

A Case Study of Enterprise Identity Management System Adoption in an Insurance Organization

Pooya Jaferian, David Botta, Kirstie Hawkey, Konstantin Beznosov

University of British Columbia, Vancouver, Canada

{pooya,botta,hawkey,beznosov}@ece.ubc.ca

ABSTRACT

This case study describes the adoption of an enterprise identity management (IdM) system in an insurance organization. We describe the state of the organization before deploying the IdM system, and point out the challenges in its IdM practices. We describe the organization's requirements for an IdM system, why a particular solution was chosen, issues in the deployment and configuration of the solution, the expected benefits, and the new challenges that arose from using the solution. Throughout, we identify practical problems that can be the focus of future research and development efforts. Our results confirm and elaborate upon the findings of previous research, contributing to an as-yet immature body of cases about IdM. Furthermore, our findings serve as a validation of our previously identified guidelines for IT security tools in general.

Categories and Subject Descriptors

K.6.5 [Management of Computing and Information Systems]: Security and Protection; H.5.2 [Information Interfaces and Presentation]: UIs—*Interaction Styles*; H.5.3 [Information Interfaces and Presentation]: Group and Org. Interfaces—*Collaborative Computing*

Keywords

Identity management, Qualitative Research, Security Tools, Organizational Factors, Case Study

1. INTRODUCTION

Identity management (IdM) comprises the processes and infrastructure for the creation and maintenance of user's digital identities and the designation of who has access to resources, who grants that access, and how accountability and compliance are maintained [5, 7]. While IdM can be studied in different contexts (e.g., the Internet), the scope of IdM in this paper is enterprise identity management. Enterprise IdM includes tasks such as managing identities of the organization's users, managing roles through their life-cycle,

assigning identities to roles, determining the resources each role has access to, and the auditing and reporting of information related to IdM in an organization.

IdM has become an important aspect of IT security infrastructure in organizations, and some consider it to be the most important solution for enabling compliance with legislative requirements [19]. Further drivers of IdM adoption include cost reduction, better security, better access to information, and better agility during mergers and acquisitions [13]. However, the practice of IdM is challenging, both organizationally and technologically [13, 19]. Identifying these challenges and studying how they can be addressed are important steps toward improving IdM systems and practices in organizations.

Despite the widespread and increasing adoption of IdM solutions, there are few available case studies that examine the practice of IdM in organizations. Two major studies of IdM are the Identity Project [19] (an academic survey of IdM practices in UK higher education institutions) and reports by the Burton Group [1] (a firm that provides IT research and advisory services to private clients). Both of these studies blend the results of their case studies into their reports and give recommendations for improving IdM practices and systems. Most recently, Bauer et al. [4] describe real life challenges in access control management as gleaned through interviews with policy professionals. Although not explicitly about IdM, Heckle et al. [9] discuss organizational challenges in implementing a single sign-on system without previously assisting end-users to develop an accurate mental model. Also, Post et al. [14] identify security controls as factor that interferes with end-users' work and propose recommendations for alleviating this problem. In order to improve the usability of IdM systems, or propose new development, more case studies are needed to illuminate nuances of the issues that are already indicated by prior research, and to reveal topics for further research.

In this paper, we present a case study of an insurance organization that has recently made two phases of a multi-phase integrated IdM system operational. IdM is an important part of this organization's security infrastructure. Our case study contributes to an as-yet immature body of cases about IdM. It provides a realistic picture of an organization during the various stages of deploying an IdM system. This picture helps identify practical problems that can be the focus of future research, and that suggest opportunities for new development. In addition, we describe the organization's challenges and discuss their IdM solution in

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHIMIT '09, November 7-8, 2009, Baltimore, Maryland, USA.

Copyright 2009 ACM 1-60558-572-7/09/11 ...\$10.00.

Table 1: Participant roles

Participant	Role
SA-lead	Security Administration group leader
IdM-lead	IdM Project leader
SecA	Security Analyst
SSecC	Senior Security Consultant

comparison to related work such as the access control management challenges identified by Bauer et al. [4] and findings and recommendations from the Identity Project [19] and the Burton Group [1]. Furthermore, we discuss and validate a framework of guidelines and recommendations for IT security tools [10]; these were selected from the literature based on their relevance to the practice of IT security management in general and we illustrate through the case study findings how they can be applied to IdM tools in particular. This information may be used by practitioners who must identify requirements for and evaluate IdM systems, as well as by developers to improve their IdM products.

The remainder of this paper is organized as follows. We first explain our methods for performing the case study and analyzing the data in Section 2, before laying out the findings of our case study in Section 3. In Section 4, we discuss the findings from our case study and compare them with those from prior research; while in Section 5, we examine how our case study findings validate our previously developed guidelines for IT security tools in general. Finally, we conclude in Section 6 with the limitations of our research and Section 7 with a summary of our contributions and a discussion of our future work in this area.

2. METHODOLOGY

We performed four semi-structured interviews with people from the Security Administration (SA) group who were involved in the selection and/or deployment of an IdM system in an insurance organization. The SA group leader (SA-lead), IdM project leader (IdM-lead), and Security Analyst (SecA) have been involved during the entire process of IdM adoption in the organization, while the Senior Security Consultant (SSecC) was only involved in the selection process (see Table 1 for the role key). Each interview lasted from one to two hours and was conducted by two interviewers in the workplace of the participant. The interviews took place at different stages in the IdM deployment process (Fig. 1). Two interviews (SecA, SSecC) were conducted in 2006 and 2007 as a part of our previous project (HOT-Admin, see [8] for an overview of the themes of analysis), and two interviews (SA-lead, IdM-lead) were conducted in late 2008 to study IdM in particular. While the focus of the two HOT-Admin interviews was on IT security management in general, the importance of IdM in the organization led our participants to answer many questions in the context of IdM. The longitudinal aspect of data collection enables us to describe the state of the organization at different stages of IdM adoption.

Interviews were audio-recorded, transcribed, and analyzed by two researchers. In addition to the interviews, the researchers analyzed the Request for Information and Qualification (RFIQ) that the organization issued when beginning the IdM selection process. This documentation of the orga-




Figure 1: Interview timeline that shows when the interviews were conducted during the course of the IdM adoption process

nization’s requirements for the IdM system allowed for triangulation at the level of data source, mitigating some bias inherent in the self-report data of interviews. Using two researchers provided triangulation at the level of analysis and increased confidence in the shared findings.

For analysis, we performed a combination of *direct interpretation* and *categorical aggregation* as recommended by Stake [15]. During direct interpretation, we focused on a single instance of data, pulling it apart and re-shaping it in a more meaningful way. Conversely, during categorical aggregation, we put together snippets of information about a phenomena from different parts of a single interview or different interviews and elaborated the meaning of the phenomena based on the different evidence. We also looked for certain *patterns* by finding *correspondence* between different categories. To perform the above mentioned analysis, we employed both high-level and open coding techniques. We first coded the data in high level categories including “Stakeholders”, “Tools”, “Challenges”, “Tasks”, and “Interactions”. We then performed open-coding inside each category and coded the data with more fine-grained codes that emerged from the data itself. After coding the data, we revised our codes by collapsing similar codes into a single code. As an illustration, to study the “Role Discovery” challenge, we examined the problem by aggregating snippets of information about role discovery. To analyze the role of managers in IdM adoption challenges, we looked for coinciding codes of “Managers” and “Challenges”. This coding practice was performed using Qualrus 2.1, a qualitative analysis software that provides the ability to perform categorical aggregation and find correspondence.

3. IDM CASE STUDY

3.1 The organization

The participating insurance organization had about 2,500 employees at the time of the study. Approximately 2,000 of them worked in the head office, and the rest in branch offices. The processing environment included a single IBM mainframe (z/OS), more than 200 Intel-based Microsoft Windows 2003 servers, and several UNIX (AIX) servers. The personal computers that were in use at both the head office and branch offices numbered around 3,000.

The responsibilities of the organization’s central IT security group included developing policies, standards, and practices related to IT security, as well as managing digital identities and computer-related access control. At the time of this study, the organization had already contracted an IdM software vendor, and the IT security group had implemented the first two phases of their IdM plan: (1) self-service password recovery, and (2) *basic* provisioning, in which all employees

had access to commonly needed services like e-mail and the Internet. Further phases were in development.

Management of digital identities was centralized, while access control was distributed. The SA group handled the creation of digital identities, primarily using Active Directory. They also controlled access to certain kinds of information on the mainframe using Resource Access Control Facility (RACF), a software product from IBM that provides access control functionality in mainframe servers. The RACF rules indicate transactions at the level of software applications, and what user groups are permitted to access those transactions. However, access to local area network (LAN) resources, the internet, databases, etc., was controlled by departmental administrative groups.

Prior to IdM deployment, the SA group developed the “Data Guardianship Policy” that governs who can access a resource in the organization. The SA group developed a list of resources in the organization. Each resource is associated with an owner, called a “data guardian”. Normally, the data guardian for a resource is the manager of the business unit to which the resource belongs. The data guardian is responsible for keeping that resource safe by deciding who can access it. As the manager may not have sufficient awareness about the necessary security requirements for a resource, or the time to respond to access requests, he or she can delegate authority to another person in the business unit. In this case, the person to whom the authority is delegated is called the “data steward.” The data-guardianship policy in our case study organization is an official policy which is documented and strictly enforced. The SA group performs periodic training for data guardians and data stewards to raise their awareness of policies and procedures.

The organization also maintained a separate IdM system that allowed its customers to track their insurance claims through a website. The separation was to ensure there would be no accidental leakage of information.

3.2 Before deployment of an IdM System

We now describe the state of the organization before deployment of an integrated IdM system, including the IdM processes in place and the challenges these processes posed.

3.2.1 The basic IdM process

The basic steps in the organization’s prior identity management workflow (before full deployment of an integrated IdM system) were as follows:

1. Human resources created an ID for a new employee.
2. Both the SA group and the employee’s manager were notified. This notification was automatic, however, the SA group manually processed the ID.
3. A security administrator provided basic permissions to access the systems that are common to all employees, such as e-mail.
4. The employee’s manager requested access to the systems that are appropriate for that employee. This means that the manager needed to know which systems to request. The request was made by means of an electronic form, and was made to the SA group. Similarly, the manager might request additional, possibly temporary, access for an existing employee.

5. If the request was related to RACF and Active Directory (Application Layer Access), a security administrator deployed the request to the data guardian of the requested data. The data might be distributed, and thereby owned by several data guardians. In this case, the security administrator made multiple requests. If the request was not related to RACF and Active Directory, the security administrator forwarded the request to the pertinent DB/LAN administrator, who could implement the access.
6. The data guardian might delegate the request to a data steward.
7. The security administrator performed a follow-up cycle, to handle non-response or lag from data guardians.
8. If the data guardian or data steward granted permission, the security administrator would implement the access in RACF and Active Directory.
9. When an employee was terminated or his status changed, the employee’s manager was responsible for notifying the SA group.

3.2.2 Challenges in IdM workflow

Challenges in the process motivated the organization to adopt an IdM system. To begin with, there could be a significant lag between when a new employee began work and when Human Resources completed processing the new employee and creating an ID (step 1), thereby reducing employee productivity. Once the ID was created in Active Directory, the SA group had to then repetitively enter it in other systems (step 2) – there was no central repository for IDs. For example, the same information had to be entered in RACF, UNIX, SQL databases, and so on (SA-lead, SecA).

In step 4, the SA group provided a detailed multi-page form for managers to identify resources and request access. The managers were overwhelmed by the amount of information that they were expected to fill in (SecA). As already mentioned, a manager who requests access for an employee must have knowledge of the systems, and not all managers knew what to ask for. Frequently, managers would ask the SA group to make a new employee’s access the same as some other current employee’s access: “A manager who hired a new employee who knew that you had the access that you needed to do the job for him or her would say, ‘Oh, make this new employee’s access just like yours.’ And so then an employee would then inherit very high levels of privileges and access based on the success of a previous employee in terms of doing that job. [...] Historically here if you were an individual who started at [the organization] in 1930 (I’m exaggerating) by the time you retired and had 40 years you would have access to every single system that you had ever used in your entire lifetime with” (SA-lead). However, the SA group disallowed this kind of generalized request as it might provide more access than required for the employee; they required the managers to indicate the individual desired systems: The managers who didn’t know what to ask for then tried to work around the security requirement by asking for the list of accesses possessed by a current employee, and requesting access for every item on that list.

Obtaining knowledge about access profiles (groups of access privileges) and presenting it to the requesting managers was in itself a significant challenge. The organization developed an application to automatically collect information about

access privileges from heterogeneous systems and put them in a database. The database grew very quickly and faced performance and availability problems (SecA). The organization had 400,000 RACF rules that were created by generations of security administrators over nearly 20 years. Each generation created rules in its own way. Succeeding generations, when failing to make sense of how the previous security administrators made rules, would invent their own way of making rules (SA-lead).

Needless to say, the managers did not necessarily understand the esoteric access rules and system names – these had to be translated into language that the managers could easily understand: “[...] and we get this huge profile - here’s all the access the user has. We have to translate that into more of an English format for the individual, saying this means this, this means this, and this means this. And then, they have to put that into a form of what they want” (SecA). Furthermore, there was no common terminology throughout the organization for requesting access, resulting in cases where it was difficult for the SA group to identify which systems in which divisions were requested: “[...] but it’s terminology we don’t know. We say what area is this in, is this in the assessment part, is this prevention, is this the claims? - sometimes the user has no idea, they just say I don’t know I just have to get access to it. So we end up going to the security coordinators for all those divisions saying are you familiar with this app? Or whatever they are talking about. Oh yeah, this is actually this. Aha. Then we know, and then we can tell if the request has come to us in a form or an e-mail and it’s approved, we can set up the access or send it to the area responsible for setting up the access” (SecA).

In step 5, when the security administrator processed the request and deployed it to various data guardians, it was often difficult to identify the correct data guardians. Also, as mentioned, some data guardians would not respond in a timely fashion, or even not respond at all (making step 7 necessary). The security administrator could end up implementing the requested access in a piecemeal fashion, thereby decreasing employee productivity as the employee would then have to wait for access: “[...] and so the security administrators would then send notes, e-mail, to the data guardians saying, ‘Are you OK if we grant Bob or Jill access to the system?’ And then there would be that sort of follow-up cycle where the data guardian would ignore them or not respond or say I’m not the data guardian or that kind of crap. So there would often be a delay in terms of getting approval... so systems access to the employee would then build over time as individual data guardians would respond” (SA-lead).

When an employee changed his or her role or department, the SA group would notify that employee’s new manager, and ask the manager to revise the employee’s access privileges. Due to a lack of accountability and/or lack of understanding the access profile of the employee, often the employee would accumulate unnecessary privileges (SA-lead, IdM-lead, SecA, SSecC). This issue was compounded by a trade union mentality: “So some people view that as an infringement upon their union rights. You can’t take things away from me. I have seniority. You can add to me, but you cannot take away from me. They don’t understand like the security concept of you’re doing this job now, you’re not doing this job, you don’t need that access anymore” (SSecC). Concerning temporary access, the SA group manually set re-

mindings in their e-mail client software about when revocation of access was due to happen. Sometimes the SA group forgot to either set a reminder or to act on a notification of a reminder (SecA).

It was the managers’ responsibility to notify the SA group about termination or changes in the status of their employees (step 9). However, as managers did not suffer negative consequences for failing to notify the SA group, terminations often went unreported. The SA group compensated by obtaining automated notifications from Human Resources. However, the automated notifications were not entirely effective, because there are various stages of termination: “We were getting status change notifications regarding employee departures... not universally 100% effective because what can happen is employees can not quite depart. Meaning that you can go on severance, which means that you have a year, two years worth of severance and you haven’t departed” (SA-lead). To solve these issues, the SA group would periodically query the system for accounts that had not been used for a while, and manually follow up on the results (SA-lead, IdM-lead).

Apart from the steps of providing or revoking access, the SA group faced challenges around both audits and troubleshooting. It was difficult to perform audits on groups, because groups in RACF were highly granular. Each access request corresponds to multiple transactions and each transaction can be accessed by multiple groups. Therefore, the SA group was required to respond to an access request by making the user a member of different RACF groups; the combination of groups provides access to all required transactions. This made it difficult for SA group to later ascertain which group was associated with some requested access, altogether resulting in the personnel sometimes creating a new group, compounding the complexity. The issue was further magnified when the SA group would send reports to the data guardians about what groups the data guardians “owned”, the users in those groups, and the transactions to which the users had access. The reports were almost impossible for the data guardians to audit as they contained a great deal of RACF information. Periodically, the data guardians would ask for a list of who had permission to access to their data, and the SA group would “give them reams of paper, and that would confuse them” (SA-lead).

Besides problems with determining who can have access to data, determining who accessed the data was almost impossible. For example, if a data guardian wanted to know who looked at a particular insurance claim, the SA group would not be able to easily identify the person. Not all the systems generated audit logs, because of the negative impact on performance (SA-lead). A similar issue occurred when someone would act as the backup for a manager and be expected to not abuse the manager’s access privileges. To double check whether or not abuse had occurred, the SA group would need to know whether the backup person actually accessed inappropriate data. Again, the lack of audit logs hampered this kind of investigation.

In our case study organization, access requests were made by e-mail, and the historical record was of poor quality. Gaps in the historical record could result in miscommunication: “Well the real serious ones are where you go to a data guardian and say okay this is what the situation is and

I wonder if this client can have access to this kind of data for this length of time and they say sure, then you go and give them the access, but you didn't get it in writing [...] and this happened to me, once, I didn't get an email [from the data guardian], I didn't say hey, by the way can you send me an email that says you approve this, I forgot that part of it. I got the access granted, and everything happened [...] then halfway through, oh why is this access granted, um I came and explained to you, no you came and explained to me but I don't remember saying yes, I was going to think about it, take it away, take it away. So we had to you know, and I got in a little bit of trouble and that was that.” (SSecC).

When a user would initiate a trouble ticket about a lack of expected access, the SA group would have to manually check RACF, Active Directory, and so on, resulting in a heavy workload. This also is related to audit logs and traces not being recorded due to performance issues. Tracing every action would result in terabytes of mostly useless information (SecA).

Concerning end-user issues, sometimes a manager would inappropriately delegate authority to an employee (i.e., by giving out his or her password) in order to have the employee perform a task, despite there being a process for delegation: *“I think some of the executive are some of the worst offenders of doing some of that stuff [bypassing the process for requesting access]. We've had some executive that have given their own accounts to their admin assistants to process transactions when they're away on vacation. Which is a clear breach of our policy which says that you should not, and will not do that”* (IdM-lead). Also, users forgetting their passwords comprised the primary call volume of the support center.

To summarize, managing accesses and identities before deployment of the IdM system was challenging. The challenges impacted the creation of IDs, access requests, provisioning of users, ongoing management of accesses, access terminations, auditing and troubleshooting, and daily tasks of users like authentication or delegation of authorities. Multiple stakeholders in the organization were impacted by the challenges, including security practitioners, managers, data guardians, and employees. Addressing these challenges was the main motive for the organization to deploy an IdM system.

3.3 The IdM deployment process

We now depict the selection, deployment, and configuration of an integrated IdM system in the studied organization. We describe the prerequisites for deploying the IdM system, the process of selection a system (including system requirements and evaluation of different IdM system options), and the challenges encountered during deployment and configuration.

3.3.1 Pre-requisites for deploying an IdM solution

Having a well-defined access control policy and a business process for executing that policy is a pre-requisite for successfully deploying an IdM system (SA-lead, SecA, IdM-lead). As the SA-lead explained: *“it's a classic thing where you buy a tool and you think it's going to solve your problems but if you don't have the staff, and you don't have the business routine already internally to grant and manage access then an identity management system isn't going to help*

you.” The IdM-lead further identified deployment prerequisites: *“there's a wide variety of things that definitely need to be in place before you even look at going down the pure technology implementation. Understanding your own processes when it comes to identities, access administration, the various identity repositories you have is another very important step as well your HR processes, because that's usually your integration point into the identity lifecycle of where they start and where they end and then understanding the tolerance or interaction from the business into your current processes as well.”*

In the studied organization, the SA group had developed a “data guardianship policy” for managing access to resources (see section 3.1). In addition, HR had a comprehensive list of jobs, which could be used as a basis for defining roles (RFIQ). The organization also had a well-defined business process for creating IDs, provisioning, and management of accesses: *“The critical piece for us was understanding our processes and understanding the user IDs and the identities in your own organization. Fortunately for us, once again, through standardization early on all our identities and all our repositories are the same; so when it came to going to explore those repositories for matches and stuff like that, it was easy to do, there's a lot of organizations that don't enforce that standardization, or through acquisition of a another company the standards are different and so they merge their information systems and they get into a bit of a problem where, you know, you've got J Smith in one system and John Smith in another and they've just got miss-matches, so there's a bit of a challenge there”* (IdM-lead). As a result, the organization was ready to introduce a new technology that would employ available business processes to support IdM more effectively.

There was also an understanding of the need to involve other stakeholders in the organization before deployment of the IdM system. In particular, it was felt that the stakeholders of the system should be made aware of the benefits of the new technology and understand what the system will offer (IdM-lead, SecA). The SA group followed two strategies to create awareness. First, through e-mail, they contacted all stakeholders who were affected by deployment of the new technology. They provided the stakeholders with a brief overview of the upcoming changes. Additionally, the organization's internal news website was used to inform employees about the new technology. The second strategy was to deploy the system incrementally. The SA group deployed and made small pieces of the system operational in each iteration. This made the transition to the new technology smoother. The incremental approach showed the benefits of the IdM system to the stakeholders early in the deployment process: *“One of the big ones right now is password self service. People today have to call the help desk to get passwords reset. Now the user can do that themselves, or will be able to within another two weeks”* (SecA).

Availability of staff that could handle the deployment of the IdM system was also a pre-requisite for the project. The SA-lead noted that they hadn't started the project until one of their staff attained the knowledge to act as a security business analyst. He defined a security business analyst to be a person with knowledge in the areas of both computer security and business. The SA-lead creatively fostered the organization's in-house development of security business an-

alysts, because he was unable to find and hire one, despite his significant resources: “[...] rare as hen’s teeth are people who are security business analysts which are people who function as business analysts, perhaps they come from applications development or some other job, but who function as business analysts in terms of information security ... you can’t hire those people for love or money. I had ads in the paper, I had ads on monster, I had ads on any kind of electronic bulletin, I put an ad in the globe and mail for [expletive], which is like going back 20 years. You cannot hire security business analysts and so you have to grow your own” (SA-lead).

The SA group was interested in having a reliable estimation of the project’s cost before starting the project. On the other hand, the vendors were reluctant to give a fixed price for their IdM system and to commit to this price until the end of the project: “We had a fairly big paranoia that a lot of the vendors said, ‘Well we can’t really give you a fixed price or any real definitive number on an implementation for what you’re asking for because we don’t know your organization, we don’t know the complexity, we don’t know various factors that may drive that cost’ ” (IdM-lead). To address this, the SA group hired a consulting firm to analyze the current state of the organization and to provide the IdM system vendor with a report of expected complexity of the project. Having this report, the vendor committed to provide its IdM system with a fixed price.

Finally, the SA group manager stated that they did not start this project earlier because IdM systems were not mature, and the cost of the solutions could not justify their benefits: “[...] the tools weren’t ready; even now, in some cases, I think we’re about 3 weeks behind the developer for the vendor, [Vendor] ‘oh yeah, you can do this now...’ oh, cool that’s what you sold us on 2 years ago, but it’s nice to know we can do it now.” (SA-lead).

3.3.2 Selecting an IdM system

The stakeholders who were involved in the selection process were the SA-lead, the SSecC, the IdM-lead, the SecA and people from the IT department (for reviewing server infrastructure) (IdM-lead). Implementation of an IdM system depends on a good understanding of the business processes [1]. Nevertheless, no one from the business side participated in the selection process. However, it is likely the presence of a security business analyst (who was groomed in-house for the job), helped to compensate for this lack.

The selection process consisted of multiple steps (SA-lead, IdM-lead). First, the selection team developed a set of requirements, on which they based their RFIQ. The requirements for the IdM system can be classified into functional and nonfunctional requirements. The functional requirements include centralized role-based access administration, role mining, integration with available infrastructure, workflow support (integrated workflow engine), and auditing and reporting. From the end-user point of view, requirements include self-serve password management, plus access request and delegation. Nonfunctional requirements of the system include customizability, scalability to about 5000 users, disaster recovery (including the ability to backup), performance, availability, and reliability (including 99.99% uptime), working over slow or unreliable links, and fail-

ing gracefully while working with different repositories (e.g. RACF and Active Directory).

After publishing the RFIQ, they received and processed paper-based proposals from vendors. The score by which they evaluated the proposals consisted of the vendor’s company profile, qualifications and experience (9%); implementation approach and timeline (10%); ability to meet functional requirements (40%); ability to meet non-functional requirements (25%); and estimated cost (16%). The SA-lead highlighted the importance of easily integrating the IdM system with current infrastructure in their decision making process when discussing an unsuccessful vendor proposal: “[...] it had to do with compatibility with our existing infrastructure. That sounds simple but for example, [IdM vendor], has an identity management system but we don’t do [vendors proprietary database technology X], we’re not an [X] infrastructure, we don’t do [X], we do [alternative technology Y] and [alternative technology Z] and so the [IdM vendor] folk we found difficult to be compatible with our infrastructure because their proposal basically said, ‘you have to convert to [X].’ and the answer is, ‘no, sorry, go away.’ ”

The scoring narrowed down their focus to three vendors, who were then invited to present their system. From those three, two were selected to provide their software for lab testing to evaluate whether they could do what their vendors claimed (IdM-lead). This testing took about two and a half months. The SA group did not test every single feature of the systems, but they did conduct a full installation and tested functions that were important to the organization.

The remaining two candidate vendors were scored on the presentation (20%) and the technical validation (80%). Distinguishing features that determined the winner were the ability to analyze and mine roles (building roles based on existing user-permission relationships), usability, and integration with current infrastructure. None of the candidates had integrated role mining feature in their products (SA-lead, IdM-lead), and this was considered a critical feature: “We actually came close to collapsing the entire bid at one point because neither vender had anything on the role management side or role analysis side, they had recommended we manage roles within their own products themselves but the way they managed roles was confusing, it wasn’t what we were trying to achieve.” (IdM-lead).

3.3.3 IdM deployment and setup process

Deployment of the IdM system was an incremental process. Two operational milestones identified were the password self-service sub-system and role based management of accesses. The SA group decided to deliver password self-service in the first stage to show the benefits of the IdM system to end-users and to obtain management support for the rest of the project.

The hardware infrastructure was provided by the organization, and they used platform virtualization to deploy the IdM. At the time of writing, the organization ran three different servers on two boxes using VMWare. The deployment was performed collaboratively by the SA group and an integrator from the vendor (IdM-lead).

Before starting the deployment phase, the vendor claimed that they could deploy and configure the IdM system in three months, but the SA group planned to finish the project

in one year. In actuality, the project took around fifteen months complete. The SA group manager (SA-lead) believed that the project went over schedule due to their lack of awareness about the full scale of the project. He also noted that the vendor's flexibility in providing support during deployment and configuration was key to keeping the overshoot down to a reasonable level. The vendor did not require the organization to pay for technical support during deployment, which helped the SA group to have technical support on demand. As a result, they could deploy the IdM system without an excessive budget for on-site support.

The IdM-lead described IdM deployment as an extremely complicated process, attributing some of the complexity to the vendor's growth: *"oh it's extremely complicated. And that's one thing actually the vendor for their next release of the product just tried to resolve; they recognize it is an issue for their clients and their customers. And this I largely believe this is through the acquisition path they've gone through to get to where they are today. But the product itself resides on multiple servers and has multiple components, it's own internal directories and databases for auditing that are completely separate [...] And that's partly just because they haven't had enough time to integrate their acquisitions further into their core product"*

To enable role based access control, a complete and correct set of roles needs to be created. The creation of the roles was the biggest aspect of setting up the IdM system for the insurance organization. The SA group decided to build roles by following both a top-down and bottom-up role engineering approach (see [16] for a discussion of role mining). They started the role engineering process by developing a set of roles from existing user-permission assignments. This was a bottom-up approach that required mining existing user-permission assignments in different access control repositories. The SA-lead highlighted the importance of discovery in role-mining: *"if you don't know how to do discovery, if your tool can't do discovery you're committing the staff to 2-3 years work of heavy lifting to do discovery. So a tool that did discovery and managed roles potentially can save you years of effort."* The role mining engine in their IdM system could analyze different repositories in each system and find users with similar accesses (SecA). Consequently, the SA group collaborated with individuals from each business area to check the differences between those similar accesses and create a single role that corresponds to a single job description (a top-down approach). Additionally, the SA group collaborated with data guardians to determine which roles should be authorized to access each resource in the organization.

3.4 Benefits and challenges

Our participants expected their IdM system to have several benefits, mainly reduced workload and role-based auditing. But the IdM system also brought its own challenges. At the time of this writing, only two components of the IdM system had been made operational throughout the organization, so the final verdict is still to come.

The automation of provisioning was expected to speed up the process and thereby reduce unproductiveness of employees who were waiting (up to a week) for access to resources. The automation was also expected to reduce the workload

of the SA group.

Upon a new employee's enrollment in the organization, HR creates a new entry in their system, and the HR system triggers an enrollment event, in which the employee is automatically (1) assigned a role (corresponding to the HR job code) in the IdM system, and (2) provided with access to basic things like e-mail. In the next phase, the automatic provisioning will include access privileges that are associated with the role. Similarly, when changing a role (e.g., an employee moving to a different department), access privileges will automatically change based on information in the HR system. Further provisioning is requested by the employee's manager through an online form, bundling many individual requests into one step, thereby reducing the workload on the the SA group. In the next phase, the form will be workflow driven, and automatically inform data guardians of requests, thereby further reducing the security administrators' workload of managing the requests. Additionally, a self-serve password feature, that enables end-users to reset their passwords by answering a set of challenge-response questions, is expected to dramatically reduce the number of calls to the help desk.

Better reporting and compliance was another expected benefit from the system. The old reports about who can access a resource were difficult for data guardians to read and understand. For example, they included cryptic RACF rules. With the new IdM system and role based access control, data guardians will be provided with a list of roles that have access to their data. Furthermore, the IdM system would generate reports by mapping the cryptic or technical terms in the rules to business terms. It was expected that this would make it easier for the organization to observe compliance.

Automated role-based access was expected to enable both end-users and security administrators to be more critical about roles. End-users could see a catalogue of potential access privileges that they could request, and security administrators could identify and correct inappropriate privileges. That is, security administrators would be "freed up" to become security analysts, engaging in questions of how roles should be built and structured, which was expected to be beneficial to the organization in terms of both effective identity management and employee retention. The SA-lead highlighted this benefit: *"[It] allows the function of a security administrator to become smarter in the sense that they are now using the IdM system to grant access and they're now able to do more consideration in terms of, 'is this the right access, can I do an investigation? do I need to build a role?' and so hopefully, I don't plan on reducing the number of the security administrators but I plan on requiring them to behave differently in the sense they're becoming more sophisticated in how they deal with things. So identity management helps us from an enterprise level in terms of managing identities better but also helps me from a business point of view that it frees up staff resources."*

We now summarize the main challenges encountered to date by our case study organization during its deployment of the IdM system. Firstly, role engineering has remained a challenge. Despite the role mining component, creating well-defined and structured roles in the organization required collaboration with the business side of the enterprise. The

collaboration comprised the major step in role engineering, far out-weighting the technical aspects: *“The two biggest areas are, depending on what you’re trying to achieve with your identity management project but for us it’s been the role analysis and the working with the business for those role definitions takes a significant period of time. Especially if you want to get it right. And coming up with a plan, organizationally of how you want to structure roles, coming up with what those common attributes are, those types of things”* (IdM-lead).

Secondly, deploying the loosely integrated components of the IdM system was a challenge. The fragmented components required multiple boxes for deployment. Configuring and running all these components was not easy. Furthermore, updates included multiple executables, each of which could require up to a day to install. This state of affairs was likely due to many acquisitions in the IdM vendor market [1]. One of the IdM vendor’s partners was acquired by a competitor, resulting in one of the components being at risk of losing support in the future.

Thirdly, the practice of rehearsal and planning has been challenging; the rehearsal environment cannot be completely identical to the production environment. Particularly, concerning importing policies into the rehearsal environment, some information (e.g., IDs in a policy database) is not preserved – the team has to make do with the same kind of objects, rather than replications of them: *“Yeah, you know, when you’re trying to import group policies and stuff like that for example the object names are slightly different than what you’ve got in production because the SIDs don’t match so it’s going to cause errors and stuff like that”* (IdM-lead).

4. COMPARISON WITH RELATED WORK

Since IdM is part of IT security, our results can be compared to prior findings in this area. Werlinger et al. [17] identified and classified challenges in IT security. Two of them – “access control” and “security culture” – were reasons for our case study insurance organization to adopt an IdM system. Our findings show that IdM related activities are distributed across the organization and require communication and collaboration between different stakeholders. This confirms prior research that shows that IT security is distributed across the organization [12, 6] and requires collaboration among different stakeholders [11, 18].

More specific to IdM, our findings provide an example of how previous research on identity and access control management [4, 19, 1] plays out in a particular case, which not only strengthens the previous findings, but also may highlight some nuances that were not previously emphasized.

Bauer et al. [4] studied challenges in access-control management in academic and non-academic organizations. The main focus of their study was on the ongoing management of accesses to file systems and physical environments. Their focus was limited to the process of access control management, while our research covers the whole process of identity management both without and with an integrated IdM system. Nevertheless, our findings about challenges in IdM process before deployment of the IdM solution corroborate their key findings. In particular, they identified a set of challenges that we also observed in our case study: management of exceptions, getting notification about policy change (em-

ployees leave the organization or change their department), getting updated information about who is responsible for a resource, verifying requests, keeping records, and choosing a usable access control management interface. As one of their implications for design, Bauer et al. propose allowing policy makers to directly edit the implemented policy in order to address the challenge that those who set the policies find it difficult to view and understand the implemented policy. Our findings reveal that this strategy may not be feasible, at least for organizations such as our case study organization that have complex system architectures and policy histories. If the policy implementers themselves (those who deploy the policies like SPs), with their domain knowledge of the systems are having difficulty, it is unlikely that the (primarily) non-technical policy makers will be up to the task of implementing it themselves.

The Identity Project [19] is a study of IdM practices in UK higher education institutions. The results of the Identity Project are based on a broad survey, which was validated and refined by 161 semi-structured interviews in the participating institutions. The Identity Project’s results are tuned to the improvement of business processes in the UK higher education sector. Being broad, the Identity Project findings can be complemented by case studies that aim at finer-grained information. While we confirm many of the results of the Identity Project, we also give a more detailed picture of some of the challenges, their cause, and the way the insurance organization coped with them. Moreover, comparison of our results with Identity Project results can show how the type of an organization can impact the challenges it faces in IdM and provide support for the generalizability of Identity Project’s challenges to non-academic organizations. For example, The Identity Project identified nine major challenges to IdM: limited consensus on defining “identity management,” heterogeneity in IT infrastructure, limited de-provisioning, lack of formal procedures, lack of both common standards and central IdM administration, lack of IdM data quality, use of non-unique user credentials, lack of policy for reuse of identifiers, and lack of adherence to a code of practice for information security. Our case study organization experienced only two of the nine challenges (limited de-provisioning and heterogeneity in IT infrastructure). We surmise that some of the other challenges may be more particular to academic organizations; Werlinger et al. [17] show that academic freedom is a barrier to enforcing policies.

The Burton Group commissioned several studies about various aspects of IdM, resulting in a body of reports for management that are summarized in a root document [1]. The Burton Group sources include their client organizations (e.g., six participating organizations in a survey and interviews about the implementation of roles), non-client organizations who use IdM (e.g., “discussions with approximately 20 medium-size to very large enterprises” concerning federation), presentations by vendors of IdM products, and discussion with consultants from other advisory organizations. The results of these studies are targeted more toward businesses who plan to decide about adopting an IdM system. Our results confirms Burton group findings about pre-requisites for IdM success and the challenges an organization might face adopting an IdM system. Our results also confirm business drivers for IdM adoption, including the need for security, the need to observe regulation, the need to reduce cost, establishing

new business models, and enhancing the user experience. In our case study, the primary driver for IdM was audits (SA-lead); that is, accountability. Concerning cost reduction, improvements are twofold. From the viewpoint of the security team, less effort is required for handling of data security by reducing labor-intensive, redundant activities, and increasing the quality of reporting. Furthermore, employee productiveness is increased as they wait less for access to resources. We identified a new business driver which is changing the role of security administrators to security analysts; this is expected to increase employee retention (SA-lead).

5. A VALIDATION OF GUIDELINES FOR IMPROVING IDM SYSTEMS

In our prior research [10], we developed a framework of design guidelines for IT security tools that classified guidelines generated by prior research according to the challenges the guidelines address. For example, the lowest layer in the framework comprises general usability guidelines for IT security tools. The next two layers contain guidelines that are necessary due to the work environment of security practitioners, which is characterized by technological and organizational complexity (including guidelines to address general communication challenges, guidelines applicable to tools used in a process that involves other stakeholders, and guidelines applicable to tools used by distributed SPs). The upper layer of the framework contains guidelines that are grouped based on task properties of the tool, such as those that require intensive configuration and deployment and those used in a process that requires intensive analysis. The framework is intended to aid researchers and developers in selecting guidelines for the security tool under consideration.

Since IdM systems are complex, involve multiple stakeholders, and involve extensive deployment and configuration, we expect guidelines that address technological complexity, diverse stakeholders, communication, distribution of IT security, and configuration and deployment to be meaningful with respect to IdM. Our case study confirms this by showing that a number of the guidelines from the framework came into play during the IdM deployment; in particular, the organization required the IdM system to be *integrated with the current infrastructure*, be *customizable*, *support workflow*, *enable flexible reporting*, *enable data to be presented in multiple formats*, *provide different methods of interaction*, *support archiving*, and *support rehearsal and planning*. We next briefly discuss how our observations shed light on how some of these more general guidelines can be applicable to the domain of IdM and where opportunities for future improvements in IdM technologies remain.

Provide integration with the current infrastructure: this guideline was a decisive requirement for our case study organization. They rejected one bidding vendor on the grounds that the vendor required them to replace their infrastructure. One of the reasons that the successful vendor was selected was because it was willing to adapt some features to the organization's way of doing things; that is, their IdM system was, in part, *customizable*, including customizing their documentation and providing customizable workflow and UI. Furthermore, the insurance organization developed much of its own software, and was able to integrate some of its own modules (such as role mining) with the system.

Provide workflow support: the insurance organization expected workflow support to reduce the workload of maintaining the data guardian framework by automating repeated manual tasks, facilitating division of responsibilities between stakeholders, and allowing effective communication and collaboration required for IdM tasks.

Afford flexible reporting and presentation of information: although the organizations' executive management was aware of security and privacy issues, nevertheless, the rationale for the implementation of the IdM system was couched in the language of business. The SA group wanted to generate reports that would show the effectiveness of the group to the executive. Likewise, they wanted to replace technical terms (like esoteric IDs) with more natural terms in reports to data guardians about who could have (or did have) access to their data. In the same vein, managers had to know what systems their employees should access and often did not know, largely because of the technical terminology.

Provide appropriate interaction methods for varying stakeholders: the insurance organization exhibited *different interaction methods* for different but related tasks. For example, managers requested access on behalf of their employees through an online, multi-page form, graphical interface (GUI), while the security administrators would deploy the requests to data guardians or network admins by e-mail. Additionally, in the RFIQ, the availability of command line interface (CLI) to allow communication with scripting tools was stated. Similarly, roles and access policies were perused differently depending on whether they were being edited by security analysts or explored by managers. Our case study shows that managers have trouble in understanding access policies. Using natural language or visualization techniques would be a viable solution to this problem.

Incorporative archiving capabilities: the insurance organization's determination to maintain accountability entailed archiving on a large scale, including through email and online forms for access requests. Our participants were not satisfied with the state of record keeping before deployment of the IdM system. In the RFIQ document, the organization stated the need for keeping time-stamped and tamper proof logs of user administration and workflow events. From a different perspective, the case study organization couldn't determine who accessed the data as they need to store logs from different systems. Log summarization could be a solution to this problem.

Enable rehearsal and planning: a main challenge for our case study organization for the IdM implementation was that the rehearsal environment could not be completely identical to the production environment. Also the IdM system was particularly difficult to update. Because of the critical nature of IT and IT security systems, planning and rehearsal are very important. The need to synchronize the rehearsal environment with the production environment may be seen as an opportunity for tool improvement.

6. LIMITATIONS

As a limitation of our work, the findings are based on the self reports of the participants. While these interviews are rich sources of information about the IdM adoption process, the data reflects the participants' interpretation of the process, which may not be accurate. We tried to alleviate this

limitation by analyzing documents related to organizational structure, the SA group structure and responsibilities, as well as the RFIQ document. Another possible approach to address this limitation would be to perform naturalistic observation of the technology adoption process; however, as yet, the insurance organization has not been willing to allow researchers to perform observation. Interview in together involved stakeholders like managers or end-users would also increase the accuracy of the findings.

Since our study was focused on one organization, our findings may not be generalizable. To address this problem, we compared our findings with related work to identify the portions of the work that strengthen or extend existing knowledge. Moreover, we tried to provide technical and organizational details about the context to facilitate the reader's *naturalistic generalization* (i.e., generalizations that are intuitive, empirical, and based on personal direct and vicarious experience [15]). Finally, our data does not cover the state of the organization after the complete deployment of the IdM system. Therefore, we can't argue whether the expected benefits are realized or not. A follow-up study is required to measure the benefits of the IdM system like service time, employee satisfaction, etc.

7. CONCLUSION AND FUTURE WORK

In this paper we described our case study of an insurance organization that deployed an integrated IdM system. We studied the organization at different stages of this project: before deployment, in the process of deployment, and after deployment of the first phase of their project. Our description of the insurance organization's stages of IdM deployment provides details about the organization's expectations and challenges. The state of the organization before IdM adoption may provide awareness in other organizations about possible costs of using legacy IdM techniques. Additionally, the experience of the case study organization and the expected benefits of IdM adoption may motivate other organizations to migrate to an integrated IdM system, while helping them know the challenges they may face. Finally, comparison of our findings with related work builds confidence that, not only are some of the findings generalizable, but they extend existing knowledge about IdM and validate that some of the recommendations for IT security tools extend to IdM.

For future work, we plan to continue interviews with various stakeholders during ongoing usage of the IdM system in the insurance organization. In particular, we are interested to learn whether the organization achieves the expected benefits. We are also performing more interviews in other organizations that have deployed an IdM system, which will help us to develop generalizable models of challenges, interactions, and recommendations for improvement of IdM solutions. Because role mining was a challenge for the insurance organization, we are interested in performing a detailed study on role mining in organizations, and identify the organizational and social barriers to role mining.

Acknowledgments

This work would not be possible without access to the HOT Admin and IdM project study participants. Members of LERSSE provided comments on preliminary versions of this paper. The HOT Admin project is supported by the NSERC

Strategic Partnership Program, grant STPGP 322192-05.

8. REFERENCES

- [1] *Root Document Enterprise Identity Management: Moving from Theory to Practice*. Technical report, Burton Group (June 2005).
- [2] Bauer, L., Cranor, L. F., Reeder, R. W., Reiter, M. K., and Vaniea, K. Real life challenges in access-control management. In *CHI*. ACM, New York, NY, USA, 2009, 899–908.
- [3] Blum, D. Identity management - concepts and definitions. Burton Group (September 2005).
- [4] Botta, D., Werlinger, R., Gagné, A., Beznosov, K., Iverson, L., Fels, S., and Fisher, B. Towards understanding IT security professionals and their tools. In *SOUPS*. Pittsburgh, PA, 2007, 100–111.
- [5] CA Corporation. How can a comprehensive identity and access management solution help me reduce security risk and achieve easier compliance? (2008).
- [6] Hawkey, K., Botta, D., Werlinger, R., Muldner, K., Gagne, A., and Beznosov, K. Human, Organizational, and Technological Factors of IT Security. In *CHI extended abstracts*. Florence, Italy, 2008, 3639–3644.
- [7] Heckle, R., Lutters, W. G., and Gurzick, D. Network authentication using single sign-on: the challenge of aligning mental models. In *CHIMIT*. ACM, Cambridge, Massachusetts, 2008.
- [8] Jaferian, P., Botta, D., Raja, F., Hawkey, K., and Beznosov, K. Guidelines for Designing IT Security Management Tools. In *CHIMIT*. ACM, 2008, 7:1–7:10.
- [9] Kandogan, E. and Haber, E. M. Security administration tools and practices. In L. F. Cranor and S. Garfinkel, eds., *Security and Usability: Designing Secure Systems that People Can Use*, O'Reilly Media, Inc., 2005. 357–378.
- [10] Knapp, K. J., Marshall, T. E., Rainer, R. K., and Ford, F. N. Managerial dimensions in information security: A theoretical model of organizational effectiveness. https://www.isc2.org/download/auburn_study2005.pdf (2005).
- [11] Microsoft Corporation. Microsoft identity and access management series: Fundamental concepts (2006).
- [12] Post, G. V. and Kagan, A. Evaluating information security tradeoffs: Restricting access can interfere with user tasks. *Computers & Security*, 26, 3 (2007), 229 – 237.
- [13] Stake, R. E. *The Art of Case Study Research*. Sage Publications Inc, 1995.
- [14] Vaidya, J., Atluri, V., and Guo, Q. The role mining problem: Finding a minimal descriptive set of roles. In *SACMAT*. ACM Press, Sophia Antipolis, France, 2007, 175–184.
- [15] Werlinger, R., Hawkey, K., and Beznosov, K. An integrated view of human, organizational, and technological challenges of IT security management. *Journal of Information Management & Computer Security*, 17(1) (2009), 4–19.
- [16] Werlinger, R., Hawkey, K., Botta, D., and Beznosov, K. Security practitioners in context: Their activities and interactions with other stakeholders within organizations. *International Journal of Human-Computer Studies*, 67, 7 (2009), 584–606.
- [17] Wright, J. Final progress reports. <http://www.jisc.ac.uk/media/documents/programmes/infrastructure/tidpfinalreport.pdf> (November 2007).