun, Per Håkon Meland, Herbjørn Andresen, **Reusable Security Requirements for Healthcare Applications**, International Conference on Availability, Reliability and Security (ARES), 380 - 385, 2009.
- Supporting paper 2: Richard Sassoon, Martin Gilje Jaatun, Jensen, Jostein, **The Road to Hell is Paved with Good Intentions: A Story of (In)secure Software Development**, International Conference on Availability, Reliability, and Security (ARES), 501 - 506, 2010

4.2 Empirical Evidence on FIM Development and Adoption (RQ2)

Key finding 4: Little empirical evidence exists on the development and adoption of federated identity management

Within this PhD project we have carried out two systematic reviews[3]. One was related to the use of model driven development (MDD) to implement more secure software, and the other to learn more about various aspects of federated identity management. Both studies included research questions to evaluate empirical research within these fields. Our first search[4] in the MDD study reported in Paper 2returned 2844 titles to be evaluated. A follow-up search a year later[5] returned yet another 27 titles. The exclusion of papers based on irrelevant titles left us with 366 papers for which we also read the abstract. 122 (plus the 27 from round two) papers were studied in full on the security in MDD topic.

[3]that turned out to be systematic mapping studies following the search and selection rigor of a systematic review, ref. our discussion in section 3.3.2

[4]First search carried out March, 2010.

[5]Second search carried out June, 2011.

We might have lost some empirical papers in the first filter process due to misleading titles. Further, we might have overseen some empirical work if the abstracts were not clear on the contributions. However, the rigor in the search and filtering process allow us to claim with high confidence that most research is based on conceptual analysis, rather than empirical evaluations of the technology.

Paper 3 and Paper 4 describe the search process for papers on federated identity management. Our initial search resulted in 684 titles. A filtering based on title and abstract left us with 113 papers that were read in full. Paper 3 includes 30 primary studies reporting the benefits of FIM, while Paper 4 includes 29 primary studies that indicate academics' views on challenges related to FIM. Again, none of the included studies presented empirical evidence. This was a bit surprising, keeping in mind that there are existing federated identity management solutions in operation, as we presented in section 2.6. One can argue that our search strategy prevented us from finding all published work related to federated identity management. We restricted our search to the phrase "federated identity management". As such, we might have overseen studies where they used other terms for the same concept, for instance identity federation or inter-organizational single-sign-on. However, in our later work we have not seen other empirical studies within this field. This make us claim with moderate to high confidence that empirical research related to federated identity management is sparse.

Source Papers

- Paper 2: Jostein Jensen, Martin Gilje Jaatun: **"Not Ready for Prime Time: A Survey on Security in Model Driven Development"**, International Journal of Secure Software Engineering, 49-61, 2011

- Paper 3: Jostein Jensen: **"Benefits of Federated Identity Management - A Survey from an Integrated Operations Viewpoint"**, Availability, Reliability and Security for Business, Enterprise and Health Information Systems, volume 6908 of Lecture Notes in Computer Science, 1-12, 2011

- Paper 4: Jostein Jensen: **"Federated Identity Management Challenges"**, Proc. Seventh International Conference on Availability, Reliability and Security (ARES '12), 230-235, 2012

4.3 Benefits and Challenges of FIM Adoption (RQ3)

4.3.1 From Academics' Viewpoint

Key finding 5: Academics expect increased privacy, security and usability for end users

In Paper 3 we show that academics have been quite optimistic about the benefits of adopting FIM solutions. From a user perspective academics claim that one should experience increased privacy protection [22] [54] by having more control over own identity attributes. Users' security should be increased [56] [22] since they only have to remember one, or at least a very limited set of username/password pairs, and as such should be able to use and remember stronger passwords. Single-sign-on and seamless access to resources should improve usability [56] [57] aspects.

Key findings 6: Academics expect businesses to benefit from reduced administrative cost and complexity, improved data quality and security, and easier cooperation

Paper 3 also presents benefits from a business side; Companies can experience **reduced cost** [74] [22] [28] related to identity management tasks by avoiding duplication of identity management efforts among the federation partners. Improved data quality [22] [33] can be experienced because user data is stored and maintained at one site, avoiding synchronisation issues. Increased security [22] [63] [18] can be achieved with FIM since security principles such as avoiding single-point-of-failure and to achieve minimal disclosure of data can be fulfilled, and fine-grained access control can be realised. Service providers can experience reduced complexity [15] by outsourcing identity tasks to specialised identity providers. Finally, FIM promises to facilitate cooperation [22] [70] among federation partners where cross-domain single-sign-on and seamless service access across company boundaries can be a result.

Key finding 7: Academics expect technical challenges related to interoperability, attribute synchronization and consistency, revocation and identity provider discovery

Key finding 5 and 6 show that the academic community has high expectations with regard to the benefits of FIM. However, the road towards federated identity environments is full of potholes, as we explain in Paper 4. Interoperability issues are among the technical challenges that must be overcome. FIM will fit well into homogeneous environments where all collaborators use the same standard [79]. Most inter-organizational collaborative environments, however, are heterogeneous in nature, which means that different organizations adhere to different standards, tools and procedures. This increases the complexity of FIM adoption. Even in homogeneous environments there might be challenges. Vendors of FIM

solutions can claim to follow a recognized standard, but due to the number of protocol options in the standard and conformance variations among the vendors, there might still be integration challenges [57]. A prerequisites for interoperable FIM solutions is a common agreement regarding the identity attributes to be used, and their semantics [76] [22]. There are also other technical challenges to consider in federated identity environments. FIM solutions must be designed to avoid inconsistencies in identity attributes within the federation, and proper mechanisms must be developed to properly synchronize replicated identity information [38] [39] [67]. Other technical challenges are concerned with how the identity provider is discovered [57], especially in environments where users can be part of multiple federations [63], and how to develop practical and effective revocation mechanisms [22] [71].

Key finding 8: Academics expect organizational challenges related to investment cost, liability issues, identity assurance, security, knowledge and trust

Paper 4 also show organizational challenges one may run into regarding FIM. At the organisational level cost [57] may be an issue. There might be a need to make considerable investments in new technology in order to make FIM work, and complexities related to the integration and deployment may demand excessive resources . A cost-benefit analysis is needed to see if the value of the promises of FIM, such as reduced administration costs, exceeds the investments. Liability [74] [25] is another issue that needs to be considered. Federations may span across organizations with different internal rules and regulations, and across jurisdictional borders, both at national and international level. Who is responsible if/when systems and processes fail and rules are broken? Related to this, there are also assurance [20] [21] issues that need to be considered. Assurance is the process of ensuring that identity management in under appropriate control. Different organizations may have different processes and requirements for their staff enrolment and vetting procedures, and to come up with agreements among the federation partners may not be trivial. An important use case of FIM is cross-domain single sign on (SSO). The fact that one ID can be used to access several services across company borders is good for usability purposes, however, it also involves a serious drawback; The SSO use case increases the attack surface, thus increases the risk of data disclosure [56] [13] of private data. Related to all of the above is trust, as we have already mentioned [74] [70] [68] [28]. However, trust issues are complex; there might be willingness to establish trust relationships between individuals and third parties, but this might not be achievable at an organizational level, and vice versa.

Source Papers

- Paper 3: Jostein Jensen: **"Benefits of Federated Identity Management - A Survey from an Integrated Operations Viewpoint"**, Availability, Reliability and Security for Business, Enterprise and Health Information Systems, volume 6908 of

Lecture Notes in Computer Science, 1-12, 2011

- Paper 4: Jostein Jensen: **"Federated Identity Management Challenges"**, Proc. Seventh International Conference on Availability, Reliability and Security (ARES '12), 230-235, 2012

4.3.2 From Practitioners' Viewpoint

Key finding 9: Practitioners perceive that FIM will improve user administration and usability, make collaboration more efficient, reduce cost, facilitate audit and lead to better protection

Paper 5 provides insight into practitioners' perceptions of federated identity management. More **effective user administration and improved quality on the identity data** are among the perceived benefits. More than half of the interviewees mention that they believe that the adoption of FIM will rationalize the user administration process, and improve the quality of recorded identity attributes. One interviewee commented: *"you could rationalize the user account systems at different companies if you have a cooperation among them."*, and another supplemented: *It is obvious that there will be less user administration if we could integrate our login systems"*. Our interviewees also mentioned **improved usability** as a consequence of single-sign-on as a perceived benefit of FIM. Especially the contractors see the advantage of being able to experience fewer login requests in order to do their job.

Representatives from both operators and contractors talked about **improved efficiency** of collaboration as a possible benefit of FIM adoption. One of the representatives from the operator drew the parallel between standardization of equipment and standardization of identity and access management within the oil and gas industry: *"The efficiency will increase and thus our cost is reduced, if [everyone] meets the same [access control] systems when they go from one operator to another."*

Efficient user administration and efficient collaboration are factors directly linked to a desire to **reduce cost**. Both the operator side and the contractor side can gain from rationalized user administration. The operator has an expectation that FIM can reduce cost due to more efficient work processes.

Some of today's systems operate with service accounts, which are shared by several engineers. While this simplifies everyday work tasks, it makes detection and audits of potential misuse difficult. A benefit mentioned by one of the interviewees is that FIM is perceived to **facilitate audit** in the systems since every user has a personal user account, and that it will be beneficial to be able to trace who does what.

With adoption of FIM the interviewees expect to have to manage fewer user accounts and passwords, which led one of them to express that *"the perceived security will increase. Fewer passwords will be written down on paper."* The interviewees also believe that the

quality of user attributes, which is used to make authorisation decisions, will increase and that the access revocation process will be more efficient. This is highly relevant to ensure **good protection** of company resources.

This text is extracted from Paper 5, where more details and examples can be found.

Key finding 10: Practitioners expect trust issues, technological challenges, investment cost and security challenges to be obstacles to adoption of FIM in industry, and there is a risk that they confuse identity management with access management

Paper 5 also documents the challenges related to FIM, as perceived by practitioners, and the following text is extracted from this paper.

Smith [74] argues that **trust** is the fundamental concept underlying federations. At the same time he points to the fact that there are challenges related to establishment of trust. Trust issues are highly relevant for the inter-organizational collaboration, which takes place in the Norwegian oil & gas industry. One of the representatives we talked with said: *"We [...] collaborate with a license partner in one oil field. [...] At the same time we are strong competitors, so it is essential that only information concerning the collaboration is available to them."* Whether you trust other companies or not is very context-dependent, and the level of trust is difficult to define. One interviewee pointed to one of the reasons: *"I think people are slightly more sceptical of the neighbouring business considering the big money that swirl around in the oil and gas industry." "The contractors will never be able to handle the processes behind federation."* This statement by one of the representatives from the operator is illustrative to the question of whether the collaborators find each other trustworthy enough to perform all identity management within each company. At the same time he said that it would be easier to trust some of their large contractors with which they have well-established cooperative frameworks.

There are several challenges with the introduction of FIM from a **technological** perspective. The complexity of IT systems in the IO domain is high, and span regular office tools to small tailored expert systems on the software side. At the network layer they operate both with traditional IP-networks and specialized process control systems. One of the interviewed security professionals stated: *"It would be a dream come true if everyone could connect to a common platform - an information bus - where all information could be shared securely [...] but it is hard to believe that it will be possible."* Further, he explained that identity federations might be possible in the future for some of their large partners and for some of their large systems. It can, however, be more difficult for smaller systems originating from small companies, who might not have the competency or economic baseline to integrate their systems with other federated systems according to standardization and interoperability needs.

Representatives from the operator told us they have tried open source federation technologies in a few cases. However, they experienced some technological challenges, especially related to the technology management. *"It is much easier to rely on technology from Mi-*

crosoft, for instance, rather than a product from a party that is not as big commercially". With this he implied that the large software companies would have to come up with solutions that fulfill the industry's needs before they will consider the technology in a larger scale. A second interviewee stated: *"The challenge with federated identities, as I understand it, is that there is no dominating standard. [...] You need a bouquet of different technologies."*

A software developer in our study argued that that despite the rapid development of technology that can have high benefit internally in a company there are still considerable challenges as soon as you get outside company borders where you meet equipment from different vendors, different security policies, firewall setting and so on. Even if software and hardware interfaces were compatible, there are still challenges with the interpretation of data originating from different systems, especially regarding semantics. These considerations are very valid when looking at integration of identity management systems. Both software and hardware interfaces must be standardized. Protocol options must be defined so that all the equipment is interoperable and the semantic meaning of identity attributes must be defined and agreed upon. Several of the interviewees mention that the industry must agree on common guidelines for FIM at a detailed level for it to be successful.

During the interviews we asked the candidates if they saw any potential showstoppers for adoption of a common FIM platform in the Norwegian IO context. *"Who's gonna pay for the fun?"* was the immediate response of a consultant in our study. He then elaborated: *"It is obvious that all the participants in such collaboration will have to make major changes to get this up and running. That is a cost I'm not sure they are willing to take."* Representatives from three of the four contractors in our study confirmed this view. *"We have to consider that we are delivering services [to oil companies] globally. It will be costly for us to implement a system for collaboration only with our Norwegian partners*, one of the interviewees said. The two other representatives were concerned about the funding for implementation of federation technology. *We don't develop anything that is not paid for by someone."* Even though the representatives from the operator did not mention funding as a factor, they all recognized that the **investment cost** of a FIM solution will be considerable.

Security issues are also a major concern in the industry. *"Our biggest fear is that someone unauthorized can get access to, and control a production process."* More than half of the interviewees were concerned that FIM will increase the attack surface of their systems. *"The drawback is that someone could authenticate as another user. She would then automatically get access [...] to all the companies where this user has access rights."* Identity theft is obviously a serious concern, but the interviewees are not only worried about hackers with sinister intentions: *"The risk of unintended errors increases. Someone can cause situations by mistake since they don't understand the consequences beyond their own company."* Some are also concerned that there will be fewer explicit barriers between systems of different criticality. They feel that they lose control when the systems are being accessed transparently. Privacy aspects related to FIM is currently a hot research topic. However, only one of the interviewees mentioned privacy as a concern. *"It would*

be fantastic to just have one digital identity to relate to, which you could use for everything. The drawback, however, is that you can trace what people are doing. [...] It might not be that important in this context, but often it is ok to know that you act anonymously so that you don't have to account for everything you do."

A major challenge related to identity management in general is to keep user databases and users' access rights up to date. One of the foreseen benefits of FIM is that it will be easier to keep the identity data up-to-date. However, more than half of the interviewees still think there will be challenges, despite identity federations. *"I'm not so sure if we will experience less administration with such a system. I guess we [...] will get fewer users to administer, but I'm not sure about this simplification."* He continues to argue that there will still need to be processes to trust each specific user before they can be authorized to access systems, and that it will be necessary to implement processes to verify the quality of federation partners' identity management processes. The interviewees emphasized the need to fully control the authorization part of the access control. The authentication service can be outsourced to a trusted third party or users can be authenticated within their home organization, but there is still a need to have strict control on which, or what type of users have access to the organisation's resources. However, here there is a risk of **confusing identity management with authorization**, since this authorization process is already being performed today, only with the added chore of local identity management in each case.

Key finding 11: Practitioners question whether there is sufficient organizational maturity to adopt new identity management solutions

We document some of the challenges the oil & gas industry is faced with today both in Paper 5 and Paper 6. The following text is extracted form Paper 6.

The oil and gas industry in Norway is collaborative and competitive at the same time. That is, the same companies that are competing to win a project bidding are often at the same time collaborating on another project. The government sometimes even require oil companies to use shared production facilities in order to make smaller oil-fields profitable. Thus, at times oil companies must utilize the infrastructure of their competitors in order to operate their own oil fields. As noted by one of the interviewees *"We see more and more of this. Many of the newly discovered oil reservoirs are small, and does not justify a full size production facility. Subsea installations, however, can be 'put to operation' and then utilize nearby production facilities."* For suppliers and contractors it is also the norm that they collaborate with their competitors. Two companies competing for a contract may at the same time depend on each other to deliver on another project. Further, there is also a constant battle for suppliers and contractors to increase their share of the contracts, such that they are also in a kind of competition. Further, the complexity and novelty of oil and gas installations require extensive collaboration and sharing of resources. Thus, the oil and gas industry needs to share access to their assets while at the same time protecting them.

The extensive collaboration between companies in the oil and gas industry makes it necessary for external users to access company-internal resources. Currently, there is no collaboration on the management of identities, and therefore companies issue user identities to individual users of external companies whenever this is required. One of our interviewees explained a complex situation where the company has approximately 20.000 employees, but administer more than 70.000 user accounts in their IT-systems. Coupling this with the fact that the oil industry at large experiences an estimated 40.000 internal and external job changes per year, it is evident that managing user identities is far from trivial.

While office systems in general utilize a centralized user database and authentication mechanism, many process control systems are not designed to be integrated with existing access control solutions. As a result, several of these expert systems have their own user database, with their own user identities and authentication mechanisms. Thus, changes to a single user's identity attributes (e.g. role in the organization) must be propagated to many systems, both internally and externally.

Offshore installations run continuous operations and require constant monitoring. As one of the interviewees stated *"one operator can't log off, and a new log in, since critical situations can appear in that timeframe"*. Consequently, many of the operational systems utilize shared user accounts for all operators, and sometimes these usernames and passwords are recorded in a book lying next to the operator terminal. Our interviewees recognize that this is a risk, but the alternative of not being able to access a system is considered worse. Shared accounts add to the challenges of maintaining identities and access rights, since it is very difficult to determine who has access to these shared accounts (i.e. who knows the password?).

Office system users are commonly granted a default set of access rights initially aimed at internal users. As a consequence, external users get access to internal services such as e-mail (with an e-mail address), access to the company management solution, access to employees' calendars and more, regardless of whether they need it or not. This is excessive for a person that *"needs to monitor our need for methylated spirit or need for cabbage"*, an interviewee explained.

Despite the fact that user accounts are considered personal, another way of handling problems with improper access rights is to share user identities and authenticators between personnel. Both representatives from the operator and contractors admit that sharing of user identities takes place. The person borrowing a co-worker's identity often has valid access, but has forgotten his security token or it has expired. The process of getting a new one is cumbersome, bureaucratic and in most cases time-consuming. Thus, while waiting for a user ID an employee might be needed to do a job and therefore borrows an ID from a colleague. Operators are aware of the situation, but explained that *"when a situation causes half a million of lost revenue per minute, this is bound to happen."*

Oil and gas installations typically remain active for decades and therefore some of the corresponding tailor-made process control systems are old. As in many industries, companies are reluctant to update or replace systems since this may negatively impact the

production. As one of the information security specialists stated in an interview: *"We have state of the art systems, but it is state of the art from 1985 in some cases"*. As the threat environment has substantially changed, security technology has to a large extent been built like fortresses around the computer systems and relies heavily on traditional network defense mechanisms. Firewalls and IP-based access control lists are used to control network activity to and from resources. *"Security is unfortunately not as integrated in our systems as we would like, but that is why firewall filtering and anti-virus scans are so important to us."*

On the question on whether there had been discussions concerning integration of user databases, on of the interviewees replied: *"That question has never been raised. Most oil companies have clear rules preventing it". "What holds [FIM] back is the same challenge we experienced when we first introduced the Integrated Operations concept. People are satisfied with the way they work today, and do not want change."* This indicates that practitioners question whether there is organizational maturity to adopt FIM solutions in the industry. At the same time the current challenges should motivate for the introduction new security concepts. *"Ten years ago, when some of our customers started [with IO], it was nearly impossible to get inside their premises with a computer. Now we get access to networks, get IP addresses, and so on."* This comment was made by one of the interviewed contractors and it show that despite a slow moving industry, there is a constant change in attitude when it comes to taking advantage of new communication technology to facilitate sharing of data.

Source Papers

- Paper 5: Jostein Jensen, Martin Gilje Jaatun: **"Federated Identity Management - We Built It; Why Won't They Come?"**, IEEE Security & Privacy, 34-41, 2013
- Paper 6: Jostein Jensen, Åsmund Ahlmann Nyre: **"Federated Identity Management and Usage Control - Obstacles to Industry Adoption"**, Eighth International Conference on Availability, Reliability and Security (ARES '13), 2013

Chapter 5

Discussion and Implications of Results

5.1 Evaluation of Contributions

In the following we will look at our key findings in regard to how they have implication for:

- IT strategy planners and solution architects
- Researchers
- IdM software developers and product manufacturers

We end the chapter with a discussion of the main limitations of our work, and suggest some directions for further work.

5.1.1 Implications for IT Strategy Planners and Solution Architects in the Oil & Gas Industry

IT strategy planners and solutions architects should understand the strengths, weaknesses, opportunities and threats (SWOT) of new technology that is considered for adoption in a given context. Our key findings can be used as input to a SWOT-analysis.

Strengths: From key finding 5 and 9 we have seen that both academics and practitioners anticipate FIM to improve usability related to access control solutions. This has to do with the single-sign-on feature, which allows end users a more seamless workflow without having to type in a username and password for each service or information request. Academics and practitioners also agree that FIM can lead to better protection. The academics expect fewer passwords to be written down on paper, and that users should be able to select stronger passwords, while the practitioners confirm the situation that the number of usernames and passwords that they must remember is so high that they currently write them down in books.

Through key finding 6 and key finding 9 we have also seen similar expectations between the academic community and industry representatives. Academics have argued that FIM will reduce administrative cost related to identity management. The practitioners share this view, and the example (from key finding 11) where one company has 20.000 employees, but administers more than 70.000 user accounts is illustrative in this regard. With such a situation one can also understand the academics and the practitioners when they expect improved quality in the recorded identity attributes. It is easier for the employer to catch changes related to the identity attributes of an employee than it is for an external company.

Academic literature explains that service providers can experience reduced complexity since identity management can be outsourced to a third party (key finding 6). Service providers can then focus on the development of functional aspects of the service. The practitioners did not talk about outsourcing of identity management as mechanism to reduce complexity, however, they mentioned outsourcing of IdM as a mechanism to establish trust related to the identity federation (key finding 10). Due to the collaborative, but at the same time competitive nature of the oil and gas industry (key finding 11) it would be easier to trust identity management to an external third party.

In key finding 5 academics say that improved end user privacy is one of the expected benefits related to FIM. One of the practitioners mentioned that FIM will facilitate system audits as every user is likely to access systems by means of personal user accounts, which is in contrast to todays' situation where common service accounts are used for many of the production systems (key finding 9). In Paper 5 we claim that although some users might be concerned about privacy aspects of FIM, we believe that the ability to perform audits will trump privacy on enterprise systems.

Weaknesses: In key finding 4 we claim that there is little empirical evidence related to adoption of FIM in industrial contexts. Consequently, decisions to adopt FIM will be made based on expert opinions, or the voice of FIM manufacturers' market divisions, rather than data from rigorous and unbiased research.

It is important to be aware of interoperability issues related to new technology, and especially technology that is meant to facilitate collaboration between organizations. Through key finding 7 and key finding 10 we found that both academics and practitioners foresee challenges related to interoperability. Different companies may adhere to different standards and tools for federated identity management in inter-organizational collaborations, and FIM vendors applying the same standard may utilize different protocol options that may prevent their solutions from co-existence within the same federation. The practitioners also add that existing services must be integrated with the FIM solutions, but that smaller product vendors might not have the competency of economic baseline to make these integrations. Further, both academics and practitioners emphasize the need to agree on the semantic meaning of identity attributes within the identity federation.

Academics and practitioners are also aware that the adoption of FIM is associated with high expenses related to investment in new technology and integration efforts (key finding

8 and 10). Will the investments be a result of a Dutch treat within the industry? Will international companies be willing to make major investments in FIM technology due to possible demands from their Norwegian partners? Within the industry it is clear that there must be agreements on who is responsible for the cost.

In key finding 10 our practitioners question whether the introduction of FIM will result in less administrative overhead, and argue that they need strict control on who has access to what. The burden of access control will not be less with FIM, they claim. Although this is most likely true, it also show that some of the practitioners confuse identity management with authorization, and that the concept of separating identity management from authorization is new in this domain. The knowledge about the technology may not be sufficient among all stakeholders to take informed decisions. This is also related to key finding 11, where we show that the collaborative picture is complex, and where the practitioners themselves question whether the industry is mature enough to adopt new identity management solutions to facilitate identity federations.

Opportunities: Academics expect FIM to facilitate cooperation among service providers, and that seamless service access across company boundaries may be a result (key finding 6). Cooperation is also important for the practitioners. Employees in the oil and gas industry will be able to work more efficiently if they meet the same access control solutions in each company (key finding 9). In Paper 5 we also explained that an onshore worker may be responsible for controlling several offshore facilities, and that this person needs to visit 15 to 20 different systems per day. Others might need to visit up to five different systems with separate logins to create production reports. FIM may lead to improved usability for these people, but also make their working day more efficient, and thus cost efficient. Key finding 11 shows that co-workers borrow each other's access credentials from time to time. The practitioners have explained that one of the reasons for this can be that they have forgotten their infrequently used password used to access external collaborator's services. It might be time consuming to wait for the bureaucratic process of resetting the password. With FIM, workers will use of their everyday credential to access all services. The chance of forgetting credentials will be limited, the inefficient waiting time for credential reset will be reduced, and the chance that credentials are borrowed from a co-worker, against company security policies, may be eliminated.

Key finding 3 shows the importance of ensuring a high quality in all parts of the identity management process. The whole identity management process may be at risk if one of the activities is flawed or insufficient. This again may leave valuable company assets at risk. FIM adds to this challenge in that weaknesses in one company's IdM lifecycle affects the security level of collaborators. This, however, can also be seen as an opportunity. FIM will require common policies for identity management to be developed within the federation. As such, the industry at large can use the introduction of FIM to clean up identity management processes and agree upon shared IdM requirements. *"The contractors will never be able to handle the processes behind federation."* (key finding 10) was one of the replies from an operator representative. Common policies for FIM may contribute to improved trust since all partners involved in the federation must adhere to them.

Threats: The concept of federated identity management involves sharing of identity data across company boundaries. Consequently, identity data will also be transferred outside the traditional defence mechanisms of a company, including the company firewall. The number of stakeholders with the means and motive to intercept and analyze the traffic, including identity data, will increase significantly. Key finding 1 and 2 clearly show that poorly implemented FIM solutions will affect the security level negatively. Identity assertions with insufficient protection, for instance, may allow adversaries to circumvent authentication procedures through, e.g., replay attacks) or elevate privileges, e.g., by modifying identity attributes.

Liability issues are pointed to by academics as one of the challenges regarding FIM (key finding 8), and they claim that federations may span across organizations with different internal rules and regulations and across jurisdictional borders, both national and international. Most of the companies within the oil and gas industry, operating on the Norwegian continental shelf, are multinational and have global operations. One of our practitioners stated: *"Today, we have a responsibility to protect our customers' data, and it is an enormous responsibility for us to ensure that it is not being misused."* In paper 5 we question whether collaboration partners are wiling to assume even more risk. Are they willing to accept liability for downtime on production facilities caused by employees' inability to access and monitor, for example, safety-critical processes, due to trouble with their internal identity management solutions? Further, what happens if unintended error occurs due to extended access based on FIM, and what happens if a contractor identity is stolen and used to access resources at a collaboration partner? Questions regarding liability must be properly discussed.

Both academics and practitioners are concerned about security issues related to use of federated identity management (key finding 8 and 10). Identity theft and impersonation attacks based on stolen identities are highlighted by the academics, and the use of identity assertions, e.g. to achieve cross-domain single-sign-on contributes to these threats. Message security and proper use of integrity and confidentiality protection is essential. The reality of the academics' concerns is illustrated through Paper 1 (key finding 1 and 2). The practitioners also mention security issues as a major concern (key finding 10). They feel that the use of FIM blurs the barriers between existing security domains, and their biggest fear is that someone unauthorized can get access to and control a production process. They point to the risk of identity theft, and also that the risk of unintended errors, which might cause consequences outside their company's premises. The introduction of poorly implemented identity management solutions will put a company's (or all companies within a federation's) assets at risk. The security specifications of the IdM product being evaluated for purchase should be studied, and one should request information about how the product was developed. Only IdM manufacturers that can prove their security awareness through the whole development lifecycle should be considered.

Last, trust issues must be given attention before one selects to adopt the technology (key finding 8 and 10). The practitioners explain a situation where companies can collaborate on one project, but compete in others. It is essential that only information relevant for a

given collaboration is available. Companies must trust the technology, and they must trust that collaborators in the identity federation maintain a high quality on their identity management processes. An expression such as *"The contractors will never be able to handle the processes behind federation"* indicates that introduction of FIM is not an overnight work task.

5.1.2 Implications for the Research Community

First, this work has provided insight into the Norwegian oil and gas industry, and their perception of FIM technology, through empirical research. The complexity of the industry and current security challenges related to identity management and access control has also been identified through Paper 5 and 6. In light of Hevner and March's design science framework [35], this contributes to building essential knowledge about an environment where FIM can be utilized in the future, and thus can be used to improve the relevance of research initiatives towards this domain.

Second, this work has identified that there is little focus on empirical research within the federated identity management community. Hevner and March [35] state that significant progress within a research field[1] will only occur when there is a balance between traditional design science and empirical research. A combination between the two is essential to improve relevance and rigor of the research. The research community should continue to develop novel security technologies, but at the same time spend more effort on empirical evaluations of their innovations.

Third, the last decade's research on identity management and especially federated identity management has clearly shown the benefits of separating identity management from authorization management, and reduce the coupling between these and the applications to be access protected. As we indicate in key finding 10, however, the industry still seems to be influenced by the "old" tradition of bundling applications with a combined authentication and authorization mechanism. The research community should join efforts to develop communication models to explain the benefits of separating identity management from authorization management.

5.1.3 Implications for IdM Developers and Product Manufacturers

Through Key finding 1 and Key finding 2 we have highlighted some of the difficulties of developing secure identity management solutions. A secure development strategy should be in place for the realization of the technical platform. Security should be an inherent part of all software development phases, from requirements specification, though design, implementation and testing. Paper 1 gave examples of what can happen if the implementation fails.

[1] their focus is on information systems

Key finding 7 is also important for developers and manufacturers of IdM solutions, especially those who aim for development of solutions for federated identity management. Most inter-organizational collaboration arenas are heterogeneous in nature. There will be variations in operating systems, office tools, collaboration software, in-house systems, out-sourced systems, cloud services and so on and so forth - all of which must be integrated with the various (federated) identity management solutions that are being used by each collaborating partner. IdM manufacturers must take interoperability issues into account. They must be willing to adhere to open standards and be open on the protocol options they use. They must not fall into temptation of implementing some conformance variations that lead to vendor lock-in and difficulties for their customers to realize new identity federations and IT-based business collaborations with new partners.

Key finding 10 documents some of the challenges the oil & gas industry sees in relation to adoption of FIM, and the concerns they have related to this concept. The practitioners confirm a complex IT environment, and the need for standardized and interoperable solutions. However, maybe the most important point for developers and manufacturers to be aware of is the industry's fear that FIM will negatively affect the security level in that they loose control of who can access what. The product manufacturers' sales departments must be explicit about the differences between identity management and access management, and that allowing an external entity control of the identity management process does not reduce their ability to control access with the same granularity as today.

Finally, Key finding 10 includes the hint of a new business strategy for IdM product manufacturers. Develop an identity management service that can act as a trusted identity provider within the oil & gas industry. As one of our interviewees stated: *"I have some trouble imagining that access to external resources can be given if [the identity management process] is to be handled within each company. I feel that it has to be organized by a common entity."*

5.2 Adoption of FIM in the Oil & Gas Industry

In Paper 5 we ask the question: Are identity federations attractive for the industry? To support the answer of this question we turn to Roger's theory on Diffusion of Innovations [65], and especially his focus on "the perceived attributes of innovations". These attributes are: relative advantage, compatibility, complexity, trialability and observability.

The relative advantage of an innovation is one of the strongest predictors of its adoption rate [65]. Our study shows that there are perceived advantages of FIM adoption for all collaborators in the Norwegian oil and gas industry (key finding 5 and 6), but that some of these are offset by either compatibility issues or increased complexity (key finding 7 and 8). At the same time, there are examples in which FIM is being tested in the industry, and from our discussions on organizational maturity, we see that there is a willingness to proceed.

FIM security technology interferes with, but isn't a part of, a primary business process. It's a preventive innovation, which according to Rogers, *"has a particularly slow rate of adoption because individuals have difficulties in perceiving its relative advantage"* [65]. So, maybe it's not strange that the adoption process is slow in the industry. Our findings are in line with Smith's claim that FIM adoption will be an evolution, rather than an overnight revolution [74].

Even though we can not give a short answer to the slow adoption rates of FIM there are three aspects that are appropriate to highlight: trust, understanding of the FIM concept and development support for FIM. *Trust* is a complex topic and a fundamental concept underlying FIM [74]. Paper 5 (and key finding 10), document evidence that the needed trust level among collaborators may not be sufficient to establish identity federations within the Norwegian oil and gas industry at current time. As an example, our informants in the case study questioned whether their collaborators could handle the processes behind FIM, and whether FIM technology could sufficiently be used as part of the asset protection strategy in the collaborative, but competitive industrial environment. We have also documented that there may be confusion within the industry regarding the understanding of the FIM concept and the possibility to separate identity management, including authentication mechanisms from authorization management. Challenges related to *understanding of the FIM concept* may influence negatively on the trust level associated with the technology, and the industry's willingness to rely on it. The industry participants fear that they will loose control over 'who has access to what' within the FIM collaboration, and such skepticism must be reduced to facilitate FIM adoption. Last, we documented the skepticism the industry has in the maturity of the FIM technology in that some of our interviewees are concerned about the complexity involved in developing new FIM solutions with existing technology. Until the last few years this has been a legitimate concern. The starting point for building a FIM solution would be to develop it from scratch with standards specifications, such as SAML2.0 as input. The few last years, however, we have seen an increasing number of maturing solutions to build FIM systems provided by large software companies, such as Microsoft. Active Directory Federation Services is an identity federation solution that can be integrated with existing user directories within an organization, and .NET frameworks[2] *provide development support* through application programmer interfaces (API's) specifically created for the purpose of developing FIM solutions. Mature identity assurance frameworks are also essential in supporting development of FIM solutions at a conceptual level. Common policies and governance mechanisms must be in place within a FIM collaboration to improve the collaborators trust in each others identity management processes.

Our interviews paint a picture of a complex industrial IT landscape, currently lacking the maturity level necessary to implement a global, ubiquitous FIM solution. There was also skepticism among interviewees as to whether systems of different criticality should be connected at all, now or in the future. The vision of a ubiquitous FIM solution that integrates with all systems might be too ambitious and certainly conflicts with Ross Ander-

[2]e.g. Windows Identity Foundation http://msdn.microsoft.com/en-us/security/aa570351.aspx

son's observation that *there are always systems that don't fit"*[16]. However, we believe that the broader industrial audience will adopt some form of federated identity management sooner or later. The fact that they've started experiments with the technology is a good indication, and the perceived benefits are clear. The challenges are complex, but being aware of them will stimulate discussions among collaborators so that palatable solutions can be found.

5.3 Limitations

This research is not without limitations. In the following we emphasize the most important limitations, which have an influence on the interpretation of the results presented in this dissertation.

In Study 1 we present experiences from the security testing of a prototype application that was developed as part of a research project, and whose purpose was not primarily aimed at solving security challenges. Our findings can as such not be generalized to development projects with a purpose to deliver solutions for the commercial market. However, the experiences and results we present indicate potential consequences of developing identity management solutions containing security vulnerabilities and flaws. History has shown that people[3] are creative, and that they will try to find software vulnerabilities to achieve their goals, not matter how banal or complex the vulnerabilities are. Identity theft, replay attacks using identity assertions, and elevation of privileges are examples of sub-goals attackers would like to achieve in order to reach their final target, which for instance may be a company's digital assets. Our results exemplifies unwanted situations, and must not be used as evidence that existing identity management solutions are troubled with security vulnerabilities and flaws.

Generalization, or external validity, is also an issue for Review 2 where we studied the academics' view of the benefits and challenges of federated identity management. The purpose of this thesis is to gain understanding regarding the concept of FIM from an industrial viewpoint. In our review work, however, we did not evaluate the context, which academics made their statements based on. Most studies involved in the review describe aspects of FIM within an Internet-based scenario where service providers on the open Internet may benefit from FIM technology. Few, however discuss FIM in an industrial context where the goal of the FIM solution is to facilitate secure sharing of sensitive company owned assets. Consequently, one can not guarantee that the results from the review are applicable within the industrial case. In section 5.1.1, we use the results from Review 2 (Paper 3 and 4) and apply them in a SWOT analysis intended for use by strategy planners or solutions architects within the industry. They, in particular, need to be aware of this limitation. However, we also include the practitioners' perceptions in the SWOT analysis, and a comparison show that both academics and practitioners share perception of the FIM technology in many areas.

[3] i.e. hackers, attackers, industrial spies or any other actor with bad intentions

In Study 2 we carried out a case study based on semi-structured interviews. The goal of the study was to collect practitioners *perception* of the FIM technology. Since the FIM technology has not been widely adopted in the industry yet, their statements will be colored by previous experiences with similar technology used in other domains, such as experiences with use of the government provided FIM solution presented in section 2.6, or FIM solutions offered by a social media providers. Further, there is a risk of bias in the data collection and the analysis phase. This is related to the study's construct validity. Yin [80] states that multiple sources of evidence should be used in case studies to strengthen the evidence, otherwise there is a risk of subjective judgements. Our study use interviews as the sole source of information. This, however, was the only solution to collect data about a technology that has not yet been utilized. Further, one researcher alone was responsible for the data analysis. This is also a weakness that may color the results. The co-author of paper 5, however, participated in the development of the research strategy, including definition of research questions and development of the interview guide. Further, he commented on and criticized the results prior to publication is an effort to avoid potential observer bias.

We do not make claims that our results from Study 2 can be generalized outside the context of the Norwegian Oil & Gas Industry. Our effort is focused on providing a richer insight into this domain. This industry may share characteristics with other industries, but new research efforts should be made to see if our results are applicable to a wider audience.

Our original plan for this research work was to carry out multiple iterations of the design science approach. Lack of existing knowledge related to adoption of FIM in industry, together with lack of insight in the complex industrial environment, however, made us deviate from this plan. Still, it is a limitation of this study that it measures perception of the FIM technology, and not real world experiences with its adoption and use.

5.4 Recommendations for Future Work

Through this work we have collected information that can be used as input to the development of a federated identity management specification[4] for the Norwegian oil and gas industry. The next step in a design science approach would be to evaluate this specification, for instance through focus groups to get the practitioners' feedback before a prototype FIM solution is built. This FIM solution should be integrated with a service of low criticality, but which is frequently used within the production environment by several of the collaboration partners within the industry. This should ease on our interviewees' skepticism that FIM will blur the security boundaries between systems of different criticality. A case study or ethnographic study could be designed in this regard; first to document implementation challenges in an operational environment, and then to obtain empirical data on the practitioners' perception of real use of FIM technology. The approach would also

[4] which should satisfy Hevner and March's requirements for relevance in the design science approach

facilitate collection of data from several data sources, such as interviews, usage logs, help desk request regarding user administration, and written agreement between the federation partners. Such a study would provide invaluable empirical results back to the research community. This strategy is in line with our recommendations on how to facilitate adoption of FIM in Paper 6, and would give real world data on the relative advantage of FIM, its compatibility with existing practices within the oil and gas industry, and allow the industry to test FIM solutions in small scale (trialability) before deploying FIM in a wider context[5].

Another natural progress in the research on FIM would be to design case studies to collect empirical data related to identity federations that already have been implemented. In section 2.6 we gave a few examples of existing solutions within the Norwegian government and educational sector. Research should be carried out to learn what worked, and what did not. Retrospects of implementation projects should be carried out to generate good practice guides on how to organize FIM implementation projects.

Siponen and Oinas-Kukkonen [73] carried out a comprehensive review of existing academic literature on the information security topics: access to information systems, secure communication, security management and development of secure information systems. They split the research into technical, conceptual and organizational contributions. Their findings suggest a bias towards technical research. Further, they found that most academics within this field lean towards mathematical approaches (including logic) and conceptual analysis to evaluate their research. Within the category organizational research they found the prevailing research approaches to be conceptual analysis and empirical research. However, they claim, the amount of empirical research is sparse. Labunets [53] recognize the lack of empirical research as a challenge within the information security research community, and argue that the lack of empirical validation is a drawback for both practitioners and researcher. Practitioners will not be able to take informed decisions as to which security technique/method/tool to use and which is the better for their context, and researchers will not gain sufficient knowledge to understand how their research results can be improved to meet industry expectations. The scarcity of empirical research on information security, and more specifically model driven development and federated identity management is also confirmed by the two reviews being part of this thesis (documented in Paper 2 [45], Paper 3 [43] and Paper 4 [42]). From this we can deduce that information security is a research field where the use of empirical research methods is still in its infancy. In Paper 6 we argue that there is a need for further research on the adoption of security technology in general, and especially that empirical data should be collected, both from organizations that have adopted a technology and those that have not. Our claim is that a thorough understanding of the factors influencing adoption can help us target development and deployment efforts towards the specific factors that are most challenging.

There is, however, an increasing awareness about the need for empirical research within

[5] relative advantage, compatibility and trialability are related to Roger's theory on Diffusion of Innovations [65]

the academic information security community. This is exemplified by two panel discussions from the last few years' international information security conferences. Niekerk and colleagues [59] held a panel on the topic at the 25th IFIP TC 11 International Information Security Conference (IFIP SEC) 2010 and the 7th International Conference on Availability, Reliability and Security (ARES) 2012[6] included a panel on the subject, moderated by Shari Lawrence Pfleeger.

We acknowledge that it is difficult to apply research methods and instruments, including those designed for empirical research, that are not specifically designed to do research on information security topics. Nyre and Jaatun's [60] research on adoption of security technology (usage control) in industry is a good example that existing research instruments on technology adoption can not automatically be applied to the information security domain. Instead, a number of existing theories must be combined and tailored to fit adoption of non-functional technologies such as security technologies. The more empirical research the information security community carries out, the more familiar we become with the methods, and the better we as community can design our (hopefully reusable) research instruments. This is not to say that empirical approaches are better than others, as also emphasized by Hevner and March [35]; *"significant progress in information systems research will only occur when the discipline as a whole recognizes and appreciates the benefits of both paradigms"*, they claim. This PhD thesis contributes with empirical research to the information security community in an attempt to even the balance between technical and empirical research contributions to federated identity management.

[6] http://www.ares-conference.eu/ares2012/

Chapter 6

Conclusion

This thesis has addressed challenges in developing secure identity management solutions for distributed service platforms, presented empirical evidence on development and adoption of federated identity management (FIM) in industry, and discussed benefits and challenges related to FIM.

First, through our findings we have seen the importance of being aware of security issues in all phases of the software development lifecycle. Software in general is burdened with security bugs and flaws. Unfortunately, this is also true for security software such as solutions for federated identity management. Our research shows examples on how identity management solutions with vulnerabilities, and especially poorly protected identity assertions in distributed (or federated) environments can be exploited in several different ways. Further, we have documented that secure identity management is more than secure technology. During technology development one must consider all phases of the identity management lifecycle to ensure that the technology provides secure interfaces and secure mechanisms to support the identity management activities. However, it is also important to focus on the organizational identity management processes to make sure these are of high quality, and mitigate threats targeted at human interactions. Secure identity management technology does not eliminate challenges related to human errors and flawed processes.

Second, we have seen that there is little empirical evidence on the development and adoption of FIM in industry. The decision to adopt FIM in inter-organizational industrial environments, such as the Norwegian oil and gas industry must therefore be based on perception of the technology, rather than documented evidence. There exit examples of successful FIM deployments, e.g. in government sectors and in educational sectors, as well as examples from the social media services on the Internet. Unbiased research on these examples would be appreciated, in order to build new solutions based on existing knowledge. Even though this thesis provides insight into industrial expectations and needs based on empirical research, it is limited to practitioners' perception of a technology that has not yet been adopted in their working environment. More research is needed to document the impact of large-scale deployments of FIM in inter-organizational contexts.

Finally, we have documented the benefits and challenges related to FIM from the perspective of the academic community and the Norwegian oil and gas industry. In Paper 5 [44] we state that our interviews with industry practitioners reveal that some of FIM's benefits are offset by its challenges, we also claim that a vision of a ubiquitous FIM solution covering all services within the Norwegian oil and gas industry might be too ambitious. We agree with the skepticism of our interviewees who fear a too tight connection between systems of different criticality. Still, we believe that the broader industrial audience will adopt some form of federated identity management sooner or later. The challenges are complex, but being aware of them will stimulate discussions among collaborators so that palatable solutions can be found.

In Paper 6 [46] we took a deeper look at factors affecting adoption of security technology, including FIM. In this paper we state that the adoption rate of FIM has been slower than expected by the academic community. Further, we conclude that there are a number of factors affecting technology adoption involving the context in which the technology is placed. We have not been able to identify a short answer to why the adoption rate remains low, but we believe that there are many ways in which it can be improved. We also advocate demonstrating how the technologies constitute an evolution of the access control system, rather than a revolution. We believe that a continuation of the original design science approach selected for this PhD project would be a good approach for this purpose. Through this thesis we have provided a knowledge base that should be used in the development of a FIM architecture tailored for the oil and gas industry, and to increase the relevance of future FIM prototypes built through research initiatives to support a soft and smooth transition from current identity management solutions to future solutions supporting identity federations.

Identity management is a fundamental prerequisite to ensure information security. Our research has shown that the Norwegian oil and gas industry experiences challenges with current identity management practices. The time is ripe reconsideration of existing practices and introduction of new solutions.

Bibliography

[1] Digital Agenda for Norway, ICT for Growth and Value Creation. Technical Report Meld. St. 23 (2012 - 2013), Norwegian Ministry of Government Administration, Reform and Church Affairs.

[2] eDrift på norsk sokkel - det tredje effektiviseringsspranget. Technical report, the Norwegian Oil Industry Association (OLF), 2003.

[3] Internet Security Glossary, Version 2. Technical Report RFC4949, 2007.

[4] WS-Trust 1.3. Technical Report ws-trust-1.3-spec-os, OASIS, 2007.

[5] Framework for Autehntication and Non-Repudiation in electronic Communication with and within the Public Sector (Norwegian title: Rammeverk for autentisering og uavviselighet i elektronisk kommunikasjon med og i offentlig sektor. Technical report, Det kongelige fornyings og administrasjonsdepartementet, Norwegian Government, 2008.

[6] Security Assertion Markup Language (SAML) V2.0 Technical Overview. Technical Report sstc-saml-tech-overview-2.0-cd-02, OASIS, 2008.

[7] National e-Authentication Framework. Technical report, Australian Government, Department of Finance and Deregulation, 2009.

[8] NGN Identity Management Framework. Technical Report IITU-T Y.2720, ITU-T, Telecommunication Standardization Sector of International Telecommunication Union, 2009.

[9] Web Services Federation Language (WS-Federation) Version 1.2. Technical Report ws-federation-1.2-spec-os, OASIS, 2009.

[10] E-Authentication Guidance for Federal Agencies. Technical Report OMB Memorandum M-04-04, 2011.

[11] Information Technology - Security techniques - A framework for identity management - Part 1: Terminology and concepts. Technical Report ISO/IEC 24760-1, ISO/IEC, 2011.

[12] Potential value of Integrated Operations on the Norwegian Shelf. Technical report, the Norwegian Oil Industry Association (OLF), April 2006.

[13] Gail-Joon Ahn and John Lam. Managing Privacy Preferences for Federated Identity Management. In *Proceedings of the 2005 workshop on Digital identity management*, DIM '05, pages 28–36. ACM, 2005.

[14] Gail-Joon Ahn, Dongwan Shin, and Seng-Phil Hong. Information Assurance in Federated Identity Management: Experimentations and Issues. In *Web Information Systems (WISE 2004)*, volume 3306 of *Lecture Notes in Computer Science*, pages 78–89. Springer, 2004.

[15] Florina Almenárez, Patricia Arias, Andrés Marín, and Daniel Díaz. Towards Dynamic Trust Establishment for Identity Federation. In *Proceedings of the 2009 Euro American Conference on Telematics and Information Systems: New Opportunities to increase Digital Citizenship*, EATIS '09, pages 25:1–25:4. ACM, 2009.

[16] Ross Anderson. Can We Fix the Security Economics of Federated Authentication? In *Security Protocols XIX*, volume 7114 of *Lecture Notes in Computer Science*, pages 25–32. Springer, 2011.

[17] Patricia Arias Cabarcos, Florina Almenarez Mendoza, Andres Marin-Lopez, and Daniel Diaz-Sanchez. Enabling SAML for Dynamic Identity Federation Management. In *Wireless and Mobile Networking*, volume 308 of *IFIP Advances in Information and Communication Technology*, pages 173–184. Springer, 2009.

[18] Sriram Balasubramaniam, Grace A. Lewis, Ed Morris, Soumya Simanta, and Dennis B. Smith. Identity Management and its Impact on Federation in a System-of-Systems Context. In *3rd Annual IEEE Systems Conference*, pages 179–182, 2009.

[19] Adrian Baldwin, Marco Casassa Mont, Yolanta Beres, and Simon Shiu. On identity assurance in the presence of federated identity management systems. In *Proceedings of the ACM workshop on Digital identity management*, pages 27–35. ACM, 2007.

[20] Adrian Baldwin, Marco Casassa Mont, Yolanta Beres, and Simon Shiu. Assurance for Federated Identity Management. *Journal on Computer Security*, 18(4):541–572, 2010.

[21] Elisa Bertino, Lorenzo Martino, Federica Paci, Anna Squicciarini, Lorenzo D. Martino, and Anna C. Squicciarini. Digital Identity Management and Trust Negotiation. In *Security for Web Services and Service-Oriented Architectures*, pages 79–114. Springer, 2010.

[22] Elisa Bertino, Lorenzo Martino, Federica Paci, Anna Squicciarini, Lorenzo D. Martino, and Anna C. Squicciarini. Standards for Web Services Security. In *Security for Web Services and Service-Oriented Architectures*, pages 45–77. Springer, 2010.

[23] Elisa Bertino and Kenji Takahashi. *Identity Management - Concepts, Technologies and Systems*. Artech House, 2011.

[24] Abhilasha Bhargav-Spantzel, Anna C. Squicciarini, and Elisa Bertino. Establishing and Protecting Digital Identity in Federation Systems, booktitle = Proceedings of

the workshop on Digital identity management, publisher = ACM, pages = 11-19, year = 2005.

[25] David Brossard, Theo Dimitrakos, Angelo Gaeta, and Stéphane Mouton. Aspects of General Security & Trust. In *Service Oriented Infrastructures and Cloud Service Platforms for the Enterprise*, pages 75–102. Springer, 2010.

[26] William E. Burr, Donna F. Dodson, Elaine M. Newton, Ray A. Perlner, W. Timothy Polk, Sarbari Gupta, and Emad A. Nabbus. Electronic Authentication Guideline. Technical Report Special Publication 800-63-1, National Institute of Standards and Technology, 2011.

[27] L. Jean Camp. Digital Identity. *Technology and Society Magazine, IEEE*, 23(3):34–41, 2004.

[28] David Chadwick. Federated Identity Management. In *Foundations of Security Analysis and Design V*, volume 5705 of *Lecture Notes in Computer Science*, pages 96–120. Springer, 2009.

[29] David Chadwick and George Inman. Attribute Aggregation in Federated Identity Management. *Computer*, 42(5):33–40, 2009.

[30] Bart Delft and Martijn Oostdijk. A security analysis of openid. In *Policies and Research in Identity Management*, volume 343 of *IFIP Advances in Information and Communication Technology*, pages 73–84. Springer, 2010.

[31] Tore Dybå, Torgeir Dingsøyr, and Geir Kjetil Hanssen. Applying Systematic Reviews to Diverse Study Types: An Experience Report. In *Proceedings of First International Symposium on Empirical Software Engineering and Measurement*, pages 225–234, 2007.

[32] Ludwig Fuchs and Günther Pernul. Minimizing insider misuse through secure Identity Management. *Security and Communication Networks*, 5(8):847–862, 2012.

[33] Jinguang Han, Yi Mu, Willy Susilo, and Jun Yan. A Generic Construction of Dynamic Single Sign-on with Strong Security. In *Security and Privacy in Communication Networks*, volume 50 of *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering*, pages 181–198. Springer, 2010.

[34] Marit Hansen, Andreas Pfitzmann, and Sandra Steinbrecher. Identity Management Throughout One's Whole Life. *Information Security Technical Report*, 13(2):83 – 94, 2008.

[35] Alan R. Hevner and Salvatore T. March. The Information Systems Research Cycle. *Computer*, 36(11):111–113, November 2003.

[36] Alan R. Hevner and Sudha Ram. Design Sience in Information Systems Research. *MIS Quarterly*, 28(1):75 – 105, 2004.

[37] Heather Hinton and Mark Vandenwauver. Identifying Patterns of Federation Adoption. In *ISSE 2006 Securing Electronic Busines Processes*, pages 151–160. Vieweg, 2006.

[38] Thorsten Hoellrigl, Jochen Dinger, and Hannes Hartenstein. A Consistency Model for Identity Information in Distributed Systems. In *Proceedings of the 34th Annual Computer Software and Applications Conference*, pages 252–261.

[39] Thorsten Hoellrigl, Jochen Dinger, and Hannes Hartenstein. FedWare: Middleware Services to Cope with Information Consistency in Federated Identity Management. In *Proceedings of the International Conference on Availability, Reliability, and Security (ARES '10)*, pages 228–235. IEEE, 2010.

[40] Martin Gilje Jaatun, Eirik Albrechtsen, Maria B. Line, Inger Anne Tøndel, and Odd Helge Longva. A framework for incident response management in the petroleum industry. *International Journal of Critical Infrastructure Protection*, 2(1-2):26–37, 2009.

[41] Martin Gilje Jaatun, Jostein Jensen, Per Håkon Meland, and Inger Anne Tøndel. *A Lightweight Approach to Secure Software Engineering*. A Multidisciplinary Introduction to Information Security. CRC Press, 2011.

[42] Jostein Jensen. Benefits of Federated Identity Management - A Survey from an Integrated Operations Viewpoint. In *Availability, Reliability and Security for Business, Enterprise and Health Information Systems*, volume 6908 of *Lecture Notes in Computer Science*, pages 1–12. Springer Berlin Heidelberg, 2011.

[43] Jostein Jensen. Federated Identity Management Challenges. In *Seventh International Conference on Availability, Reliability and Security (ARES '12)*, pages 230–235, 2012.

[44] Jostein Jensen. Identity Management Lifecycle - Exemplifying the Need for Holistic Identity Assurance Frameworks. In *Information and Communication Technology*, volume 7804 of *Lecture Notes in Computer Science*, pages 343–352. Springer, 2013.

[45] Jostein Jensen and Martin Gilje Jaatun. Not Ready for Prime Time: A Survey on Security in Model Driven Development. *International Journal of Secure Software Engineering*, 2(4).

[46] Jostein Jensen and Martin Gilje Jaatun. Federated Identity Management - We Built It; Why Won't They Come? *IEEE Security & Privacy*, 11(2):34–41, 2013.

[47] Jostein Jensen and Åsmund Ahlmann Nyre. Soa security - an experience report. In *The Norwegian Information Security Conference (NISK)*, pages 185–196, 2009.

[48] Jostein Jensen and Åsmund Ahlmann Nyre. Federated Identity Management and Usage Control - Obstacles to Industry Adoption. In *Eighth International Conference on Availability, Reliability and Security (ARES '13)*, pages 31–41, 2013.

[49] Jostein Jensen, Inger Anne Tøndel, Martin Gilje Jaatun, Per Håkon. Meland, and Herbjørn Andresen. Reusable Security Requirements for Healthcare Applications. In *Fourth International Conference on Availability, Reliability and Security (ARES '09)*, pages 380–385, 2009.

[50] Audun Jøsang and Simon Pope. User Centric Identity Management. In *AusCERT Asia Pacific Information Technology Security Conference*, 2005.

[51] Barbara Kitchenham. Guidelines for performing Systematic Literature Reviews in Software Engineering. Technical report, EBSE Technical Report, School of Computer Science and Mathematics, Keele University, 2007.

[52] Heinz K. Klein and Michael D. Myers. A Set of Principles for Conducting and Evaluating Interpretive Field Studies in Information Systems. *MIS Quarterly*, 23(1):67–93, March 1999.

[53] Katsiaryna Labunets. Empirical Validation of Security Methods. In *Proceedings of the Doctoral Symposium at the International Symposium on Engineering Secure Software and Systems (ESSoS-DS 2013)*, pages 55 – 61. Springer, 2013.

[54] Susan Landau, Hubert Le Van Gong, and Robin Wilton. Achieving Privacy in a Federated Identity Management System. In *Financial Cryptography and Data Security*, volume 5628 of *Lecture Notes in Computer Science*, pages 51–70. Springer, 2009.

[55] Susan Landau and Tyler Moore. Economic Tussels in Federated Identity Management. In *Proceedings of the 10th Workshop Economics of Information Security (WEIS 11)*.

[56] Paul Madsen, Yuzo Koga, and Kenji Takahashi. Federated Identity Management for Protecting Users from ID Theft, 2005.

[57] Eve Maler and Drummond Reed. The Venn of Identity: Options and Issues in Federated Identity Management. *Security & Privacy, IEEE*, 6(2):16–23, 2008.

[58] Per Håkon Meland and Jostein Jensen. Secure Software Design in Practice. In *Third International Conference on Availability, Reliability and Security (ARES 08)*, pages 1164–1171, 2008.

[59] Johan F. Niekerk and Rossouw von Solms. Research Methodologies in Information Security Research: The Road Ahead. In *Security and Privacy - Silver Linings in the Cloud*, volume 330 of *IFIP Advances in Information and Communication Technology*, pages 215–216. Springer, 2010.

[60] Åsmund Ahlmann Nyre and Martin Gilje Jaatun. Usage Control in Interorganisational Collaborative Environments - A Case Study from an Industry Perspective. In *Multidisciplinary Research and Practice for Information Systems*, volume 7465 of *Lecture Notes in Computer Science*, pages 317–331. Springer, 2012.

[61] Geraint Price. The Benefits and Drawbacks of Using Electronic Identities. *Information Security Technical Report*, 13(2):95 – 103, 2008.

[62] Ying Qian, Yulin Fang, Martin Gilje Jaatun, Stig Ole Johnsen, and Jose J. Gonzalez. Managing Emerging Information Security Risks during Transitions to Integrated Operations. In *43rd Hawaii International Conference on System Sciences (HICSS)*, pages 1–11, 2010.

[63] Sebastian Rieger. User-Centric Identity Management in Heterogeneous Federations. In *Fourth International Conference on Internet and Web Applications and Services (ICIW '09)*, pages 527–532, 2009.

[64] Colin Robson. *Real World Research*. Wiley, third edition, 2011.

[65] Everett M. Rogers. *Diffusion of Innovations*. Free Press, 5th edition, 2003.

[66] Richard Sassoon, Martin Gilje Jaatun, and Jostein Jensen. The Road to Hell is Paved with Good Intentions: A Story of (In)secure Software Development. In *Fifth International Conference on Availability, Reliability, and Security (ARES '10)*, pages 501–506, 2010.

[67] Frank Schell, Jochen Dinger, and Hannes Hartenstein. Performance Evaluation of Identity and Access Management Systems in Federated Environments. In *Scalable Information Systems*, volume 18 of *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering*, pages 90–107. Springer, 2009.

[68] Jonathan Scudder and Audun Jøsang. Personal Federation Control with the Identity Dashboard. In *Policies and Research in Identity Management*, volume 343 of *IFIP Advances in Information and Communication Technology*, pages 85–99. Springer, 2010.

[69] Carolyn B. Seaman. Qualitative methods in empirical studies of software engineering. *IEEE Transactions on Software Engineering*, 25(4):557–572, 1999.

[70] Arvind Kumar Sharma and Chattar Singh Lamba. Survey on Federated Identity Management Systems. In *Recent Trends in Networks and Communications*, volume 90 of *Communications in Computer and Information Science*, pages 509–517. Springer, 2010.

[71] Dongwan Shin, Rodrigo Lopes, and William Claycomb. Authenticated Dictionary-Based Attribute Sharing in Federated Identity Management. In *Proceedings of the Sixth International Conference on Information Technology: New Generations*, pages 504–509. IEEE, 2009.

[72] Guttorm Sindre and Andreas L. Opdahl. Eliciting Security Requirements with Misuse Cases. *Requirements Engineering*, 10(1):34–44, 2005.

[73] Mikko T. Siponen and Harri Oinas-Kukkonen. A Review of Information Security Issues and Respective Research Contributions. *ACM SIGMIS Database*, 38(1):60–80, 2007.

[74] Don Smith. The challenge of federated identity management. *Network Security*, 2008(4):7–9, 2008.

[75] Colin Soutar and Joni Brenan. Identity Assurance Framework: Overview. Technical report, Kantara initiative, 2010.

[76] Marc Speltens and Patrick Patterson. Federated ID Management - Tackling Risk and Credentialing Users. In *ISSE/SECURE 2007 Securing Electronic Business Processes: Highlights of the Information Security Solutions Europe/SECURE 2007 Conference*, pages 130–135. Vieweg+Teubner Verlag, 2007.

[77] Inger Anne Tøndel, Martin Gilje Jaatun, and Jostein Jensen. Learning from Software Security Testing. In *International Conference on Software Testing Verification and Validation Workshop (ICSTW '08)*, pages 286–294, 2008.

[78] Knut S. Tungland, Einar Landre, Svein G. Johnsen, Frode Myren, and Paul Carr. Reference Architecture of IT systems for OLF's IO G2. Technical report, March 2008.

[79] Martin Wolf, Ivonne Thomas, Michael Menzel, and Christoph Meinel. A Message Meta Model for Federated Authentication in Service-Oriented Architectures. In *Proceedings of the International Conference on Service-Oriented Computing and Applications (SOCA)*. IEEE Computer Society.

[80] Robert K. Yin. *Case Study Research, Design and Methods*. Applied social research methods series. Fouth edition, 2009.

Appendix A

Selected Papers

Paper I

Jostein Jensen, Åsmund Ahlmann Nyre, "SOA Security - an experience report", Proceedings of The Norwegian Information Security Conference (NISK), 185-196, 2009

SOA security – an experience report

Jostein Jensen, Åsmund Ahlmann Nyre
Department of Computer and Information Science, NTNU
{josteije, nyre}@idi.ntnu.no

Abstract

Service Oriented Architectures are gaining popularity, and are used to realise corporate networks, including healthcare systems. This paper summarises the security standards that are defined for web services, which are the building blocks used for most SOA systems, and security challenges related to developing such systems. Further, a case study is presented and the results related to the security evaluation of this case are given. The security evaluation was performed using grey-box testing techniques on a prototype implementation of a web services based healthcare platform. The tests revealed severe security weaknesses in the platform. The experiences presented in this paper illustrate some of the complexity of developing secure SOA-based systems, and show the importance of having a focus on security throughout the entire software development lifecycle.

1. Introduction

Service Oriented Architectures (SOA) are gaining popularity and in a 2009 survey performed by Computer Economics [1], 58% of the responding organisations said they were making a transition to the service oriented model. Also, healthcare organisations are considering service orientation. One example of such is the National ICT, a Norwegian institution coordinating ICT initiatives in the specialised health services, who suggest SOA as a means to achieve a common platform for those services [2]. New technologies and new concepts often introduces new security concerns, and as pointed out by Epstein et al. [3] this is also true for systems built after the principles of SOA.

Healthcare systems are subject to extensive formal regulations concerning the security (i.e. confidentiality, integrity and availability) of data [4], and studies show that privacy related personal health information is important. In a Norwegian study [6], a total of 85% replied that they feel protection of personal health information is important to them, and a nationwide U.S. survey [7] showed that 70% percent said they were concerned that personal information could be leaked from electronic medical records because of weak security.

The MPOWER project[1] is a research project funded by the European Commission, with the objective to develop a home care service platform. The platform is based on SOA principles and realised through web services. Therefore, the security concerns and requirements related to both SOA and health care systems apply. In this paper we use the MPOWER platform as a case study to illustrate a selection of security challenges faced when developing secure web service-based SOA systems, and discuss how to mitigate these threats.

2. Web service security

The following sections give an overview of security standards that are of relevance for the development of secure web service-based systems, and challenges faced when implementing such systems.

[1] *sourceforge.net/projects/free-mpower/*

This paper was presented at the NISK 2009 conference.

Security standards

Security must be built in as an inherent part of a secure service platform. However, the original standards used to define Web Services, i.e. XML, HTTP, WSDL, SOAP and UDDI do not originally address security issues; they were designed to provide connectivity [8]. Consequently, various organisations have come up with a number of security standards to meet prevalent security challenges.

Some of the standards are particularly interesting for achieving message level security for Web Services: *XML-encryption* [9] can be used to achieve confidentiality protection, while *XML-signature* [10] can be used for creation of digital signatures, and thus integrity protection. Both XML documents as whole and single XML elements can be secured by means of these two. The *WS-Security* standard [11] specifies a set of extensions to the Simple Object Access Protocol (SOAP) to help building secure Web Services. This standard defines three main mechanisms: methods for providing message integrity, methods for providing message confidentiality and descriptions of how security tokens can be included in SOAP messages [12]. Confidentiality and integrity are achieved by means of XML-encryption and XML-signature respectively. *WS-SecureConversation* [13] is a standard defining extensions to the WS-Security standard, and provides a framework for establishing and sharing security contexts, as well as session key derivation. While the WS-Security authentication model is focused on message authentication, WS-SecureConversation is aiming at establishing a security context between two endpoints. Thus, a series of messages can be authenticated, which increases the performance and efficiency of the security mechanisms.

With respect to access control there are other web service security standards that are of interest. *SAML* [14] is an acronym for Security Assertion Markup Language and is a standard, that is *"…an XML based framework for communicating user information, entitlements, and attribute information"* [15]. SAML assertions can be issued by identity authorities and then used e.g to uniquely identify a service requester in a trusted manner. SAML can be used for authorisation purposes, and thus to achieve single-sign-on (SSO) [16] in a web service environments. To define service level authorisation and to implement a security policy for authorisation purposes *XACML* [17] can be utilised.

Trust is an important interoperability issue, and to allow communication between services and actors from different trust domains *WS-Trust* [18], defines extensions to WS-Security. These extensions enable issuance and exchange of trusted security tokens, and additionally the standard defines procedures to broker trust relationships between different trust domains to allow trusted SOAP message exchange.

The implementation of Public Key Infrastructures is described by O'Neill et al [19] to be complex, and the management may be difficult. *XML Key Management Specification (XKMS)* [20] specifies protocols for distributing and registering public keys. The purpose of XKMS is to transfer this complexity to specific, more easily managed services that specialise in PKI management.

SSL [21] and TLS [22] are point-to-point protocols used to provide confidentiality and integrity protection of data between two communicating entities. An initial handshake protocol between the two entities is used to negotiate security context, such as cryptographic algorithm and establishment of a common symmetric key. While all the above mentioned protocols operate at the application layer to support secure Web

Services and secure Web Service communication, SSL/TLS both operate at the transport layer[2]. That is why these protocols cannot provide end-to-end security through intermediary nodes in web service environments. However, web service environments can benefit from SSL/TLS if secure communication is to be achieved between two entities with direct communication.

Security challenges

Web Service infrastructures introduce new threats to web-based applications as well as new challenges when it comes to securing them [19] [23] [24] [25]. Although we acknowledge that general threats to web applications, such as the OWASP top ten list[3], also apply to web services, we restrict our study to only include threats that are characteristic of web services. In this paper we therefore focus on threats and challenges stemming from firewall traversals, publicly accessible business logic, identity management and end-to-end security.

Traditional **network firewalls** operate at the network and transport layer and will not be able to discover threats inside the application layer SOAP messages sent to and from a web service's WSDL interface. Consequently, the application's attack surface is expanded from the user interface alone, to also include the service interfaces for each service/module of the application.

The fact that **business logic becomes publicly accessible** is seen as a challenge**.** A service's WSDL schema defines how to use and interact with the given web service. Yunus and Mallal [24] explains that attackers gain detailed and valuable knowledge of how to attack applications through these schemas. The verbose information given by the WSDL interfaces simplifies an adversaries' process of identifying platform vulnerabilities and weaknesses. Vorobiev and Han [23] also mention this verbosity as a challenge, and that the detailed information about available functions and function parameters easily can be misused.

Identity management is described by O'Neill et al [19] as yet another challenge in Web Service environments. In traditional server-client architectures there is a direct link between users and the web applications. To take the SSL/TLS authentication procedure as an example, users can authenticate directly towards the web application, e.g., by use of certificates. However, in web service environments users will interact through web portals, which are responsible for communicating with the web services on users' behalf. This implies challenges related to how provider services can trust that they admit access to authorised users. User information must be passed between services, and possibly even through a chain of intermediary services in a secure fashion to allow the final provider service to make the correct authorisation decisions. Ideally, a user should only need to authenticate to the system once, and not for each service he wants to connect to.

Transport layer security mechanisms such as TLS or SSL are commonly used to provide data integrity and confidentiality between to communicating entities. The problem arises when the concept of intermediary nodes are introduced. The point-to-point principle implies that messages have to be decrypted at the intermediary node and then

[2] Referred to the OSI model

[3] OWASP top 10: http://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project

re-encrypted before it is forwarded to the final destination. This model cannot guarantee that the intermediary service does not manipulate the data either intentionally or unintentionally, thus **end-to-end security** is not ensured. This limitation is recognised by Anzböck and Dustdar [26], among others, who argue that the security needs to be integrated in web services to provide a necessary level of security for end-to-end transmission of messages.

3. The test case

In this section we will present how the MPOWER authentication and authorisation procedures were specified before the implementation started. In the following figures, the participants in the sequence diagrams represent individual web services.

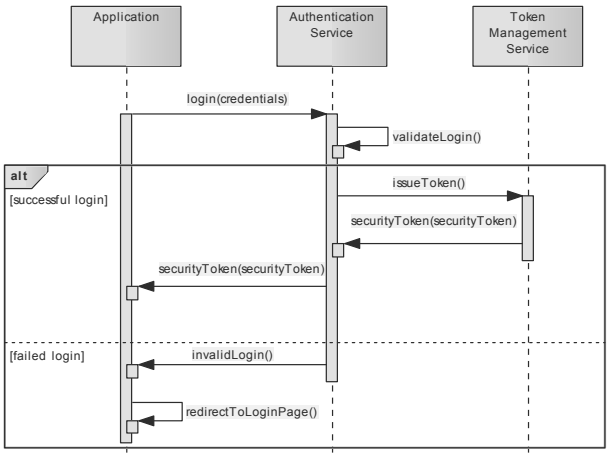

Figure 1: Specified authentication procedure

Figure 1 illustrates the authentication procedure designed for the MPOWER middleware platform. As the message sequence chart illustrates, all users have to perform an initial login before MPOWER services can be utilised (via an end-user application). The authentication process is performed as follows: Users and user sessions are handled by the applications tailored to use the MPOWER platform. These applications forward login requests from users to the authentications service. If the validation process accepts the authentication request, a security token is issued. The Authentication service calls the Token Management service, which generates and digitally signs the security token. The security token is then sent back to the user via the Authentication service. Once the user is authenticated she can use the security token as proof of identity to other services.

The issuance of a security token as described is meant to provide a single-sign-on solution. The security token is used as a session identifier, and a unique identification of the user. By including it in service requests, MPOWER services are able to validate the authenticity of the requesting actor without having to perform the authentication procedure for each new request. The lack of a security token indicates that the user is not authenticated and should not be given access to MPOWER services' resources. The validity of the security token should be time limited to mitigate risks of unauthorised use if it is stolen or somehow leaked to a malign third party.

Figure 2 illustrates the authorisation procedure at the service level. The principle is adapted from the OASIS SAML SSO profile [16]. The illustration is based on the assumption that a user is previously logged in and has a valid security token. The following describes the illustrated authorisation process: A user tries to access an MPOWER middleware service via an application, which sends a service request, including the security token. The middleware service needs to validate whether the user is authorised to access the requested service and operation or not before access is granted. An isAuthorised request including the security token, service name and operation is sent from the middleware service to the Authorisation service. The security token identifies the requesting user, and the roles she holds. It must therefore be checked for validity before further authorisation decisions can be made. An isValid request is sent to the Token Management service to perform this validation. The token's expiry date and its digital signature are central in this validation procedure. If the token validation is successful, the Authorisation service will be notified and then contact the Access Management service to get the access policy defined for the role/s the requesting user holds. Based on the retrieved access policy the Authorisation service checks the permissions for the relevant role/s and performs a permission check that can result in two alternatives: A) Either the request is accepted, and the Authorisation service grants access and notifies the middleware service which again provides the requested content to the application, or B) the request is denied, and the Authorisation service notifies the middleware service of the denied access which sends an access denied message back to the user. A failed token validation either means that a security token was not presented, or that the presented security token was not valid. In either way, the Authorisation service notifies the user via the middleware service and presents a message saying that token validation failed (or that a new login is required).

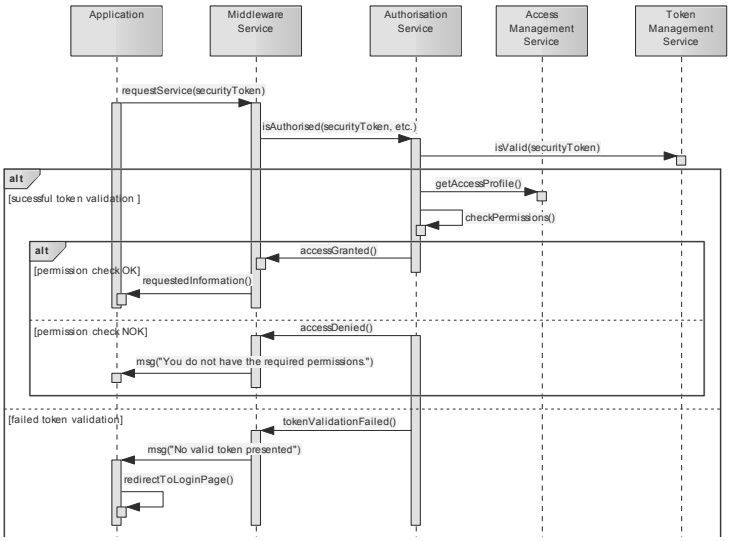

Figure 2: Specified authorisation procedure

Testing methodology

A security assessment of the service platform was performed using grey-box testing [27], which is a combination of pure interface testing (black-box) and software model

testing (white-box). That is, we have used our knowledge of the software model and implementation to find vulnerabilities and then attempted interface testing based on this knowledge. Our tests do not cover every aspect of the system and were not intended for validation of correct behaviour. Our aim was to review how security was handled on a basic level; hence the more advanced test procedures have been left out.

The server providing all web services was running the Ubuntu Linux operating system and GlassFish application server[4]. On the client side we developed our own Python scripts and Java classes for SOAP message generation and sending, while the Wireshark[5] network protocol analyser was used for sniffing and analysing transmitted SOAP messages

Although the prototype implementation includes the security services that were specified in the design, several vital features have been left out. In this section we point at some of the problems of the current implementation, while acknowledging that it is merely a prototype and not an operational system.

Security testing

In this section we present the procedures and results of the tests we have conducted as well as pointing to the potential impact the vulnerabilities that were found may have.

Test scenario 1: Message integrity, confidentiality and authentication

Protecting the integrity, confidentiality and authenticity of messages is particularly important for health care information systems, both due to privacy concerns and risk of incorrect treatment. However, using a network sniffer to look at the SOAP request/response messages sent between a client application and the MPOWER web services, we saw SOAP messages such as those depicted in Figure 3. These clearly show that no cryptographic protection is currently implemented. Adversaries may therefore edit messages at will, without fear of being detected by the system. Entity authentication is done solely through a login procedure, where the user provides a username and password to prove her identity. The login procedure is performed only once per session, and any subsequent proofs of identity are implicit by providing a security token. Although password-based authentication is both widely used and recognised, the lack of password confidentiality makes the process highly insecure. With passwords transmitted in clear text as shown in the figure, password sniffing is trivial.

Test scenario 2: Replay attacks

Replay attacks include situations where an adversary records and resends messages, possibly without knowing their contents. Thus, a message, e.g., containing a prescription for drugs, may be duplicated so that the patient receives more than what the original prescription said, or it could be as simple as to flood an electronic patient record with identical messages in order to hamper reading. Since messages do not carry any cryptographic protection nor sequence number or other data that conveys uniqueness, adversaries may replay any message or sequence of message in order to manipulate services or users.

[4] Available from https://glassfish.dev.java.net/ (version 2.0)
[5] Available from http://www.wireshark.org/ (version 1.2)

We conducted a test where we repeatedly made the same request to add a new activity to a patient's calendar. All repetitions were accepted, which makes the calendar entries difficult, not to say impossible, to read.

Test scenario 3: Access rights elevation

The security token returned during the login procedure, which was shown in Figure 1, contains a list of services and functions available to the user entitled "serviceIDs". Our tests reveal that it is sufficient for a user to simply add services' serviceID in order to gain access to them. Test input was generated through a normal login procedure and then a subsequent request to a service for which the user was not authorised. After manipulating the security token by adding the service's serviceID, the request was retransmitted. While the service correctly denies access in the original request, the manipulated security token is accepted and access granted for the second request (see Figure 4). Although service ID's are not published by services, it is easy to compute them as they normally are formed by concatenation of service and function name.

```
<?xml version="1.0" ?>
<S:Envelope>
  <S:Body>
    <ns2:authenticateUserPassResponse>
      <return>
        <message>Login succeeded!</message>
        <boolValue>true</boolValue>
        <securityToken>
          <userID>16</userID>
          <sessionID>caRSU7cJa7XgQAB/AQFxmQ==</sessionID>
          <primaryRoleName>HealthProfessional</primaryRoleName>
          <authenticationTime>2009-08-21T11:16:21.670</authenticationTime>
          <serviceIDs>CalendarManagement__rejectActivity_A</serviceIDs>
        </securityToken>
        <status>
          <messageId>0</messageId>
          <result>0</result>
          <errorCause>Login succeeded!</errorCause>
          <timestamp>0</timestamp>
        </status>
      </return>
    </ns2:authenticateUserPassResponse>
  </S:Body>
</S:Envelope>
```

```
<?xml version="1.0" ?>
<S:Envelope>
  <S:Body>
    <ns2:authenticateUserPass>
      <username>Anna.Nowacka</username>
      <password>1234</password>
    </ns2:authenticateUserPass>
  </S:Body>
</S:Envelope>
```

Figure 3: SOAP messages for authentication (request in the top right corner)

Our testing revealed that any field in the *securityToken* may be changed, without affecting access control, except the user ID. Through code inspection it is clear that the user ID is looked up in the database and if it exists, and the time since last authentication has not expired, the token is accepted. The values needed for this check are collected from the database, rather than the token. It therefore has no effect, e.g., to change the authentication time.

Although perhaps not the most common of attacks, it is noted that the opposite is also possible; that is to reduce access rights by removing *serviceIDs* from a user's token.

Test scenario 4: Omitting procedures

The MPOWER platform contains no mechanisms for authentication and authorisation of services requiring access to another service. In Figure 1 the authentication service calls a token management service, which in turn issues a security token for the user. The authentication service will only perform the call if the authentication process is successful. However, the token management service does not verify the authenticity of the request, i.e. that it originates from the authentication service. Thus, it is possible for a user to request a token directly from the token management service, thereby omitting

the entire authentication procedure. The only required input in order to successfully issue a token is a valid user ID, see Figure 5. Such background services would of course normally not be published in a directory service, but they are nevertheless accessible and therefore possible to contact. Furthermore, in this case user IDs are integer values that are created using auto-increment in the database. It should therefore not require many attempts in order to find a valid user, for which access rights can be manipulated afterwards.

```
<?xml version="1.0" ?>
<S:Envelope>
    <S:Body>
        <ns2:isAuthorized>
            <token>
                <userID>16</userID>
                <sessionID>aGFsbG9kdWph</sessionID>
                <primaryRoleName>HealthProfessional</primaryRoleName>
                <authenticationTime>2009-08-21T11:16:21.670</authenticationTime>
                <serviceIDs>MyOwnService__MyOwnMethod</serviceIDs>
            </token>
            <serviceID>MyOwnService</serviceID>
            <methodName>MyOwnMethod</methodName>
        </ns2:isAuthorized>
    </S:Body>
</S:Envelope>
<!---------------    Response ------------------>
<?xml version="1.0" ?>
<S:Envelope>
    <S:Body>
        <ns2:isAuthorizedResponse>
            <return>
                <message>Access is granted.</message>
                <boolValue>true</boolValue>
                <status>
                    <messageId>0</messageId>
                    <result>0</result>
                    <errorCause>Access is granted.</errorCause>
                    <timestamp>0</timestamp>
                </status>
            </return>
        </ns2:isAuthorizedResponse>
    </S:Body>
</S:Envelope>
```

Figure 4: SOAP message showing a manipulated security token is accepted by the authorisation service

With the four fundamental security flaws revealed in the above sections, any additional security testing serves no purpose. It is hard to imagine other vulnerabilities that would be easier to exploit or provide more functionality, and we therefore concentrate on these when we discuss possible solutions.

```
<?xml version="1.0" ?>
<S:Envelope >
    <S:Body>
        <ns2:IssueToken >
            <userid>16</userid>
        </ns2:IssueToken>
    </S:Body>
</S:Envelope>
```

```
<?xml version="1.0" ?>
<S:Envelope>
  <S:Body>
    <ns2:IssueTokenResponse >
      <return>
        <token>
          <userID>16</userID>
          <sessionID>cZAwxVvrOJPgQAB/AQEjlQ==</sessionID>
          <authenticationTime>2009-08-20T11:15:19.356</authenticationTime>
        </token>
        <status>
          <messageId>0</messageId>
          <result>0</result>
          <errorCause>Successful operation.</errorCause>
          <timestamp>0</timestamp>
        </status>
      </return>
    </ns2:IssueTokenResponse>
  </S:Body>
</S:Envelope>
```

Figure 5: SOAP message for issuing a token (request is in the top right corner)

4. Discussion

In the section about security challenges related to the use of SOA we listed four key challenges, where one of them are related to publicly accessible business logic, and the verbose information that can be obtained by studying a services WSDL file. Our limited tests can be used to illustrate three related misuse scenarios. A) The issued security tokens contain a reference to all the services and operations a specific user has authorisation to use (see Figure 4), as illustrated in test scenario 3. Since WSDL files are publicly accessible, an attacker can easily collect information about all available services and operations and thereby modify the token to obtain full access to the entire service platform. A simple script scanning WSDL files and modifying the security token is everything an attacker needs to perform this task. B) The WSDL file of the Token Management service revealed that this service had an operation called issueToken with userID as input parameter (test scenario 4). Consequently, it was easy to generate a script calling this service's operation with an arbitrary user ID. Even though this feature was only intended to be used via the authentication service, the service responded to direct calls from a tailored client. The fact that the service's WSDL file was published made the attack easier. C) Even non-existing information in the WSDL file can be used by attackers. In our case the configuration files actually reveal that WS-Security is not enabled at all, as there are no signs of a WS-Security header at all in the WSDL specification.

One can argue that all of these attacks are trivial, and that they could never have been exploited if the security functionality had been successfully implemented. Security by obscurity is not a good design choice, meaning that WSDL verbosity in itself should not be a problem. However, that does not change the fact that the more information an attacker can obtain, the easier it is for him to find the weak spots of a system. These findings confirm the concerns expressed by Younus and Mallal [24] and Vorobiev and Han [23].

A second challenge was related to the problem of Identity Management. In the MPOWER platform the concept of security tokens are introduced to identify the user who initiates service requests and thereby to be able to make access control decisions based on that. The MPOWER solution is inspired by the SAML security standard [15], yet the implementation illustrates several challenges related to managing identities. Test scenario 1 shows the importance of protecting confidentiality of user credentials to provide identity theft, and together with test scenario 3 we see that without proper protection, users can manipulate their digital identities to elevate their privileges. The latter scenario would have been avoided if the security token had been properly integrity protected, e.g., by digitally signing the entire token. Yet, test scenario 2 reveals another problem. Adversaries are able to replay messages into the system, meaning that they can record valid messages including security tokens and feed them into the system to obtain information at a later time. Even if the token was integrity protected a malign user could record it from one message and paste it (including the signature) into a crafted message to obtain access. A possible solution to this would be to protect the integrity of the entire message, including the token so that messages without a valid signature would be rejected by the service platform. This illustrates the complexity of building secure SOA-based systems with single-sign-on functionality.

Our test results do not give indications related to the challenges of obtaining full end-to-end security or problems with network firewall traversals.

The MPOWER middleware platform is only at the prototype stage and therefore by no means operative and fully functional. However, for illustrating common web service vulnerabilities and the challenges emerging from making the shift from monolithic applications to service oriented architecture, the prototype implementation is ideal.

Now that the test scenarios are used to illustrate security challenges in SOA environments, we will look at how the implementation could be improved.

Test scenario 1 - Integrity, confidentiality and authentication: Integrity protection or authentication can be provided through a Public Key Infrastructure (PKI) and XML-based signatures [10]. Also, the WS-Security standard specifies how such XML-based signatures can be included in SOAP messages for web services. Similarly, XML-encryption [9] can be utilised to encrypt messages or parts of messages between web services. The fundamental problem of these approaches is that they either require the deployment of a PKI or that every communicating party holds pre-shared keys, to allow secure hash functions to be used for message integration. History has shown that establishing and using certificate infrastructures is not trivial. The XKMS specification [20], however, aims at simplifying this task.

Test scenario 2 - Message replay attacks: Perhaps one of the main challenges is to prevent replay attacks (record and resend) so that adversaries are unable to repeat requests. There are several available mechanisms to mitigate the risk of such attacks, including: A) Require that all messages carry timestamps, so that messages become invalid after a short period of time. B) Require that all messages carry a counter or sequence number that is strictly increasing, so that messages with equal or lower counter value than the last recorded will be discarded. C) Implement a challenge-response protocol, so that users must prove ownership of a credential for all requests.

While these mechanisms all protect from replay attacks, there are several practical problems in using them. Timestamps require clock synchronisation between all communicating parties, which is not trivial when the amount of users and organisations increase. Additionally, clock synchronisation must also be secured to prevent adversaries from manipulating the time in order to gain access. The counter mechanism is relatively simple, but requires that services (security services) store the latest counter value for all users. Additionally, user clients must be aware of their counter as well, which might be problematic for portable usage (mobile clients, web clients, etc). So scalability is an issue for this mechanism. The challenge-response mechanism is widely used and proven in many secure communication protocols (e.g., SSL/TLS). Users and services verify the authenticity of the other party using randomly generated challenges. Replayed messages will therefore not be accepted, since the contained challenge will not match. The problem with this approach is that it requires additional messages to be passed between communicating parties and hence may introduce considerable overhead in terms of time and bandwidth consumption.

Test scenario 3 – Access rights elevation: As already mentioned, the problem of elevating privileges by modifying the security token could have been eliminated if, e.g., XML-signature [10] had been used for integrity protection of tokens and the messages they are included in. This presupposes that the signature is properly validated when signed messages are received.

Test scenario 4 - Omitting procedures: This problem stems from the fact that normal internal messaging has become public through the use of Web Services. That is, interaction between services in order to fulfil a given task (e.g., login) may be denoted internal communication, but with the advent of web services, this kind of communication is as public as any other. Hence, in order to ensure that procedures are followed, access to secondary (i.e. underlying) services must be restricted to include only those parts of a valid procedure. Again, Public Key Infrastructures or shared secret keys may be used.

Building secure software systems is complex and as indicated in this paper, service orientation and distribution of services adds to the complexity. McGraw [28] and Apvrille and Pourzandi [29], among others, suggest keeping a security focus throughout the entire software development lifecycle. Our results illustrate the importance of their message. Even though this paper does not discuss where and why the errors were introduced, we see that the final security assessment reveal weaknesses that can be removed in the next prototype.

5. Conclusion

In this paper we have given a brief overview of the available standards for securing web services, and outlined some of the challenges faced when implementing secure service oriented software systems. We have exemplified some of these challenges through grey-box testing of a prototype implementation of SOA-based home care system. Our examples illustrate the complexity of building secure web service systems and highlight the importance of integrating security in the entire software development lifecycle.

6. Acknowledgement

Initial security testing of the MPOWER platform, inspiring us to do more tests, was performed by Anja Svartberg in a minor project at the Institute of Telematics, NTNU.

7. References

[1] "SOA Adoption Surges," Computer Economics, January 2009. http://www.computereconomics.com/article.cfm?id=1423&tag=rbspot

[2] "Tjenesteorientert arkitektur i spesialhelsetjenesten," Nasjonal IKT, 2008. http://www.nasjonalikt.no/Publikasjoner/Tjenesteorientert_arkitektur_i_spesialisthelsetjenesten_hovedrapport_full_v1_0e.pdf

[3] J. Epstein, S. Matsumoto, and G. McGraw, "Software security and SOA: danger, Will Robinson!," *Security & Privacy, IEEE*, vol. 4, pp. 80 -83, 2006.

[4] J. Jensen, I. A. Tøndel, M. G. Jaatun, P. H. Meland, and H. Andresen, "Reusable Security Requirements for Healthcare Applications," presented at Availability, Reliability and Security, 2009. ARES '09. International Conference on, 2009.

[5] "Personvernrapporten 2005 - Som søvngjengere inn i et overvåkningssamfunn?," Norwegian Data Inspectorate, 2005, http://www.datatilsynet.no/upload/Dokumenter/publikasjoner/aarsmeld/Personvernrapporten_2005.pdf

[6] I.-A. Ravlum, "Pinning our faith on Big Brother ... together with all the little brothers? ," TØI rapport 789/2005, 2005, http://www.toi.no/getfile.php/Publikasjoner/T%D8I%20rapporter/2005/789-2005/789-2005.pdf.

[7] HarrisInteractive, "Health Information Privacy (HIPAA) Notices Have Improved Public's Confidence That Their Medical information Is Being Handled Properly," 2005, http://www.harrisinteractive.com/news/allnewsbydate.asp?NewsID=894
[8] E. Pulier and H. Taylor, *Understanding Enterprise SOA*: Manning Publications Co., 2005.
[9] W3C, "XML Encryption Working Group," http://www.w3.org/Encryption/2001/
[10] IETF/W3C, "XML-DSig Working Group," http://www.w3.org/Signature/
[11] OASIS, "OASIS Web Services Security (WSS) TC," http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wss
[12] OASIS, "Web Services Security: SOAP Message Security 1.1 (WS-Security 2004)," wss-v1.1-spec-os-SOAPMessageSecurity, 2006,
[13] OASIS, "WS-SecureConversation 1.3," ws-secureconversation-1.3-os, 2007,
[14] OASIS, "OASIS Security Services (SAML) TC," http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security
[15] OASIS, "SAML V2.0 Executive Overview," sstc-saml-exec-overview-2.0-cd-01, 2005,
[16] OASIS, "Profiles for the OASIS Security Assertion Markup Language (SAML) V2.0," saml-profiles-2.0-os, 2005,
[17] OASIS, "OASIS eXtensible Access Control Markup Language (XACML) TC," http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=xacml
[18] OASIS, "WS-Trust 1.3," ws-trust-1.3-spec-os, 2007,
[19] M. O'Neill, P. Hallam-Baker, S. M. Cann, M. Shema, E. Simon, P. A. Watters, and A. White, *Web services security*. McGraw-Hill/Osborne Media, 2003.
[20] W3C, "XML Key Management Specification (XKMS 2.0)," 2005, http://www.w3.org/2001/XKMS/Drafts/xkms-req.html
[21] "The SSL Protocol Version 3.0," 1996, http://wp.netscape.com/eng/ssl3/draft302.txt
[22] "The TLS Protocol Version 1.0," 1999, http://www.ietf.org/rfc/rfc2246.txt
[23] A. Vorobiev and J. Han, "Security Attack Ontology for Web Services," presented at Second International Conference on Semantics, Knowledge and Grid, 2006. SKG '06. , 2006.
[24] M. Yunus and R. Mallal, "An Empirical Study of Security Threats and Countermeasures in Web Services-Based Services Oriented Architectures," in *Lecture Notes in Computer Science*, vol. 3806: Springer Berlin / Heidelberg, 2005, pp. 653-659.
[25] A. Singhal, T. Winograd, and K. Scarfone, "Guide to Secure Web Services, Recommendations of the National Institute of Standards and Technology," National Institute of Standards and Technology Special Publication 800-95, 2007,
[26] R. Anzböck and S. Dustdar, "Modeling and implementing medical Web services," *Data & Knowledge Engineering*, vol. 55, pp. 203-236, 2005.
[27] G. V. Bochmann and A. Petrenko, "Protocol testing: review of methods and relevance for software testing," presented at Proceedings of the 1994 ACM SIGSOFT international symposium on Software testing and analysis, Seattle, Washington, United States, 1994.
[28] G. McGraw, *Software Security - Building Security In*: Addison-Wesley, 2006.
[29] A. Apvrille and M. Pourzandi, "Secure Software Development by Example," *IEEE Security and Privacy*, vol. 3, pp. 10-17, 2005.

Paper II

Jostein Jensen, Martin Gilje Jaatun, "Not Ready for Prime Time: A Survey on Security in Model Driven Development", International Journal of Secure Software Engineering, 49-61, 2011

Is not included due to copyright

Jostein Jensen, "Benefits of Federated Identity Management - A Survey from an Integrated Operations Viewpoint", Availability, Reliability and Security for Business, Enterprise and Health Information Systems, volume 6908 of Lecture Notes in Computer Science, 1-12, 2011

Paper III

Is not included due to copyright

Jostein Jensen, "Federated Identity Management Challenges", in Proceedings of the Seventh International Conference on Availability, Reliability and Security (ARES '12), 230-235, 2012

Paper IV

Is not included due to copyright

Jostein Jensen, Martin Gilje Jaatun, "Federated Identity Management - We Built It; Why Won't They Come?", IEEE Security & Privacy, 34-41, 2013

Paper V

Is not included due to copyright

Jostein Jensen, Åsmund Ahlmann Nyre, "Federated Identity Management and Usage Control - Obstacles to Industry Adoption", in Proceedings of Eighth International Conference on Availability, Reliability and Security (ARES '13), 2013

Paper VI

Is not included due to copyright

Paper 7

Jostein Jensen, "Identity Management Lifecycle - Exemplifying the need for Holistic Identity Assurance Frameworks". Information and Communicatiaon Technology, volume 7804 of Lecture Notes in Computer Science, pages 343-352. Springer, 2013

Paper VII

Is not included due to copyright

Appendix B

Secondary Papers

SP1:
Jostein Jensen, Inger Anne Tøndel, Martin Gilje Jaatun, Per Håkon Meland, and Herbjørn Andresen, *"Reusable Security Requirements for Healthcare Applications"*. Fourth International Conference on Availability, Reliability and Security (ARES '2009), 2009

Abstract:
Healthcare information systems are currently being migrated from paper based journals to fully digitalised information platforms. Protecting patient privacy is thus becoming an increasingly complex task, where several national and international legal requirements must be met. These legal requirements present only high-level goals for privacy protection, leaving the details of security requirements engineering to the developers of electronic healthcare systems. Our objective has been to map legal requirements for sensitive personal information to a set of reusable technical information security requirements. This paper presents examples of such requirements extracted from legislation applicable to the healthcare domain.

SP2:
Richard Sassoon, Martin Gilje Jaatun, Jostein Jensen, *"The Road to Hell is Paved with Good Intentions: A Story of (In)secure Software Development"* Fifth International Conference on Availability, Reliability, and Security (ARES '10), 2010.

Abstract:
Model driven development (MDD) is considered a promising approach for software development. In this paper the results of a systematic survey is reported to identify the state-of-the-art within the topic of security in model driven development, with a special focus on finding empirical studies. We provide an introduction to the major secure MDD initiatives, but our survey shows that there is a lack of empirical work on the topic. We conclude that better standardisation initiatives and more empirical research in the field is necessary before it can be considered mature.

Appendix C

Statements from Co-Authors

NTNU

Encl. to application for assessment of PhD thesis

STATEMENT FROM CO-AUTHOR

(cf. section 10.1 in the PhD regulations)

Jostein Jensen applies to have the following thesis assessed:

Federated Identity Management in the Norwegian Oil & Gas Industry

Statement from co-author on article:

I hereby declare that I am aware that the work in the paper:
Jostein Jensen, Åsmund Ahlmann Nyre: "SOA Security - an experience report", Proceedings of The Norwegian Information Security Conference (NISK), 185-196, 2009 of which I am a co-author, will form part of the PhD dissertation by PhD student Jostein Jensen.

The paper is written as a collaborative effort between the two authors, with an equal contribution during the research and writing phases.

Trondheim 29.10.13
Place, date

Åsmund Ahlm
Signature co-author

Statement from co-author on article:

I hereby declare that I am aware that the work in the paper:
Jostein Jensen, Åsmund Ahlmann Nyre: "Federated Identity Management and Usage Control - Obstacles to Industry Adoption", Eigth International Conference on Availability, Reliability and Security (ARES '13), 2013 of which I am a co-author, will form part of the PhD dissertation by PhD student Jostein Jensen.

The paper is written as a collaborative effort between the two authors, with an equal contribution during the research and writing phases. Jostein Jensen was in lead on the federated identity management research, and Åsmund Ahlmann Nyre on the usage control research.

Trondheim 29.10.13
Place, date

Åsmund Ahlm
Signature co-author

Statement from co-author on article:

I hereby declare that I am aware that the work in the paper:
Jostein Jensen, Martin Gilje Jaatun: "Not Ready for Prime Time: A Survey on Security in Model Driven Development", International Journal of Secure Software Engineering, 49-61, 2011 of which I am a co-author, will form part of the PhD dissertation by PhD student Jostein Jensen.

Jostein Jensen made a major contribution to the work in the research phase, and a proportional contribution to the work in the writing phase.

Trondheim 29/10-13
Place, date

Signature co-author

Statement from co-author on article:
I hereby declare that I am aware that the work in the paper:
Jostein Jensen, Martin Gilje Jaatun: "Federated Identity Management - We Bulit It; Why Won't they Come?", IEEE Security & Privacy, 34-41, 2013 of which I am a co-author, will form part of the PhD dissertation by PhD student Jostein Jensen.

Jostein Jensen made a major contribution to the work in the research phase, and a proportional contribution to the work in the writing phase.

Trondheim 29/10-13
Place, date

Signature co-author