8-1-2011

Assured Identity for the Cloud

Jeff Daniels
Indiana State University

VITA

JEFF DANIELS

EDUCATION

Indiana State University 2011
Ph.D.: Technology Management, Digital Communications

Rensselaer Polytechnic Institute 1999
M.S.: Management of Technology

University of Central Florida 1997
B.S.B.A.: Major in Management Information Systems

TEACHING EXPERIENCE

Florida Southern College 2001-2005
CSC302 Database Design
CSC306 Systems Analysis & Design
CSC405 Database Programming

University of Central Florida 1999-2001
ISM 4113, Information Systems Analysis & Design
ISM 4212, Database Management Systems

RECENT PUBLICATIONS and PRESENTATIONS

Daniels, Jeff and Omar Channer. "Sustainable Systems Architecture." *Lockheed Martin Architecture Workshop Proceedings*. July, 2010. Stevens Institute of Technology. Hoboken, NJ.

Daniels, Jeff. "Server Virtualization Architecture and Implementation." *ACM Crossroads* 16, 1 (September 2009), 8-12. DOI=10.1145/1618588.1618592

EMPLOYMENT HISTORY

Lockheed Martin Corporation 1997-Present
Systems Engineer/Systems Architect

Florida Southern College 2001-2005
Adjunct Faculty

University of Central Florida 1999-2001
Adjunct Faculty

ASSURED IDENTITY

FOR THE CLOUD

A Dissertation

Presented to

The College of Graduate and Professional Studies

College of Technology

Indiana State University

Terre Haute, Indiana

In Partial Fulfillment

Of the Requirements for the Degree

Doctor of Philosophy

---

by

Jeff Daniels

August 2011

Keywords: Information Technology, Digital Communication, Cloud Computing, Technology Management

COMMITTEE MEMBERS

Committee Chair: Dr. David P. Beach

Professor, Electronics & Computer Engineering Technology

Indiana State University

Committee Member: Dr. Edward R. Kinley

Associate Vice President for Academic Affairs and Chief Information Officer

Faculty Affiliate, Electronics & Computer Engineering Technology

Indiana State University

Committee Member: Dr. William E. Croft

Professor, Electronics & Computer Engineering Technology

Indiana State University

Committee Member: Dr. Robert E. English

Associate Vice President of Academic Affairs

Professor, Electronics & Computer Engineering Technology

Indiana State University

Committee Member: Dr. Yuetong Lin

Assistant Professor, Electronics & Computer Engineering Technology

Indiana State University

ABSTRACT

It has been widely reported the largest security concerns with cloud computing design and implementation are centered on identity and access management. Pearson (2009) identifies open security challenges such as where processing takes place, auditability of transactions, and data sensitivity in distributed systems. Cloud computing builds on prior research in virtualization, distributed computing, utility computing, networking, and web services (Vouk, 2008).

A recent study conducted by the Office of Homeland Security found that cyber security is a national problem (Homeland Security, 2009). The study recommended that "managing identities" must be part of a comprehensive national cyber security strategy. The Department of Defense Cyber, Identity, and Information Assurance Strategic Plan calls for systems and security to be united.

In this research project, an approach to enabling assured identity and access management controls specifically in cloud computing environments was evaluated. The research designed and implemented the Assured Identity Management Systems (AIMS) using the systems engineering process (SEP). The evaluation of use cases and sequence diagrams demonstrated the capability for identity assurance with lifecycle events in cloud computing environments.

The dissertation study designed an e*xtensible* model including requirements, use cases, context diagrams, sequence diagrams, *reusable* components to further the adoption of cloud

computing, and a prototype built using interoperable cloud and virtualization technologies. The research supports the 2011 U.S. Federal Cloud Computing Strategy as well as the Standards Acceleration to Jumpstart Adoption of Cloud Computing (SAJACC) initiative. The dissertation research contributes to the body of knowledge in systems management, security, cloud computing and virtualization.

DEDICATION

I dedicate this research to my family for their contributions, commitments, stories, creativity, and ideas.

Mackenzie and Cambrie, this is for you:

"You're braver than you believe, and stronger than you seem, and smarter than you think. But the most important thing is even if we're apart; I'll always be with you."

- *Christopher Robin to Winnie the Pooh (Geurs, 1997)*

ACKNOWLEDGMENTS

In his book *Magnificent Desolation*, Apollo 11 astronaut Buzz Aldrin recalls a moment of peace and solitude on launch pad 34 before he boarded the Saturn V launch vehicle. Dr. Aldrin contemplated the historic journey ahead. He reflected on his life and how "all the facets and experiences had worked out along the way" to put him *in the right place, at the right time*.

The first quarter of 2011 was the toughest period of my life both physically and mentally. I was fighting for my life for two weeks in February; upon full recovery, in April, I successfully defended the research study. The message is one of peaks and valleys, staying the course in life, and finishing strong.

Like Aldrin, I spent time reflecting on all the opportunities, experiences, and people that helped lead me to this point. I would like to thank my parents, Howard and Dr. Debbie Daniels for encouraging me to reach for the stars, setting the bar high, and creating a passion for learning. Thanks to my friend and colleague Joel Johnson for supporting this study and advocating a healthy work-life balance. Thanks to my friend Cori Zuppo for being a sounding board, sharing information, and helping focus my efforts. I also want to thank my late grandfather E.P. Cochran and the late Rev. Earl Tyson, who always believed and encouraged me. Thanks to my committee members, Dr. English, Dr. Croft, Dr. Lin, and Dr. Kinley for their participation in the research and their valued individual and collective contributions to improve the study. Finally, I want to extend my sincere appreciation to Dr. David P. Beach for his mentorship and guidance through the years.

To my beautiful wife Jennifer and daughters Mackenzie and Cambrie, thank you for your sustaining support throughout this journey. You were always with me, even when my research took me far away.

Each of you helped me reach this threshold *in the right place, at the right time*.

Always believe, for with God all things are possible.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION AND CONTEXT OF PROBLEM

It has been widely reported the largest security concerns with cloud computing design and implementation are centered on identity and access management. The National Institute of Standards and Technology (NIST) released draft publication SP 800-125, Guide to Security for Full Virtualization Technologies in July 2010. The focus of the publication is security for virtualization and cloud computing environments. The purpose of the publication is to discuss the security concerns associated with full virtualization technologies and to provide recommendations for addressing these concerns.

SP 800-125 is reflective of the current state of security within the information technology domain, specifically with virtualization and cloud computing systems. Organizations are increasing initiatives in cloud computing driven by simplicity, affordability, and sustainability factors, but remain cautious with implementations as security risks are evaluated and analyzed. Pearson (2009) identifies open security challenges associated with cloud computing such as where processing takes place, auditability of transactions, and data sensitivity in these computing systems. In 1995, The European Union published Directive 95/46/EC which calls for the protection of personal data with respect to movement of data, transparency, and processing of data. Seven years later, the EU followed up with Directive 2002/58/EC on privacy and electronic communications, a continuation of the 1995 directive. The United States has also

enacted laws controlling the security and confidentiality of "Personally Identifiable Information" (NIST, 2008) described as

> *information which can be used to distinguish or trace an individual's identity, such as their name, social security number, biometric records, etc. alone, or when combined with other personal or identifying information which is linked or linkable to a specific individual, such as date and place of birth, mother's maiden name, etc. (p. C-1)*

The 1995 Health Insurance Portability and Accountability Act (HIPAA) focused on protecting patient health records and the Anti-Phishing Act of 2005 was centered on protecting personal data during a malicious phishing attempts. A number of legislative proposals have been proposed in Congress including consumer protection, wireless security, and social security number protection. The US and EU have partnered on the Safe Harbor Deal which seeks to prevent inadvertent data disclosure or loss. Cloud computing providers may opt in Safe Harbor, but are not required to.

In a recent interview Sun Microsystems executive Susan Landau (2009) recognizes the need for security in cloud computing environments given the distributed nature of these computing environments. Landau mentions the need for security control mechanisms for protection as well as the complexities associated with federated identity management within the cloud. Upon evaluation of cloud service providers, Salmon (2008) found that security is the responsibility of the customer. The customer and/or data owner must evaluate the physical and logical security controls deployed in the cloud environment. Depending on security requirements customer controls may or may not be adequate. Pearson identifies fluid changes in the systems life cycle whereby the security controls may change, or the customer security requirements may change, thus prompting constant security evaluation and re-evaluation.

A June 2010 survey of 308 technology professionals conducted by a Gartner Group research team found that identity management was the top security priority (Messmer, 2010). The survey found other security priorities include data-loss prevention and intrusion prevention. Another Gartner Group survey focused on global CIO strategic directions found that most CIOs were interested in virtualization technologies followed by cloud computing.

Software maker Microsoft is combining cloud and identity management services in their latest offering. Active Directory Federated Services 2.0 increases interoperability between private, public and hybrid clouds according to JG Chirapurath, Microsoft's director in the identity and security business group (Messmer, 2010). Chirapurath states that "identity" is the glue that will make cloud computing work.

Despite increased awareness of identity and access controls within cloud computing environments, there remains significant questions. There may be a dilution of control for patching, maintenance, and general operations of a cloud system. Responsibility among customer, end-user, and cloud services provider for the maintenance and security of applications configuration, data validity, and data integrity may be unclear given the expanded nature of the cloud. Data breaches and public exposure of data loss, compromises, and intrusions are valid threats in the cloud.

This study involved identity assurance for cloud computing systems. Comprehension of how user credentials are provisioned through the lifecycle process is vital to understanding *who* is accessing data in public and private clouds. The research addressed interoperability among cloud computing systems, specifically where common credentialing methods used in legacy systems are compatible with newly deployed cloud systems.

Problem Statement

Cloud computing systems have security attributes that must be considered; no longer are systems stand-alone and deployed solely behind corporate firewalls. Cloud systems present security interdependencies as a result of collaborative benefits which are critical to business strategy. Systems designers and systems administrators must address identity assurance in terms of provisioning, access control, authorization, and non-repudiation. The problem this research addressed is that identity management technologies are not mature and cause security concerns among cloud computing adopters.

Definition of the Theoretical Terms Used in the Research

Virtualization and cloud computing technologies have established themselves in the enterprise architecture environment. Hardware vendors are packaging systems tuned to support virtualization. Software vendors are developing virtual server tools for migrations, performance, and high-availability. Customer IT organizations have defined a virtualization strategy and have begun deploying *virtualized* data centers and integrating cloud computing environments.

The virtual machine concept has been around for many years. The recent revolution in virtualization technology, hypervisors, and paravirtualization has allowed servers using the popular x86 architecture to operate efficiently and effectively with virtual machines. Virtual machine technology is an enabler for service oriented architectures, isolated secure systems, and flexible deployment.

An architectural style is a "named collection of design decisions applied in a particular development context and intended to elicit known beneficial qualities" (Dashofy, 2007). The architectural style helps to validate design decisions and demonstrate system-wide adherence to guiding architectural principles in part by implementing visual methods. The architectural style

can show component attributes and connectors or adapters used to facilitate communications among system interfaces.

Cloud computing is a best described as a collection of computing resources and platform services. Characteristics of cloud computing systems are scalability of system resources on demand, minimization of start-up costs and infrastructure, and improving business process flows by using common compute platforms. Some of the challenges with cloud computing include security for partitioned environments, interoperability among data and applications, applications portability into and out of the cloud and governance of the cloud computing environment (Cloud Computing Manifesto, 2009). Cloud computing is the next generation of virtualized infrastructure and application services. The National Institute of Standards and Technology Information Technology organization provides a framework for cloud computing (see Figure 1). NIST depicts deployment and service models with essential characteristics to highlight distinctive features of the environment.

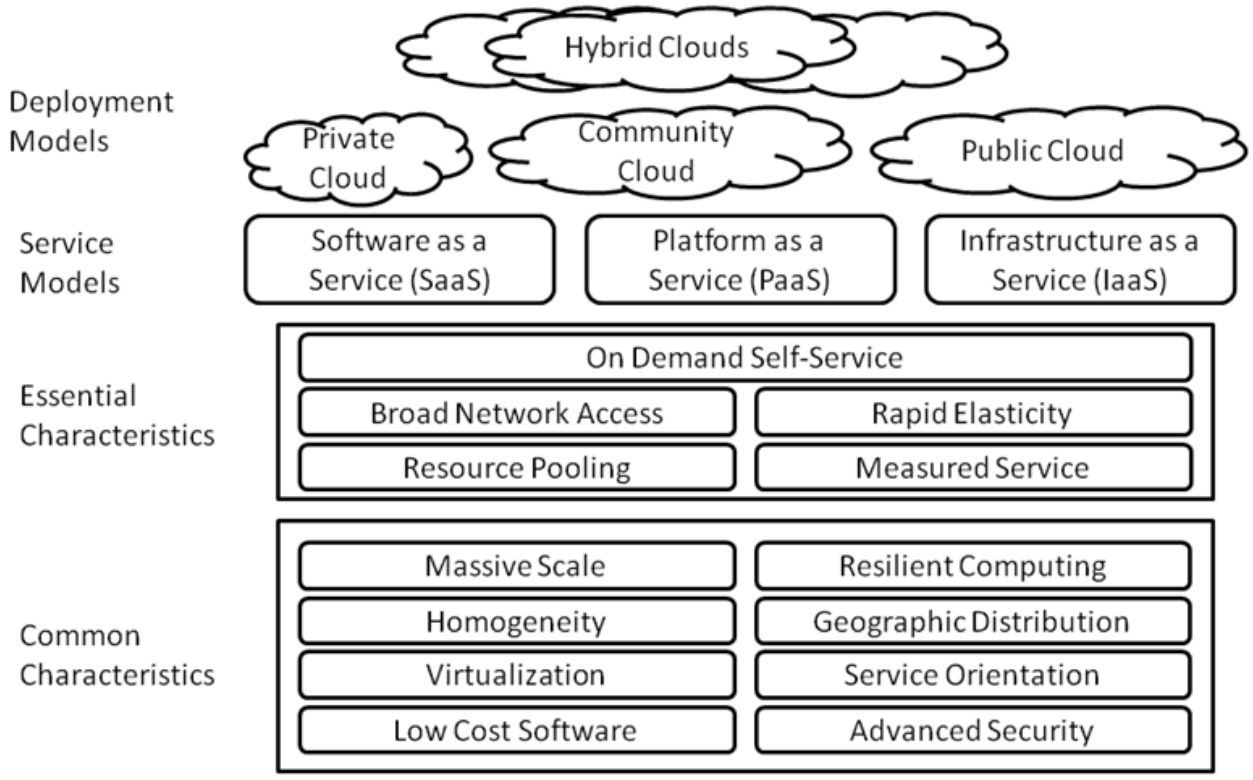

Figure 1.

NIST Cloud Definition Framework (NIST, 2009)

Research Goals and Objectives

The goal of the research project was to *study systems that enable assured identity specifically in cloud computing environments.*

The research project was guided by the following specific objectives:

- Build prototype environment to validate assured identity in cloud computing
- Demonstrate capability for organizations to exchange secure data from assured identities through cloud computing.
- Create repeatable, extensible use cases for identity assurance and cloud computing

- Provide artifacts that support the adoption of cloud computing

Trusted computing within the cloud is dependent on identity assurance. To provide assurance of credentialed identities, trusted secure services must provision attributes and entitlements. The study provided integration with cloud-based provisioning systems by demonstrating validation of adequate privileges and rights within the cloud environment. The system also provided authentication and authorization service management based on privileges, authorized actions, and entitlements.

Research Ethics

The research principal feels it is important that research ethics are adhered to and all possible steps be used to avoid potential conflicts of interest. With conflict, the success and authenticity of the research project is potentially jeopardized. There is an obligation on behalf of the research principal and research team to understand all organizational policy, ethics, and related programs to ensure his ethics and standards meet certain criteria.

This research project adhered to the following ethics principles:

Do What's Right

The research team is committed to the highest standards of ethical conduct in all that we do. We believe that honesty and integrity engender trust. We abide by the laws of the United States and other countries in which we do research. We strive to be good citizens and take individual responsibility for our actions.

Respect Others

We believe that respect - for colleagues, customers, partners, and all those with whom we interact - is an essential element of all positive and productive relationships.

Perform With Excellence

We understand the importance of the research project and the trust placed in us to perform. With this in mind, we strive to excel in every aspect of research and approach every challenge with the determination to succeed.

Ethical behavior is vital for the protection academic, corporate, and national defense research. Classified and unclassified programs much operate under ethical principles. To behave otherwise is costly, wasteful, and perhaps catastrophic to the interests of a free society and representative form of government.

Scope

In this research project, the scope encompassed systems architecture modeling, process modeling, systems engineering, virtual machines, cloud computing, and utility computing concepts.

The research focused on conceptual enterprise class systems commonly found in the Department of Defense, universities and research institutions, and the federal government. These systems included cloud computing technologies within the architecture.

Justification for research

A 2008 survey by *CIO Research* suggested that while "58 percent say cloud computing will cause a radical shift in IT" (McLaughlin, 2008), nearly half stated that there are security problems as the greatest concern with cloud computing. Gardner (2010) found that cloud infrastructures are not secure enough for mission critical applications according to users. Recent studies have focused on security measures such as authentication and federated identities for virtual organizations (Gemmill, 2006), grid computing security models (Kanaskar, Topaloglu, Bayrak, 2005), and access control systems for grid infrastructures (Oo and Naing, 2007). Lombardi and Di Pietro (2010) discuss integrity protection with a "transparent layer of security"

at the virtualization monitor level. Christodorescu, et al. (2009) proposes a specialized virtual machine to provide infrastructure security.

There are various other possibilities where cloud-based technologies could create opportunities or cost savings. As business goals are defined and objectives determined by the business, virtualization technologies should be considered as one of the ways IT can help meet those goals. Cloud computing is a key component in several planning areas including:

- expanding business lines such as shared and dedicated hosting
- faster deployment, time to market
- increased standardization, leading to lower TCO (total cost of ownership)
- consolidation efforts
- increased utilization of computing capital

Enterprise architecture is *"the organizing logic for business process and IT infrastructure capabilities reflecting the integration and standardization requirements of the firm's operating mode"* (Ross, 2007). Enterprise architectures seek to align business goals and organizational needs with technology. The idea is to plan, deploy, and manage technologies to meet business objectives. Similar to the IT strategic plan, virtualization technologies have their place in the enterprise architecture model.

Ross mentions two important concepts in her definition of enterprise architecture: integration and *standardization*. Cloud computing offers increasingly flexible methods of systems integration. Hot failovers, highly available systems, real-time relocation of virtual systems, dynamic reallocation of system resources, and even wide-area network disaster recovery (backup) are features of the virtualized cloud computing environment. The expanded data-center concept may integrate physical and virtual servers with public and private cloud systems that perform functions such as routing, messaging, and directory services.

Cloud systems go a long way towards standardization for infrastructure operations. Servers can be commoditized using the "gold image" model where a virtual machine with the latest compliant system configuration is used to build new servers, ensuring standardization and change control. This also reduces risk of mis-configuration or non-configuration of features that may occur due to human error when building and rebuilding physical systems. Common platforms serve as an enabler for business objectives and other enterprise architecture components. Initiatives such as ERP implementations and service oriented architecture applications rely on infrastructure being available, standardized, and usable. Cloud technologies can be used as a building block in the standardization and integration in enterprise architectures.

Significance of the Study

The research project contributes to the body of knowledge in systems management, systems architecture, virtualization and cloud computing. The collaboration between assured identity and cloud systems is relatively new, hence the need to demonstrate and design security systems and manage user identities and credentials throughout the cloud environment. Cloud computing systems are gaining momentum in the enterprise, federal government, and corporate environments.

Research Questions

Using a research question format, the general research question of this study is:

How can identity and access management controls be designed to support cloud computing systems?

As a result, the following questions are addressed in this research project:

1. *How do we determine who is authorized to be on the cloud?*
2. *What mechanisms exist to provide the identity management and access function?*

3. *How is assured identity provided in cloud computing environments?*
4. *How do we interoperate with different identity and access mechanisms in a global enterprise?*

Statement of Assumptions

In conducting the research study, a number of assumptions were made. The design and implementation of the Assured Identity Management System (AIMS) were heavily influenced by security considerations with one distinction: the underlying cloud layers Infrastructure-as-a-Service (IaaS) and Software-as-a-Service (SaaS) were not explicitly trusted. These layers may be hosted by cloud service providers and have specific trust and security mechanisms. The focus of the AIMS system is Identity as a Service (IdaaS) and considered a component of the Platform-as-a-Service (PaaS) deployment model. To that end, the tools used in study are considered to be representative of cloud-based platform service technologies. The final assumption is that all cloud-based integrated systems within the study supported RESTful (Representational State Transfer) protocols such as HTTP (hypertext transfer protocol), SOAP (Simple Object Access Protocol) and APIs (application programming interfaces) such as Sun's Cloud API.

Statement of Limitations

The study was conducted with several known limitations that must be acknowledged. Within the cloud framework, some functions such as administrative privileges were prohibited; the accounts and rights granted were a general user category. Given the increased vetting process and complexity of provisioning administrative privileges at the infrastructure level, these were omitted from the study. The cloud deployment tools are based primarily on open source technologies where possible with a combination of commercial off the shelf (COTS) products as required. The use cases demonstrated in the study are representative of industry scenarios, but

may not encompass all possibilities. One of the drivers for this research study was to define use cases and scenarios to further the use of cloud computing due to the fact that standards are still maturing. The use cases are reflective of proposed cloud computing interoperability, with recognition that standards change and cloud computing business uses will evolve. The intent of these use cases is to capture the steps and associated functionality for identity assurance in the cloud. The use cases may be extended and/or tailored to fit customized deployments or industry specific cloud-based designs.

Statement of Methodology

This research studied the assured identity management controls in one specific type of deployment model: Platform-as-a-Service (PaaS). The term *platform* abstracts the lower level infrastructure services such as network attached storage, virtual machines, and network services. To deliver identity services, the application was built on a cloud computing fabric. The fabric of services is made available over hypertext transfer protocol (HTTP) through simple representational state transfer (REST) application programming interfaces. The cloud-based system called AIMS, assured identity management system, was configured to be the platform that provides identity lifecycle management services in a distributed computing model. The AIMS application relied on two functional concepts: service bus and access control.

The service bus concept allowed services from remote or disparate networks to integrate. Rather than manage specific port and connectivity information, the service bus exposed an endpoint in the cloud that consuming applications may call. Consuming applications used metadata attributes brokered by the service bus to provide identity assurance. The second concept was access control which allowed *federated* authentication where claims-based requests were fulfilled up proper identification from third party identity providers.

The Systems Engineering Process (SEP) as defined by the International Council on Systems Engineers was employed for this study. Within the SEP, the systems life cycle model has six stages including concept, development, production, utilization, support, and retirement (INCOSE, 2004). The AIMS system followed the SEP from concept through support for purposes of the study.

INCOSE defines systems engineering as an interdisciplinary approach to enable successful systems. Ramo describes systems engineering is a discipline that concentrates on the design and application of the whole (system) as distinct from the parts. It involves looking at a problem in its entirety, taking into account all the facets and all the variables and relating the social to the technical aspect. (1998). Systems engineering has emerged a way to design and manage complex systems.

A requirements specification for the AIMS system including functionality, security, and interface requirements was generated based on industry working group focus, market analysis, and literature review. The requirements framework served as the basis for a series of use cases that validated the functionality of AIMS and the integrated of assured identity credentials.

Summary of Contributions

The study evaluated potential cloud architectures for *assured identity and access control* using a representative benchmark of security attributes from the security standards community. The study implemented these security measures in an integrated cloud-based system called AIMS, the assured identity management system and test it on several candidate architectures involving client, server, and cloud resources. The study contributes to the body of knowledge in the following specific areas:

- The study presents a candidate cloud configuration for assured identity and access control;
- The study describes a comprehensive set of integrated, tested, and repeatable use cases based on AIMS, the Assured Identity Management System.

CHAPTER 2

REVIEW OF LITERATURE

What is Cloud Computing

Cloud computing is a term within coined within the past several years. The mere mention of the term may cause confusion or disagreement on the scope of cloud computing. Industry specialists and academics may differ in their definition of cloud computing, however most agree on the conceptual framework that cloud computing consists of. "'Cloud' computing builds on decades of research in virtualization, distributed computing, utility computing, and, more recently, networking, web and software services" (Vouk, 2008).

Peter Marks at Google says "The idea of cloud computing comes from the early days of the Internet, where we drew the network as a cloud. We didn't care where the messages went—they came in one side and out the other, and we didn't have to worry about the network [because] the cloud hid it from us. [It's] a 'cloud' around [network] buckets" (Farber, 2008). Tim O'Reilly, founder of O'Reilly media, best known for its technical manuals, observes the cloud is the foundation of Web2.0 technology, and offers his thoughts on cloud computing, ""Cloud computing is a network of networks. [It's] a great way to think about how we will be delivering computing systems in the future." (Slack, 2009). Microsoft's Ray Ozzie views cloud computing as "a personal mesh of devices – a means by which all of your devices are brought together, managed through the web as a seamless whole" (Economist, 2008).

Researchers Peter Mell and Tim Grance of the National Institute of Standards and Technology Information Technology Laboratory define cloud computing with the following (Mell and Grance, 2009):

> *Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models. (p. 1)*

The NIST definition is commonly referred to, however, Mell does admit, "everybody has confusion on this topic" (Talbot, 2010). NIST has produced over 15 revisions to the standard document that defines cloud computing. In the latest draft, a disclaimer states cloud computing is still an evolving paradigm (Mell and Grance, 2010) and that spirited debate on the topic will evolve and change the definition of cloud computing over time.

Perry (2008) describes cloud computing as a style of computing where services are delivered on a massive scale through the Internet technologies. The Introduction to Cloud Computing Architecture (Sun, 2009), Sun Microsystems has this perspective on cloud computing:

> *…it's using information technology as a service over the network. We define it as services that are encapsulated, have an API, and are available over the network. This definition encompasses using both compute and storage resources as services. Cloud computing is based on the principle of efficiency above all —*

> *efficiency that produces high-level tools for handling 80% of use cases so that applications can be created and deployed at an astonishing rate. (p. 1)*

Clearly all three definitions include services as the core of cloud computing.

Pachner (2010) defines the cloud this way: "In a nutshell, cloud computing means tapping software, hardware or storage over the Internet, then using and paying for it on an as-needed basis." She uses Facebook and Google's Gmail service as examples of cloud-based applications.

From an educational perspective Siegel (2010) describes cloud computing as "a computing technology that uses the Internet and central remote servers to maintain data and applications." Students and teachers would be able to use applications without installing them on their local computers and save files *from* any connected computer. A technology research team from Purdue University recently created a classroom companion tool called Mixable that resides on top of Facebook (Kolowich, 2010). The tool takes advantage of Facebook's cloud presence and creates a collaborative environment for students to interact, discuss, and share educational materials.

Vouk (2008) identifies cloud computing by implying the cloud characteristics, "a service oriented architecture, reduced information technology overhead for the end-user, great flexibility, reduced total cost of ownership, on-demand services, and many other things."

Erik Brynjolfsson, Paul Hofmann, and John Jordan (2010) define cloud computing from two perspectives. The practitioner point of view is "cloud computing is on-demand access to virtualized IT resources that are housed outside of our own data center, shared by others, simple to use, paid for via subscription, and accessed over the Web." Brynjolfsson, et al. look to Armbrust's (2010) definition of cloud computing for the academic perspective: "Cloud

computing refers to both the applications delivered as services over the Internet and the hardware and systems software in the data centers that provide those services. … The data center hardware and software is what we will call a cloud. When a cloud is made available in a pay-as-you-go manner to the public, we call it a public cloud; the service being sold is utility computing."

In her 2009 report "Cloud Computing: Not So Cloudy Anymore," research analyst Jennyfer Valez uses the definition: "flexible and scalable shared environment in which third-party suppliers use virtualization technologies to create and distribute computing resources to customers on an as-needed basis, through the Internet browser."

Cloud computing have varying definitions that will continue to change as the technologies mature and adoption of cloud computing increases. The definitions indicate a strong emphasis on flexible, scalable computing services. The interconnectedness of compute resources in the cloud is evident; however none of the definitions mention security, encryption, or access.

History of Virtualization and Cloud Computing

J.C.R. Licklider was a distinguished engineer and visionary computer scientist, often referred to as the "Johnny Appleseed" of computing for his contributions by planting the seeds for the Internet and World Wide Web. In 1963, Licklider was appointed head of the Behavioral Sciences and Command and Control programs at the Department of Defense Advanced Research Projects Agency (ARPA). "Lick" as his colleagues affectionately referred to him, addressed members of the Intergalactic Computer Group later that year with a memo calling for a network of computers that would allow scientists to collaborate irrespective of distance or computer compatibility issues. Lick referred to the system as an intergalactic computing network, describing it conceptually, "….we could have at least four large computers, perhaps six or eight

small computers, and a great assortment of disc files and magnetic tape units-not to mention remote consoles and teletype stations-all churning away" (Waldrop, 2000).

Lick's vision for connected computing served as the foundational concept for the ARPANET, which led to the creation of a series of military networks including MILNET, SIPRNET, and eventually the Internet and later World Wide Web. In Segaller's 1998 book "Nerds: A Brief History of the Internet", Roberts describes Lick and his vision:

> *"Lick had this concept of the intergalactic network which he believed was everybody could use computers anywhere and get at data anywhere in the world. He didn't envision the number of computers we have today by any means, but he had the same concept-all of the stuff linked together throughout the world, that you can use a remote computer, get data from a remote computer, or use lots of computers in your job. The vision was really Lick's originally. None of us can really claim to have seen that before him nor {can} anybody in the world. Lick saw this vision in the early sixties. He didn't have a clue how to build it. He didn't have any idea how to make this happen. But he knew it was important, so he sat down with me and really convinced me that it was important and convinced me into making it happen" (Roberts in Segaller, p. 40).*

J.C.R. Licklider died in 1990 having worked on core components of UNIX development, network computing, time sharing operations (Project MAC), and professor emeritus at the Massachusetts Institute of Technology. His vision for intergalactic computing continues to inspire new platforms and tools for collaboration and information sharing.

The Virtual Machine

Virtual machines have been in the computing community for over 40 years. Early in the 1960's systems engineers and programmers at MIT recognized the need for virtual machines. In her authoritative discourse, "VM and the VM Community: Past, Present, and Future", Melinda Varian (1997) introduces virtual machine technology starting with the Compatible Time-Sharing System (CTSS). IBM engineers had worked with MIT programmers to develop a time sharing system to allow project teams to use part of the mainframe computers. Varian (1997) goes on to describe the creation, development and use of virtual machines on the IBM OS/360 Model 67 to the VM/370 and the OS/390. Varian's paper covers virtual machine history, emerging virtual machine designs, important milestones and meetings, and influential engineers in the virtual computing community.

In 1973, Srodowa and Bates demonstrated how to create virtual machines on IBM OS/360s. In "An Efficient Virtual Machine Implementation", they describe the use of IBM's Virtual Machine Monitor, a hypervisor, to build virtual machines and allocate memory, storage, and I/O effectively. Srodowa and Bates touch on virtual machine topics still debated today: performance degradation, capacity, CPU allocation, and storage security.

Goldberg (1973) concludes "the majority of today's computer systems do not and cannot support virtual machines. The few virtual machine systems currently operational, e.g. CP-67, utilize awkward and inadequate techniques because of unsuitable architectures." Goldberg proposes the "Hardware Virtualizer", in which a virtual machine would communicate directly with hardware instead of going through the host software. Nearly 30 years later, industry analysts are excited about the announcement of hardware architectures capable of support virtual

machines efficiently. AMD and Intel have revealed plans to introduce Pacifica and Vanderpool chip technologies in 2006.

The 1980's and early 1990's brought distributing computing to data centers. Centralized computing, and virtual machine interest was replaced by standalone servers with dedicated functions: email, database, web, applications. After significant investments in distributed architectures, renewed focus on virtual machines as a complimentary solution for server consolidation projects and data center management initiatives has resurfaced.

Recent developments in virtual machines on the Windows x86 platform merit a new chapter in virtual machine history. Virtual machine software from Virtuozzo, Microsoft, Xen, and EMC (VMWare) has spurred creative virtual machine solutions. Grid computing, computing on demand, and utility computing, technologies seek to maximize computing power in an efficient, manageable way.

The virtual machine was created on the mainframe and only recently has been introduced on the mid-range, distributed, x86 platform. Technological advancements in hardware and software make virtual machines stable, affordable, and offer tremendous value given the right implementation.

Types of Virtualization

Virtual machines are implemented in various forms. Mainframe, open source, paravirtualization, and custom approaches to virtual machine have been designed over the years. Complexity in chip technology and approaches to solving x86 limitations of virtualization have led to three different variants of virtual machines. Figure 2 shows three virtualization approaches: software virtual machines, hardware virtual machines, and virtual operation system or containers.

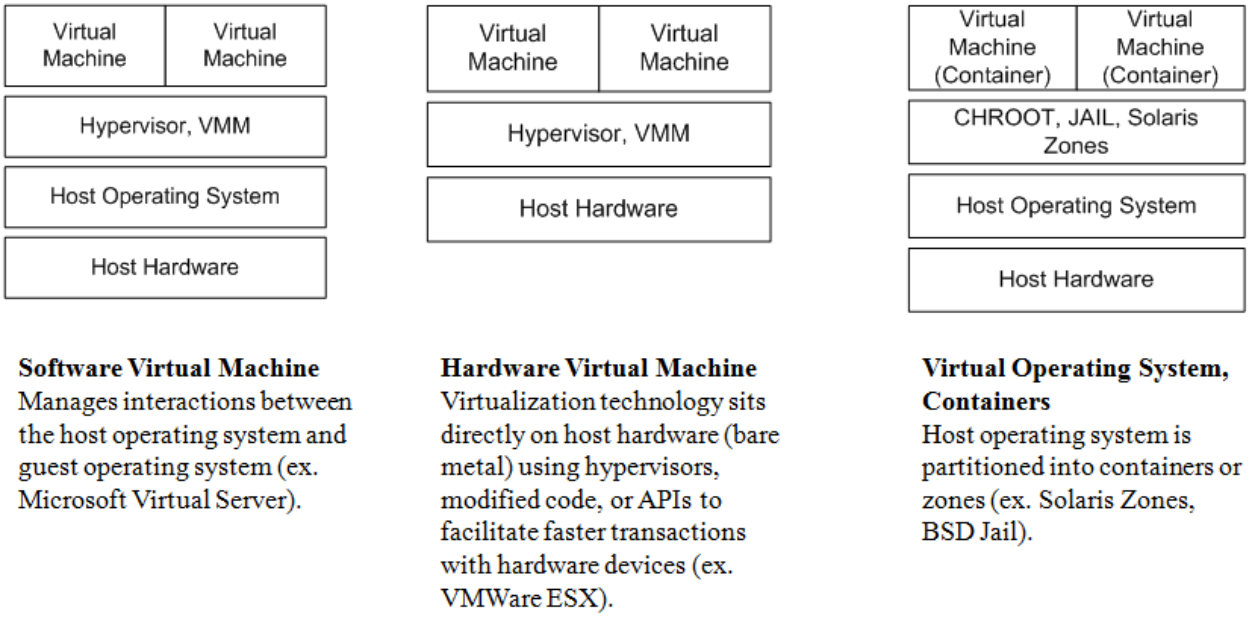

Figure 2.

Types of Virtualization

Software virtual machines create a management layer that emulates a guest operating system that resides on a host operating system. A distinction from hardware virtual machines is the host operating systems requirement which may vary by virtual machine software. This implementation offers flexibility to reside on an existing host operating system. Popular tools in this domain are VMWare Desktop, VirtualBox, VirtualPC, Parallels, and Xen.

The hardware virtual machine is sometimes referred to as "bare metal" given that host operating systems are not required. This virtualization implementation has a modified hypervisor kernel that allows direct communication with the host hardware (rather than through a host operating systems layer). One benefit of this model over software virtual machines is higher performance.

A simple UNIX implementation called *chroot* allows an alternate directory path for the root file system may be created in the virtual operating systems or container model. This creates

a “jail”, or sandbox for new applications or unknown applications. Isolated processes in chroot are best suited for testing and applications prototyping. They have direct access to physical devices, unlike emulators.

SUN Microsystems’ “Solaris Zones” technology is an implementation of chroot, similar to the FreeBSD jail design, with additional features. Zones allows multiple applications to run in isolated partitions on a single operating system (Tucker, Comay, 2004). Each zone has its own unique process table and management tools that allow each partition to be patched, rebooted, upgraded, and configured separately. Distinct root privileges and file systems are assigned to each zone.

Symmetric multiprocessing or SMP was introduced on RISC platforms such as SUN Sparc and DEC Alpha chipsets before being adopted on the x86 Intel Xeon and AMD Athlon processors. SMP allows multiple, identical chipsets to share one memory bank. Instructions can be shared among the processors or isolated to a dedicated processor on the system. The system can share a workload and with increased efficiency. A variation of SMP is AMD’s Opteron technology which allows dual-processors per chip. The Opteron uses DDR SDRAM memory dedicated to each processor as opposed to a single shared memory bank. The multiprocessing nature of numerous virtual machine guest servers on one host makes dual-core Opteron chips an attractive platform.

Paravirtualization is a variant of full-operating system virtualization. Paravirtualization avoids “drawbacks of full virtualization by presenting a virtual machine abstraction that is similar but not identical to the underlying hardware” (Whitaker, 2002). This technique allows a guest operating system to be “ported” thru a special API (application programming interface) to run. A paravirtualization research project called Xen at the University of Cambridge is a virtual

machine monitor (hypervisor) that allows commodity operating systems to be consolidated and effectively mobilizes guests across physical devices. Xen currently supports only open source guest systems; a Windows XP port is being developed. Denali is another paravirtualization implementation but requires significant modification to host system binaries and focuses on high performance virtual machines.

EMC's VMWare technology is the market leader in x86 virtualization technology. VMWare ESX server uses a special hypervisor to "dynamically rewrite portions of the hosted machine code to insert traps wherever VMM intervention might be required" (Barham, et al., 2003). The VMWare solution is more costly, but provides a robust management console and full-virtualization support for an array of guest operating systems including Solaris, Linux, Windows, and DOS.

On the heels of virtualization, cloud computing arose as a conceptual framework of services. The concept builds on years of previous work in information technology and computer science. The various types of cloud computing use virtualization technologies at different levels such as application containers, dynamic virtual machines, and hypervisors. It is important to understand the virtualization technologies behind cloud computing that help provision and deprovision elastic services such as storage, CPU, memory and other system resources. Access to these resources requires identity assurance to securely manage the cloud computing virtual systems.

Drivers for Cloud Computing

"The cloud will do for government what the Internet did in the '90s," Federal Government Chief Information Officer Vivek Kundra says. Other technologists agree, cloud computing represents a fundamental change to the way our government and organizations operate. Rather

than owning the infrastructure, there are potential cost savings and cost avoidance opportunities with the cloud (Nagesh, 2008). While these reasons are attractive and even exciting, many organizations are approaching cloud computing with caution given security concerns. Regardless, a number of drivers are pushing cloud computing to the forefront of the information technology field.

Affordability

Affordability is emerging as a driving factor for cloud computing. The state of geopolitical economics and organization budgets, particularly within information technology organizations, has influenced organizations to take an affordability challenge.

The 2010 National Defense Authorization Act Section 805 calls for all Department of Defense technology programs to "maximize value to the Department of Defense by providing the best possible product support outcomes at the lowest operations and support cost." The Weapon Systems Acquisition Reform Act (WSARA) 2009 established a new cost assessment and program evaluation office, put a stronger emphasis on cost and cost estimates, and recommended to "make life cycle affordability a cost business process for all communities and stakeholders involved in system acquisition and sustainment."

According to Reza Malekzadeh, senior director of products at EMEA, the majority of an IT budget is allocated to 'maintenance' activities, and very little on innovation projects (Hayes, 2009).

A 2008 Department of Defense study on technology acquisition and the systems life cycle revealed that just over a quarter of the cost of a system is encompassed by the initial stages of the life cycle. The stages of the systems life cycle including planning, design, development, and deployment accounted for 28% of the total cost of the system. Maintenance accounted for

approximately 72% of the total system cost. These findings, independent of cloud computing, support Malezadeh's statements on the majority of the budget taken up by system maintenance.

Independent of the DoD findings, a recent IBM study estimates that 70% of the cost of a system is spent on maintaining current information technology infrastructures. The study found that 70¢ per $1 is spent on maintenance versus adding new capabilities. Vouk (2008) indicates that unless IT is the primary business of an organization, less than 20% of its efforts not directly connected to its primary business should have to do with IT overhead, even though 80% of its business might be conducted using electronic means. Supporting existing systems is a costly endeavor and organizations are challenged to reduce the heavy footprint of maintenance costs.

Intel, for example, has a strategic plan to consolidate its data centers from more than 100 eventually down to about 10 facilities. In 2008 the total fell to 75, with cost savings of $95 million. According to Intel's co-CIO Diane Bryant, 85% of Intel's servers support engineering computation, and those servers run at 90% utilization (Brynjolfsson, et al., 2010). Contrast Intel's 10% idle computing average with an IBM study that revealed most servers in distributed environments sit 85% idle. Intel is maximizing the investment in computing resources by allocating more processing to their cloud systems, which represents cost savings or cost avoidance from purchasing additional (less utilized) computing systems. This decision also frees up funds to invest in innovation projects or research programs.

Innovation

Innovation might be the hidden strength of cloud computing according to researchers and analysts. Through efficient use of computing systems, availability of affordable computing resources, and quicker time to market, the real strength of cloud computing is as a catalyst for innovation (Brynjolfsson, et al., 2010). Cloud computing offers ubiquitous, cheaper access to

systems which in turn free up resources to focus on innovation projects and other strategic business objectives.

The potential of the cloud to be a catalyst for innovation is echoed by Harvard economist and information technology specialist Dale Jorgenson, "Many of these [software] applications are going on at a blistering pace, and cloud computing is going to be a great facilitative technology for a lot of these people" (Talbot, 2010). Convenient computing resources made available to try out new application functionality helps cloud users focus on their core role as a software developer or systems analyst. These cloud users are no longer required to support capital approval processes for new computing systems or walk request through procurement cycles, they can focus on analysis and design and not worry about the availability of computing resources when needed to test or demonstrate new features.

Cloud computing saves time and helps meet organization affordability targets. Schools are also finding cloud computing to be educationally beneficial. Students can learn a variety of computer production skills in a collaborative learning environment. Projects started at school can be continued at home without having to transfer files or download compatible software (Siegel, 2010).

To summarize, organizations spend more on maintenance and keeping systems running than they do on core business functions, new projects, and innovation. Cloud computing is a catalyst for innovation by freeing up individuals to focus on their core job function rather than non-core functions such as locating system resources, procuring new hardware, etc.

Efficiency

Cloud computing offers an efficiency factor that acts as a driver to adopt cloud computing systems. A recent survey (Wittmann, 2010) found that the main reason organizations

are moving to clouding computing is not cost savings, but the ability to get a system or application up and running and in the market quickly. Cloud computing allows new ideas, business concepts, and prototype systems to rapidly move to market, which represents a competitive advantage.

Cloud computing offers efficiencies in terms of operations and maintenance support. Using cloud and virtualization technologies, a systems administrator can maintain 1,000 servers in a large data center rather than the average 140 servers in a medium-sized data center (Hamilton, 2008). Cloud systems can be managed remotely and in many cases, the end user is able to determine the systems configuration by editing the level of resources allocated to a system hosted in the cloud.

Data Robotics' CEO Dr Geoff Barrall feels that data backup is the next big step for cloud computing. "In 1990 a high-end technical consumer would have had about 100MB of data on average, and a high-speed data link would have been a 28.8k modem," Barrall explains by way of example. "Now home bandwidth is up to (say) Mbps – approximately 174× growth. Actually data however has grown to about 1.3TB for the same user – a 13,000× data growth that massively outpaces the growth in bandwidth." (Hayes, 2010).

There is a convenience factor with cloud computing. Knowledge workers and users spend hours working from their web browsers. There is a comfort level with the interface, the browser application and how it behaves. A 2008 survey found 51% of Internet users who have done a cloud computing activity say a major reason they do this is that it is easy and convenient (Pew, 2008). In the same survey, another 41% said they enjoy the ability to access to data from any computer workstation.

Characteristics of Cloud Computing

Cloud computing has several key characteristics that differentiate it from other forms of computing such as standalone computing, mobile computing, grid computing, and other computing methods. The cloud may be combined or used in conjunction with other types of computing environments such as mobile computing, whereby a mobile device would access cloud-based data.

Werner Vogels, Chief Technology Officer for Amazon is a pioneer and early adopter of cloud computing. Amazon's cloud computing services have been available since 2006, making it one of the early market leaders. Amazon's S3 storage service stores roughly 18 billion objects in the cloud. Vogels identifies key characteristics to enable cloud computing:

- *Security*
- *Scalability*
- *Availability*
- *Performance*
- *Cost-effective*
- *Acquire resources on demand*
- *Release resources when no long needed*
- *Pay for what you use*
- *Leverage other's core competencies*
- *Turn fixed cost into variable cost*

(Farber, 2008)

On demand self service is a key characteristic of cloud computing whereby the user would be able to request system resources near real-time and limit interaction with the service provider. Virtualization technologies cut down the provisioning time for delivery of systems. In the development lifecycle, there was often a gap between design to implementation of a system due to procurement constraints such as purchase orders, order fulfillment, and time to build the systems in a computer lab. Cloud computing gives the end-user (i.e. consumer) greater autonomy to specify required services and quickly provision the appropriate resources. The time to market is reduced, ultimately resulting in a shorter schedule, less management components, and possibly a competitive advantage be establishing an early market presence.

Cloud computing systems have broad network access capabilities that allow connectivity through standard protocols and from a variety of network, personal, and mobile devices. Resource pooling is another important characteristic of cloud computing systems. Various resources including network devices, security services, storage, application servers, and other components are pooled together in a fashion not limited by geography, location or geopolitical boundaries such as country, region, state, and municipality. Resources may be combined to provide for efficiency and priority scheduling.

Rapid elasticity is a distinctive feature of cloud computing systems. Sometimes referred to as “cloud bursting”, this capability allows the infrastructure to scale up quickly to meet peak demands in a near real time manner. For example if a widely publicized “go-live” for an application is planned, the cloud computing system can allocate additional resources to meet the increased demand. Capacity planning becomes easier with ability to scale up and down based on usage requirements. Disk, CPU, memory, storage, and other features can be adjusted as needed in a cloud computing system.

Measured services are a component of cloud computing that serves cost and performance functions. First the utilization of cloud services must be metered or measured in order to charge the cloud service consumer or user. One of the key features of the cloud is pay for what resources are used and only those resources. In other words idle compute resources are not billed to the consumer. Performance such as disk usage, CPU utilization rates, and memory consumption are key factors in planning for cloud bursting or appropriate billing levels to maintain the system. In summary, measured services are a cornerstone of cloud computing and enable accurate billing and appropriate performance management for cloud computing system resources.

Types of Cloud Computing

The term cloud computing encompasses a variety of different styles, configurations, and types. Research organization Blakely and Reeves (2010) identify five deployment models of cloud computing:

- A *public cloud* offers IT capabilities as a service to any consumer over the public Internet.
- A *private cloud* offers IT capabilities as a service to a select group of consumers.
- An *internal cloud* is a private cloud by which an IT organization offers an IT capability as a service to its own business.
- An *external cloud* is an IT capability offered by a service provider to a third-party business.
- A *hybrid cloud* is an IT capability offered as a service using both internal and external IT resources.

The public cloud is a consumer service that is widely available for use from service providers such as Amazon, Google, and IBM. The public cloud offers infrastructure services, platform services, storage services, and computing services that require no initial investment by consumers to begin using. The public cloud model is based on economies of scale and demand for resources varies within this model.

The private cloud offering is based on services similar to the public cloud, but accessible by a limited number of consumers. An example of the public cloud might by a consortium of universities or businesses with a common interest or purpose. The public cloud would serve to further their computing needs, but not be widely accessible by the general public. The United Kingdom government created the "G-Cloud" in June 2009. Analyst Philip Hunter describes the new government private cloud configuration as an "infrastructure dedicated to a related group of organizations, with the economies of scale and flexible provisioning coming from sharing of resources among different agencies and groups" (Hunter, 2010).

Internal clouds are limited to providing services within one organization, program or team. This might be considered a collaborative tool for proposal services or workgroup repository for a project team. Another example might be testing services for the internal organization where services would be deployed in the cloud, tested, and then deprovisioned.

External clouds are hosted by services providers and used by third-party businesses or organizations. These could be cloud services that are used external to the organization, but not openly available to the public. For example, a corporate conglomerate might use external clouds to communication between other business units or holding organizations. Another concept is federal agencies that need to communicate through an external cloud, but not make the services publicly available for data security, transaction type, or other contractual requirements.

Comparisons of Cloud Computing

Cloud computing invokes many comparisons to "utility" computing given the concepts of elasticity and ubiquity of services. In reality there are distinctive technical considerations for cloud computing and information technology in general that require evaluation to support the cloud computing to utility comparison. Brynjolfsson, et al. (2010) found "an overly simplistic reliance on the utility model risks blinding us to the real opportunities and challenges of cloud computing."

The utility model builds on economies of scale. Industry analyst Geva Perry (2008) contrasts utility computing and cloud computing:

> *Utility computing is seen as a business model, such as Amazon.com's "Amazon Web Service" (AWS) that rents storage space and access to companies, cloud computing refers more broadly to a computing architecture. This architecture links computers in a grid and allows users to buy access to data and software stored on the grid or processing power that is harnessed for specific purposes by the grid of computers. (Perry, 2008)*

Obstacles for Cloud Computing

Cloud computing presents a new way of processing data, a new computing environment, and a different way to manage information. The cloud offers advantages in terms of cost avoidance, cost reduction, and quicker time to market for certain business models. However, there are obstacles to cloud computing including security concerns in terms of trusted computing, information protection and identity management. Other obstacles include cloud standards, reliability, and vendor lock-in from cloud providers.

In a recent survey, respondent Barry Wellman, professor of sociology and Netlab director at the University of Toronto (O'Dell, 2010), summed up a number of factors with the cloud "Trust not the cloud for reliability, security, and privacy." Another respondent predicted a "huge blow-up with terrorism in the cloud" which would lead to a severe lack of confidence in cloud computing. John Chambers, CEO of Cisco, stated in his keynote speech at the 2009 RSA conference, [cloud computing] is a security nightmare, and it can't be handled in traditional ways" (Greene, 2009).

Analyst David Talbot stresses "What nobody has yet solved….is the security problem inherent in the size and structure of clouds" (Talbot, 2010). In 2009, three computer scientists exposed a security vulnerability by placing malicious virtual machines in Amazon's EC2 cloud system. While they did not steal any data (they were proving a point), the scientists were successful in their malicious attempts 40% of the time. State University of New York at Stony Brook computer scientist Radu Sion observes the current situation, mammoth-sized cloud hosting organizations which provide services to thousands of companies who co-host with them. The potential for a single data breach or malicious attacks in the cloud could potential affect thousands of companies and many more customers.

Internet connectivity is often cited as the Achilles' heel for cloud computing. If the Internet is not accessible from a location, cloud-based systems may not be accessible to the consumer. Julien St John-Dennis, head of business products at ntl:Telewest Business observes, "critical applications still relying on the UK's ageing legacy communications infrastructure could suffer downtime. As such, potential cost savings have to be balanced against a drop performance efficiency" (Hayes, 2009). Analysts argue that internal-based enterprise networks incur outages as well. Applications hosted internally have regular maintenance windows for

reboots, recycles, patching, and other outages. This may offset the "connectivity problem" for cloud computing, but must be coupled with other concerns such as security, affordability, and scope of control.

Data Protection

90% of cloud application users say they would be very concerned if the company at which their data were stored sold it to another party (Pew, 2008). Many users cite "efficiency and convenience" as reasons to use cloud computing, however Gartner research analyst Daryl Plummer states that "moving *corporate* data that requires frequent updates out into a cloud is time-consuming, risky and impractical." There may be complexities in identifying, coding, and exporting data from corporate managed systems to the cloud. In addition to the time, effort, cost, and complexity of migrating from internal systems to cloud systems is the sensitive information protection issue. Plummer observed that "some [cloud computing] vendors don't even make it clear if you still own the data or where it's located" (Pachner, 2010).

Reliability is another key concern for potential adopters of cloud computing. There have been a number of recent data breaches and failures in cloud computing systems. In a 2009 filing with the SEC, Google identified vulnerabilities with data, data centers, and cloud systems. While some of this is legalese, it does address concerns with reliability in the cloud. The following excerpt is from SEC form 10-Q filing:

> *"The availability of our products and services depends on the continuing operation of our information technology and communications systems. Our systems are vulnerable to damage or interruption from earthquakes, terrorist attacks, floods, fires, power loss, telecommunications failures, computer viruses, computer denial of service attacks, or other attempts to harm our systems.*

> *"Some of our data centers are located in areas with a high risk of major earthquakes. Our data centers are also subject to break-ins, sabotage, and intentional acts of vandalism, and to potential disruptions if the operators of these facilities have financial difficulties. Some of our systems are not fully redundant, and our disaster recovery planning cannot account for all eventualities," the company writes.*
>
> *"The occurrence of a natural disaster, a decision to close a facility we are using without adequate notice for financial reasons, or other unanticipated problems at our data centers could result in lengthy interruptions in our service. In addition, our products and services are highly technical and complex and may contain errors or vulnerabilities.*
>
> *"Any errors or vulnerabilities in our products and services, or damage to or failure of our systems could result in interruptions in our services, which could reduce our revenues and profits, and damage our brand." (Google, p. 47)*

Of particular interest to cloud computing users is the statement "*Some of our systems are not fully redundant."* Google points out its Google Apps cloud offering was the first to receive Federal Information Security Management Act (FISMA) accreditation and certification. Public law 107-347, also known as the Federal Information Security Management Act of 2002or e-Government Act of 2002, requires each federal agency "to develop, document, and implement an agency-wide program to provide information security for the information and information systems that support the operations and assets of the agency…" (FISMA, 2002).

Google continues to offer 99.9% uptime with no-scheduled outages despite the SEC filing disclaimers.

Another cloud computing service provider, Amazon issued a similar statement in its 2010 10-K SEC filing, stating:

> *"Our computer and communications systems and operations could be damaged or interrupted by fire, flood, power loss, telecommunications failure, earthquakes, acts of war or terrorism, acts of God, computer viruses, physical or electronic break-ins, and similar events or disruptions. Any of these events could cause system interruption, delays, and loss of critical data, and could prevent us from accepting and fulfilling customer orders." (Amazon, p. 11)*

Hype and Confusion in the Cloud

Vivek Kundra, CIO of the U.S. Federal Government is an advocate of cloud computing, "I believe it's the future," he says. "It's moving technology leaders away from just owning assets, deploying assets and maintaining assets to fundamentally changing the way services are delivered" (CIO, 2008). Larry Ellison, CEO of Oracle famously quipped in 2008, ""The interesting thing about cloud computing is that we've redefined cloud computing to include everything that we already do." The term invokes confusion and a pessimistic view due to the fact so many technology companies are presenting themselves as leaders in cloud computing and the combination of changing definitions (15 revisions by NIST) and the emerging standards bodies within the field.

According to Kundra cloud computing is definitely not hype (CIO, 2008). Kundra's statement on cloud computing hype is contrasted by Dunan Stewart, technology analyst with research and consulting firm Deloitte Canada. According to Stewart, "There's this idea that this brave new world of cloud computing will sweep the planet, we'll throw out our computers and just use them in the cloud. That's goofy" (Pachner, 2010). Stewart feels cloud promoters are

underestimating financial and logistical barriers associated with moving to the cloud. Still Slack (2009) advises “those who dismiss it [cloud computing] as "just another trend" are likely to miss out on the opportunities that cloud computing provides for organizations of all sizes.” Brynjolfsson, Talbot, and others recognize the potential upside to cloud computing: innovation catalyst, affordability targets, and increased efficiency in the data center. The key to unlocking cloud computing is addressing the security and trust challenges.

Information Assurance

Information Assurance (IA) is critical to the protection of organizational assets and systems security. The goal of IA is to protect information as a critical resource through policies, procedures, practices, and implementation guidelines. IA provides capabilities in form of security controls, access mechanisms, credentialing, and connectivity to protect, defend, and provide a reasonable amount of data integrity. The Department of Defense, in specification DODD 8500.1 (2002), defines Information Assurance as:

> *Measures that protect and defend information and information systems by ensuring their availability, integrity, authentication, confidentiality, and non-repudiation. This includes providing for restoration of information systems by incorporating protection, detection, and reaction capabilities.*

Given the breadth of IA responsibilities, “supporting IA structures” are required to help in the monitoring and reporting of events. Supporting IA structures encompass are the infrastructure components primarily designed to alert systems managers of potential compromise of data, malicious attempts, sabotage, and intrusions. Reporting mechanisms are used to maintain system integrity, provide compliance status, and identify anomalies or exceptions.

One part of the Information Assurance (IA) framework is *identity* assurance. Simply put, identity assurance is the ability to determine a user's identity with a level of certainty; in other words the user presenting the identity credential is who he claims to be. Proofing of an identity has a direct relationship to the overall IA goal of protecting and defending information systems by ensuring *integrity, authentication, confidentiality, and non-repudiation*. Identity assurance centers on specific functionality and technologies associated with user identities, profiles, accounts, and credentials. The identity lifecycle includes issuance of a credential, storage of the credential, and ultimately disposition. Secure storage and protection of identity-related data is on component of identity assurance. Typically storage and retrieval for identity attributes are facilitated by relational databases, lightweight directory access protocol (LDAP), and various meta-directories or virtual directories. Often referred to as "AAA", authentication, authorization, and audit of identity and identity credentials issued or recognized by trusted systems is within the realm of identity assurance. Finally, federation and single-sign on capabilities are elements of identity assurance. As systems migrate into the cloud, identity assurance will play an increasingly critical part of the IA framework.

Former Deputy Under Secretary of Defense Donald C. Latham, points out vulnerabilities of civil and public sector critical information systems including financial and telecommunications sectors. To help mitigate the risks associate with attacks against these systems, Latham (2005) suggests a "multifaceted solution that addresses more thorough vetting of employees with critical access to telecommunications, networks, computers, servers, and other related equipment and software." Latham describes the need for identity assurance, including provisioning, vetting, and issuance for accounts and access to systems. With identity assurance, proper controls are deployed as enforcement mechanisms. Identity management spans the

lifecycle from creation, issuance, management, storage, processing, changes, and disposition of user information and credentials.

Information Assurance in Practice

In practice, information assurance is about managing risks to specific data assets and systems within a defined tolerance. IA practitioners follow several steps to secure systems and assets: 1) conduct a risk assessment, 2) assign assurance level, 3) select IA technology, and 4) validate the controls meet assurance levels. Many will build IA programs based on an inventory of critical assets and assign risk levels based on factors such as sensitivity of the information, financial loss, harm to the organization, or liability. To apply the appropriate security measures and controls, each asset can be logical placed in an assurance level. U.S. National Institute of Standards and Technology (NIST) 153 Special Publication 800-63 version 1.0.2 [NIST800-63] has defined four levels of assurance based on the Office of Management and Budget (OMB) E-Authentication Guidance for Federal Agencies [OMB M-04-04]. The Kantara Initiative summarizes this framework for information assurance in Table 1.

Table 1.

Assurance Levels, Kantara Initiative

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

The Carnegie Mellon University Software Engineering Institute has produced the Advanced Information Assurance Handbook for technical guidelines and best practices. The handbook covers IA techniques to harden systems, monitor systems, and implement intrusion detection systems (IDS) on a network. The handbook serves as a reference for specific systems and components that might be found within an IA program's scope. For example, the handbook covers systems hardening for the Windows 2000 and Red Hat Enterprise Linux operating systems, router configurations, logging practices, and computer forensics approaches. In terms of identity management and authentication, the handbook suggests technologies to verify an identity should be implemented within an organization. The handbook's assurance concepts are applicable to the realm of identity assurance and cloud computing making it a useful asset for the IA practitioner.

The Air Force Research Laboratory Information Directorate in Rome, N.Y., recently announced it seeks to establish a University Center of Excellence (UCoE) in Assured Cloud Computing. The researchers "want to develop ways to assess and influence the predictability of heterogeneous Air Force communication networks that assure data transfer, computations, and assured operations in hostile, contested, and high-interference environments" (Keller, 2010). European Network and Information Security Agency (ENISA), under the direction of the European Union (EU), created the Cloud Computing Information Assurance Framework (2009).

The document applies IA principles to cloud computing. Recommendations found in the framework for cloud identity providers include risk assessment analysis, cloud provider comparisons, and identity management related questions related to registration of identities, levels of assurance, and a number of de-provisioning controls.

Related Work

The Cloud Computing Use Case Discussion Group has produced 4 versions of the Cloud Computing Uses Cases White Paper (2010). The authors identify a cloud computing taxonomy, developed casual use cases, and various scenarios at a very high level. Version 3 of the document focuses on security and access controls. The security use cases include a customer scenario section, how the customer solved the problem, what requirements and controls were used in the solution, federation patterns, and roles. No preconditions, post conditions, extensions, flow of events, key scenarios, diagrams, or UML artifacts are included in the document. The standards to which these use cases are written do not meet the fully dressed criteria defined by Cockburn (2002). The document serves as a cloud computing primer, however, the use cases are not written to include sufficient detail for repeatability.

The MIT Kerberos Consortium released a series of use cases in using the term Kerberos-in-the-cloud (KITC) to describe various Kerberos authentication models in the cloud. The use cases in this document are best described as casual use cases with a paragraph or two describing a general scenario. No preconditions, post conditions, extensions, key scenarios, or UML/SysML artifacts are included in the document. One figure, with a list of event steps, is referenced by three use cases. Three additional use cases are identified, but listed as “to be determined”, thus the document may be considered a work in progress. The MIT-KC working

document serves as starting foundation for cloud computing with Kerberos technology. In future revisions, more detailed use cases will most likely be written.

CHAPTER 3

METHODS

Methodology is defined as "the study of scientific methods" whose objective is the "improvement of procedures and criteria employed in the conduct of scientific research" (Ackoff, 1962). The tools and techniques used in the research methodology were chosen based on current systems engineering practices for systems design and modeling. The systems engineering methodology and the design tools used were a combination of ArchStudio 4, Topcased, and Visio for SysML modeling.

Two research characteristics for this study focused on repeatability and clarity. Through use cases and systems models, repeatability was achieved by capturing the design artifacts and making them available for future research or systems design and implementation by research teams. The use cases were written with clarity, accurately capturing multiple actors, subsystems, conditions, and data flows with the intent that improvements and extended experiments may be conducted.

This study employed the use case and systems modeling approach to validate the following questions in this research project:

1. *How do we determine who is authorized to be on the cloud*
2. *What mechanisms exist to provide the identity management and access function?*
3. *How do we truly provide assured identity in cloud computing environments?*

4. *How do we interoperate with different identity and access mechanisms in a global enterprise?*

The use case method was chosen based on the need to define user interaction and identify changes in the way cloud computing affects end users. Talbot (2010) admits confusion surrounding cloud computing as well as security in the clouds. A series of use cases presented repeatable, testable scenarios that clarify and improve the processes for identity assurance, access control, and interoperability.

A *use case* is a prose description of a system's behavior when interacting with the outside world (Cockburn, 2002). First introduced by Jacobson as "usage scenarios", these informal artifacts give a general idea of how a system works. Jacobsen identified the Swedish term "anvendningsfall" which translates into "usage case" in English, which eventually became the shortened "use case." Use cases may vary in formality, detail, and depth, but their primary purpose is to depict a functional level action within a system. They include a purpose or goal of the use case and describe in detail how the system responds to an actor's input or interaction with the system.

This study used a "fully dressed" use case to depict technical systems processes that involve identity assurance in cloud computing environments. The fully dressed use case includes more detail than casual or brief use cases. A template to capture the use case was created based on similar work in U.S. Department of Health and Human Services Homeland Security Presidential Directive-12 (HSPD-12) Program Office Standard Operating Procedures and Use Cases. The HSPD-12 document describes standard procedures and use cases for a variety of scenarios involving federal employees using Personal Identity Verification (PIV) requirements for personal identification that meets minimum security objectives defined by the Office of

Management and Budget and the Federal Information Processing Standard 201 (FIPS 201) directive. The purpose of this document is to establish *repeatable processes*.

Following the principle that graphics reveal data (Tufte, 2001), each use case in this research included a SysML (Systems Modeling Language) diagram graphically showing the actors, subsystem components, and process initiation. Each use case featured an overview that summarizes the events with a brief description. Next a basic flow of events was used to capture decision points and actions. Alternate flows were included where use cases have multiple options or divergent paths. Sub flows were included if additional detail was needed, for example, if data attributes were needed from a subsystem process. Preconditions and post conditions are identified in the next two sections. These included dependencies on systems, data, input, output, and other use cases. Finally an extension point identified the point in the base use case where the behavior of an extension use case could be inserted (IBM, 2004). The use case format included sufficient detail in the sub sections to be considered fully dressed artifacts rather than brief or casual use cases. Equally important was the use case purpose specifically geared towards repeatable processes which are complimentary to the intent of this study.

The National Institute of Standards and Technology (NIST) recognized a need to jumpstart the adoption of cloud computing. In May 2010, Badger and Grance introduced the Standards Acceleration Jumpstarting Adoption of Cloud Computing (SAJACC) to address critical cloud features such as interoperability, portability, and security *while* standards were being created, often a time consuming process. To promote cloud computing in the near term, while standards are being created, they developed a process to test import cloud system requirements, the SAJACC. The SAJACC communication strategy is built on creating cloud computing use cases that are known to work and can easily be used by cloud users, providers,

and extensible. Badger and Grance indicate that "use cases will provide insight on how clouds can work" (Badger and Grance, 2010).

Languages

The AIMS system components were designed in a modeling language called Systems Modeling Language or SysML for a couple of reasons. SysML is an extensible language that helps visualize system constraints, integration points, and the interaction between AIMS system components. The model allows for repeatability by using a standards based extensible language. The language was chosen based on the researcher's familiarity with the language as well as the commonality with industry standards and practice. SysML is commonly used to demonstrate systems models and widely compatible with open source and third party modeling tools.

SysML was adopted by the standards body OMG (Object Modeling Group) in 2008 to visual depicts integration among architecture components. SysML descriptors were used to capture attributes for virtualized and cloud-based systems. xADL 2.0 is modular framework described as "a software architecture description language (ADL) developed by the University of California, Irvine for modeling the architecture of software systems. Unlike many other ADLs, xADL 2.0 is defined as a set of XML schemas. This gives xADL 2.0 unprecedented extensibility and flexibility, as well as basic support from the many available commercial XML tools." xADL 2.0 is customizable and offers an extensible framework for compatibility with numerous systems architecture disciplines including virtualization. The extensions include variants that define connector types, version types, structure types, Java source code implementation, lookup function, message extension, and access control extension. xADL 2.0 includes the following items:

- Components (the loci of computation),

- Connectors (the loci of communication),
- Interfaces (the exposed entry and exit points for components and connectors), and
- Configurations (topological arrangements of components and connectors as realized by links)

xADL 2.0 is an XML-based language, easily programmable with any XML editor such as Microsoft FrontPage, Adobe Dreamweaver, and other common web editing software. The text and tag structure should be instantly recognizable to systems architects, developers, and systems engineers as displayed in Figures 3 and 4.

```
<types:component xsi:type="types:Component" types:id="xArchADT">
  <types:description xsi:type="instance:Description">xArchADT</types:description>
  <types:interface xsi:type="types:Interface" types:id="xArchADT.IFACE_TOP">
    <types:description xsi:type="instance:Description">xArchADT Top Interface</types:description>
    <types:direction xsi:type="instance:Direction">inout</types:direction>
    <types:type xsi:type="instance:XMLLink" xlink:type="simple" xlink:href="#C2TopType" />
  </types:interface>
  <types:interface xsi:type="types:Interface" types:id="xArchADT.IFACE_BOTTOM">
    <types:description xsi:type="instance:Description">xArchADT Bottom Interface</types:description>
    <types:direction xsi:type="instance:Direction">inout</types:direction>
    <types:type xsi:type="instance:XMLLink" xlink:type="simple" xlink:href="#C2BottomType" />
  </types:interface>
  <types:type xsi:type="instance:XMLLink" xlink:type="simple" xlink:href="#xArchADT_type" />
</types:component>
```

Figure 3.

xADL 2.0 Component Description

xADL 2.0 relationships are maintained within a schema structure and accessible through an artifact such as the one in Figure 6.

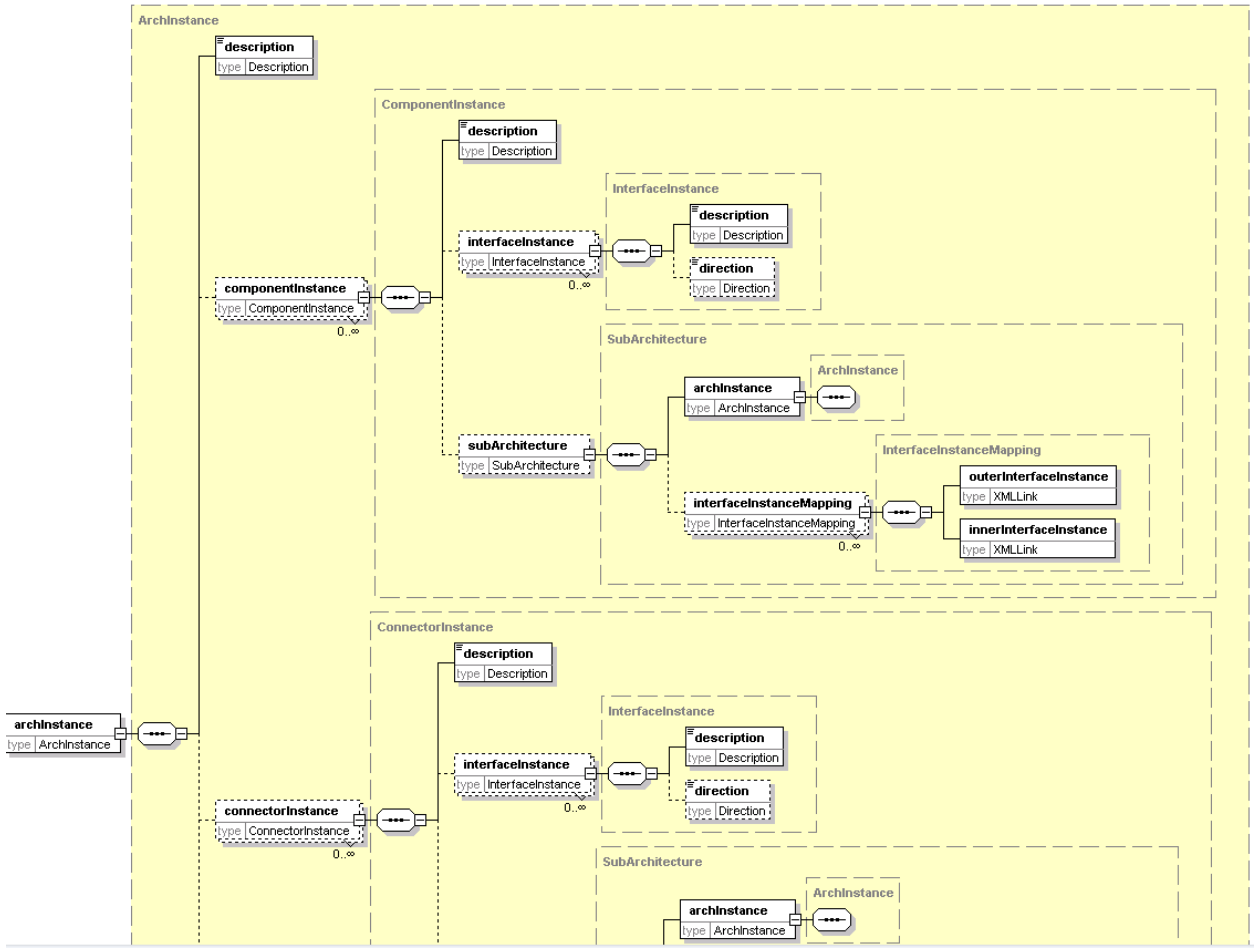

Figure 4.

xADL 2.0 XML Relationship (UC Irvine)

Modeling Tools

ArchStudio 4 is an open source architecture development and modeling tool based on xADL 2.0 architecture description language specifications. ArchStudio 4 is an environment of integrated tools for modeling, visualizing, analyzing, and implementing software and systems architectures (de Lemos, 2008). The software includes objects for design and developing highly complex, dynamic system architectures and models including the focus of this study, cloud computing and identity assurance models. These objects include connectors, interfaces, product-

lines, and various system components. ArchStudio 4 requires three components for the system to operate:

1. A Java 2™ Standard Edition (J2SE) version 5.0 virtual machine;
2. Eclipse version 3.5 (Galileo); and
3. ArchStudio 4 application

ArchStudio 4 is a meta-modeling tool that allows various stakeholders to tailor their own view of the architecture, based on semantics, and for their specific requirements. This allows various stakeholders to explore the architecture model from their perspective without having to re-create project artifacts, design documentation, and interface specifications. ArchStudio 4 provides visualization capabilities that not only show architecture diagrams, but allows users to edit the object properties with minimal recoding and programming.
ArchStudio 4 has a number of different components including Archipelago, ArchLight, and Type Wrangler.

The Archipelago feature of ArchStudio allows users to build visual diagrams with connectors and customizable properties. One of the strong features of ArchStudio is that when a change is made to an object, the architecture model is dynamically updated to reflect the new object property throughout the system (see Figure 5).

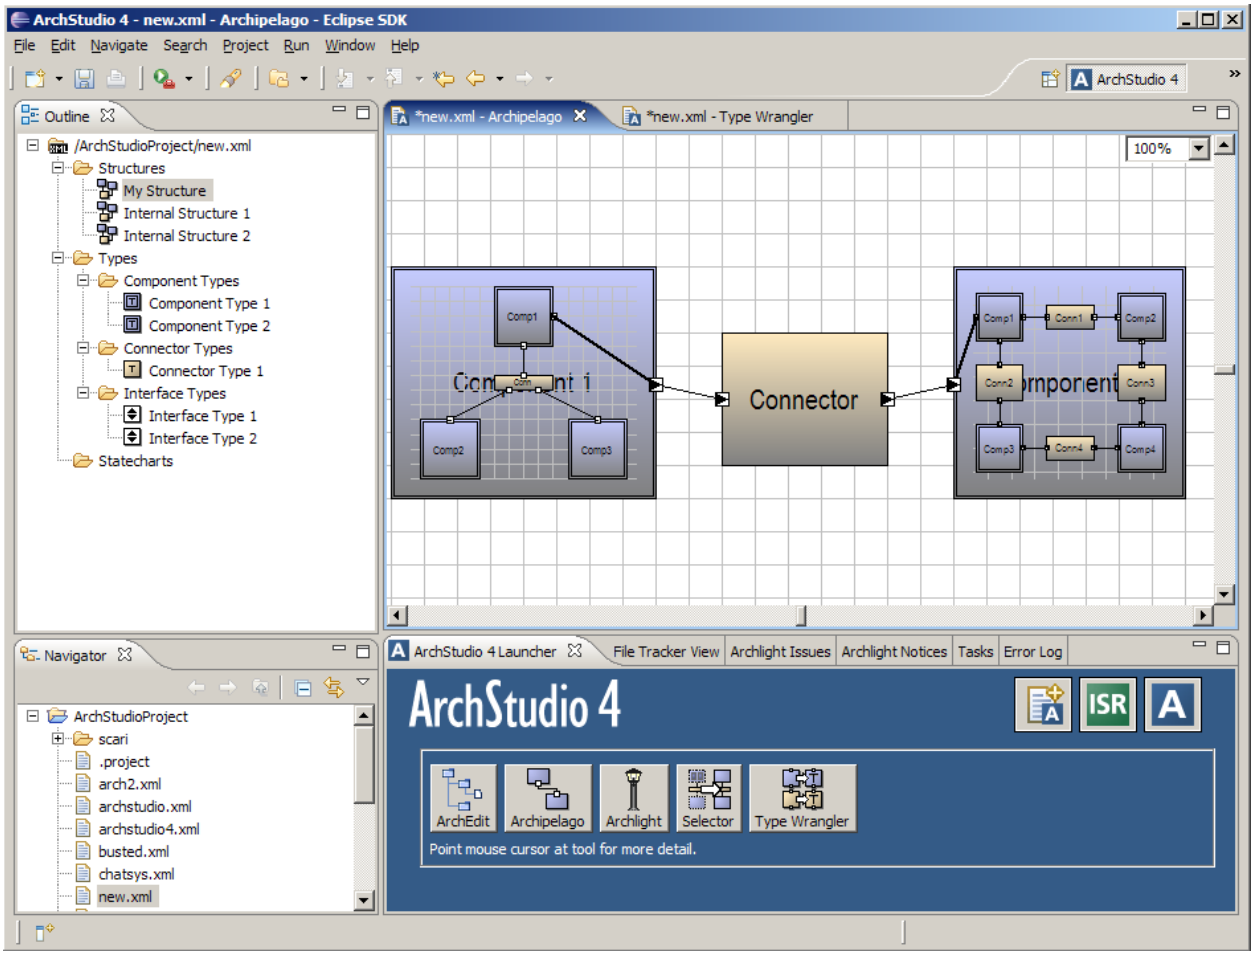

Figure 5.

Archipelago visual tool in ArchStudio 4 (UC Irvine)

Archedit is a component of ArchStudio 4 that allows XML syntax to be edited. Other editors may be used, but Archedit is built into the integrated framework and syntax changes made in Archlight are immediately reflected throughout the architecture model (see Figure 6).

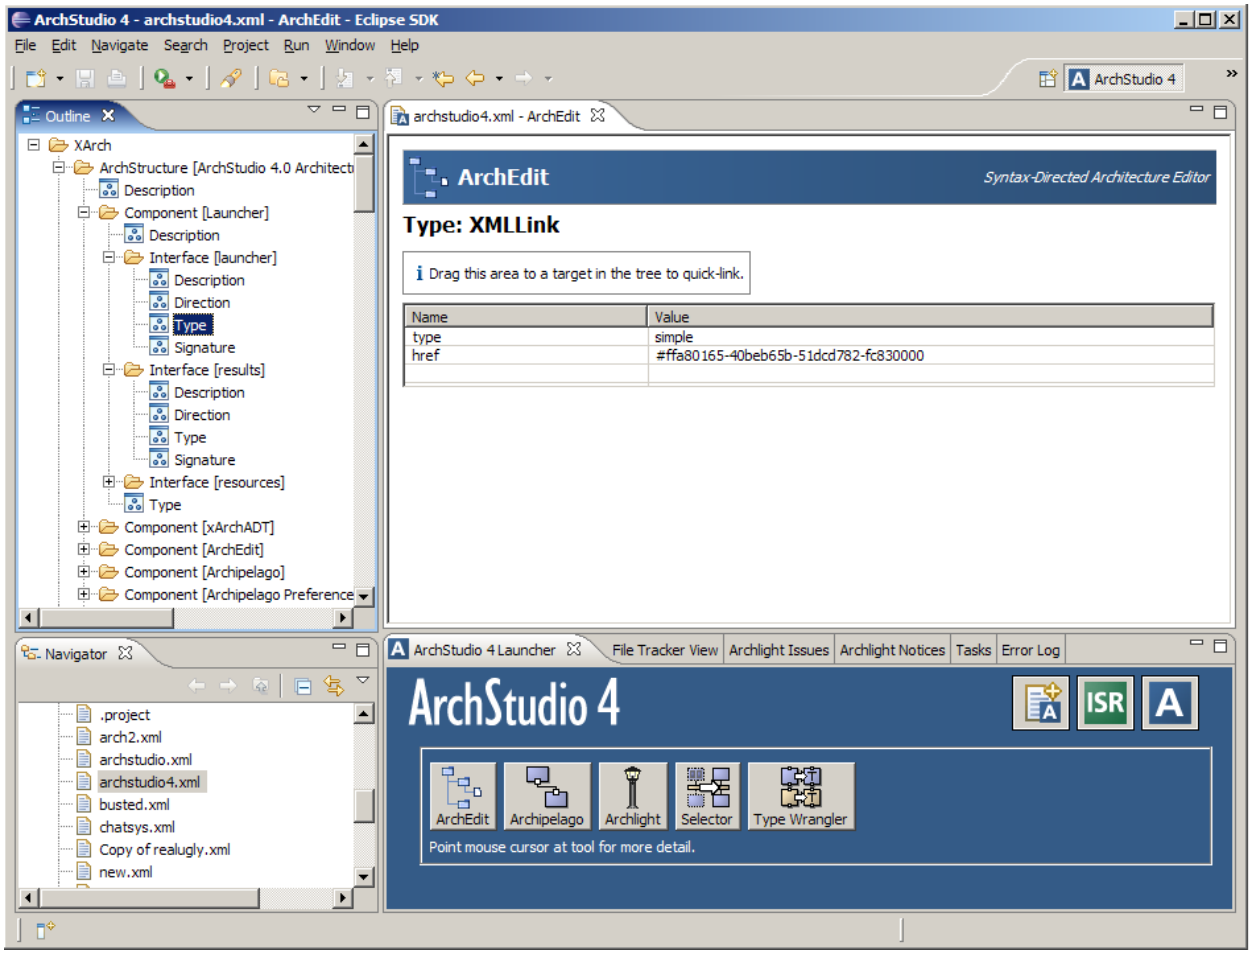

Figure 6.

Archedit XML syntax editor, ArchStudio 4 (UC Irvine)

Archlight is an analysis tool included in ArchStudio 4. Archlight has the capability to run tests and present findings such as anomalies, inconsistencies, discrepancies, and failed components. The Type Wrangler tool checks for and enforces consistency throughout the systems architecture model (see Figure 7).

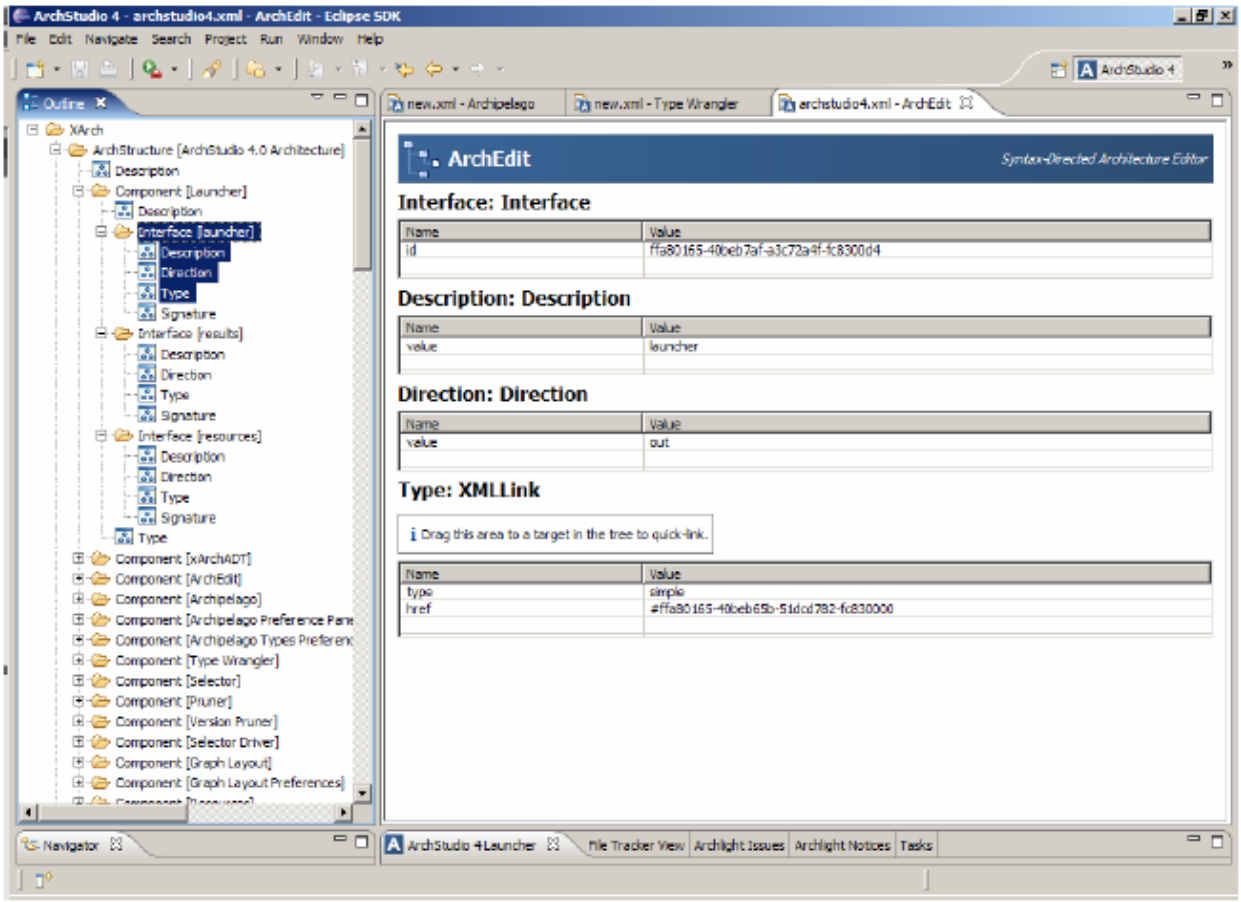

Figure 7.

ArchStudio 4 Interface deployed in Eclipse framework (UC Irvine)

ArchStudio 4 is deployed on the Eclipse integrated development framework. Eclipse is an open source tool that allows users and developers to collaborate together. Eclipse has the ability to allow "plug-in" applications for specific implementations. For this research project ArchStudio 4 is a plug-in software component used with the Eclipse framework. Eclipse is a good fit for ArchStudio 4 given its advantages in component reuse, trust, simplicity. ArchStudio 4 is a logical extension of the Eclipse collaborative built on data exchange and quality interaction. The Eclipse Platform is made available under a Common Public License (CPL).

Microsoft Visio was used as a modeling tool as well. Visio objects do not offer the relationship integrity and architecture descriptors, but the interface is easily navigated and the product is widely available as part of the Microsoft Office Suite of tools. A set of SysML stencils were created to accurately model package diagrams, requirements, sequence diagrams, and other SysML artifacts. The deliverables were SysML notation and included a variety of views and process scenarios. Figure 8 shows the Visio 2010 product user interface with the SysML stencils imported into the available "shapes" (see Figure 8).

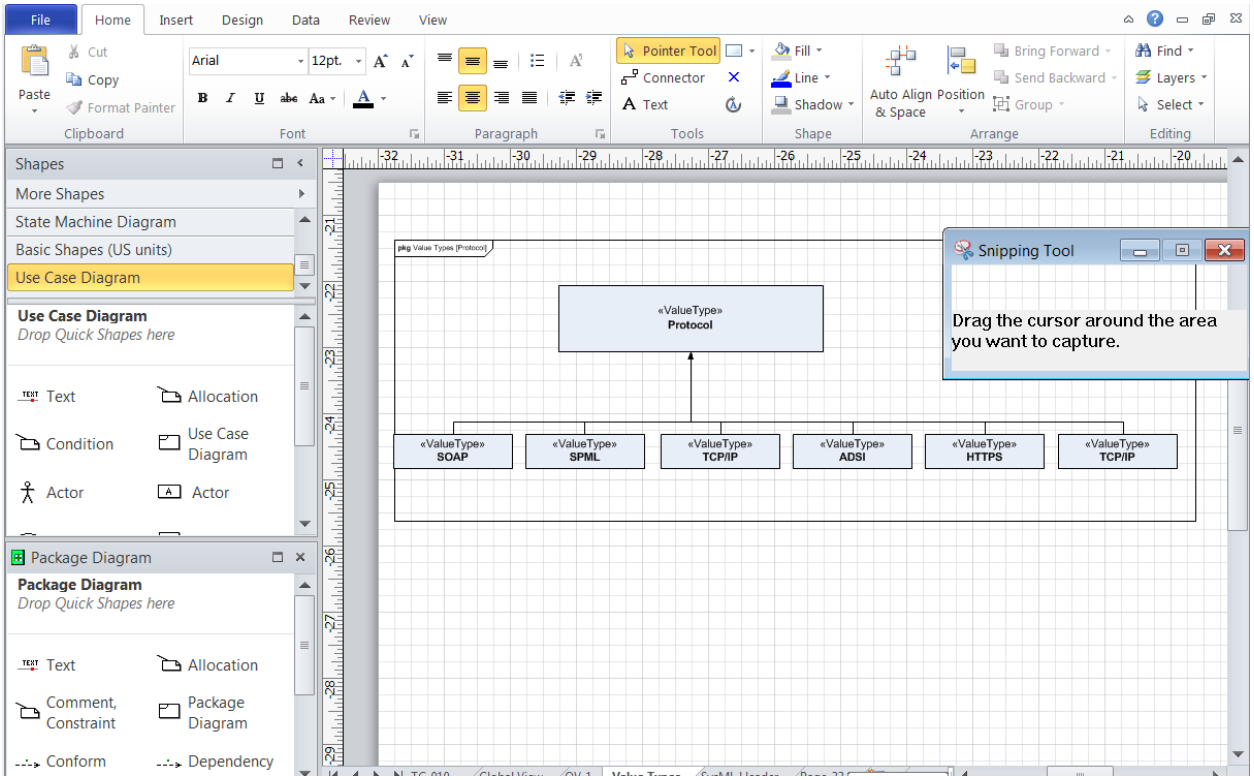

Figure 8.

Visio SysML Stencils

Topcased is an open source systems design and modeling tool based out of France. Topcased offers advanced modeling capabilities in terms of behavior, design attributes,

classification of system objects, and other features. Topcased offers an easy to use, pre-configured Eclipse deployment and many helpful support resources including tutorials, customer presentations, and examples of SysML work products (see Figure 9).

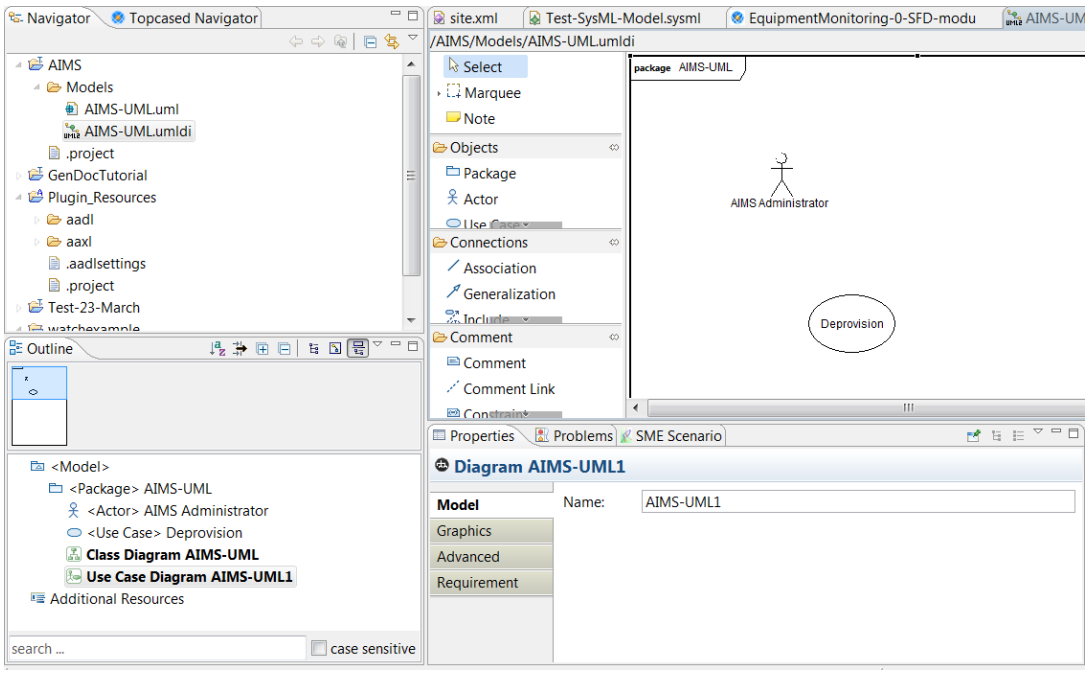

Figure 9.

Topcased User Interface with AIMS package diagram

Topcased has a similar look and feel to ArchStudio 4 given that both are deployed in the Eclipse framework. The figure above shows the similarities between "panes" in the user interface including the outline, navigator, and properties type objects. SysML descriptors and relationships for the use cases were modeling in Topcased.

To summarize the methodology tools strategy, cloud computing use cases focusing on identity assurance in the cloud were created using a word processor desktop productivity tool such as Microsoft Word 2010. The use cases were modeled using the SysML language, in a combination of ArchStudio 4, Topcased, and Visio applications. The end product was a

combination of fully dressed (detailed) use cases and SysML models that represent tested, repeatable processes. The work products and system artifacts tested the questions posed by this research study including how to identify who is on the cloud, what mechanisms are used for identity and access management, identity assurance and interoperability for global enterprise organizations.

CHAPTER 4

RESULTS

A prototype system was created to validate the use cases developed in the research investigation. The system revolves around two primary subsystems, the Assured Identity Management System (AIMS) and the Cloud Identity Management System (CIMS). The two subsystems communicate with each other facilitating the exchange of identity records and information according to the use cases. The first two research questions posed at the beginning of this study provide context for investigating records in the cloud and what mechanisms provide identity and access management. The latter two questions pose cloud interoperability and assurance questions. The main question of the study is:

Research Question: How can identity and access management controls be designed to support cloud computing systems?

The approach to this question was to analyze process controls by creating use cases to identify control points, interfaces, and communication paths among subsystem components. The concept is to define a process, then apply technology solutions in a prototype system to validate the use cases. Breaking down the question into four specific research questions, the investigation answered the following:

Research Question 1: How do we determine who is authorized to be on the cloud?

Research Question 2: What mechanisms exist to provide the identity management and access function?

Research Question 3: How do we truly provide assured identity in cloud computing environments?

Research Question 4: How do we interoperate with different identity and access mechanisms in a global enterprise?

Systems Concept

The Assured Identity Management System (AIMS) facilitates account lifecycle events for cloud computing environments. AIMS provides consistent, accurate, up-to-date identity-related data to manage accounts and credentials (user IDs, passwords, attributes, etc.) across enterprise clouds. AIMS integrated with the Cloud Identity Management System (CIMS) to demonstrate exchange of identity data from on-premise credentialing to a cloud environment. The Cloud Identity Management System was used to extend user accounts, define application roles, and store custom identity attributes as required for use by various computing applications. AIMS' role was to provide identity assurance to the CIMS, thereby ensuring a level of trust that the data provided from AIMS was accurate and delivered in a near-real-time fashion.

The OV-1 operational view diagram in Figure 10 shows the conceptual framework of the system that was designed for this study. This diagram is commonly found in the Department of Defense Architecture Framework (DODAF) reference model. The principal investigator chose to use this diagram style from DODAF because of the simple depiction of services and components from a high-level. The SysML diagrams found later in the study tend to be detailed with less creative illustrations. Systems engineers involved in designing enterprise-class systems

use conceptual diagrams as a way to simplify highly complex system of systems. The diagram also serves as a launching point to understand the system components from a visual perspective.

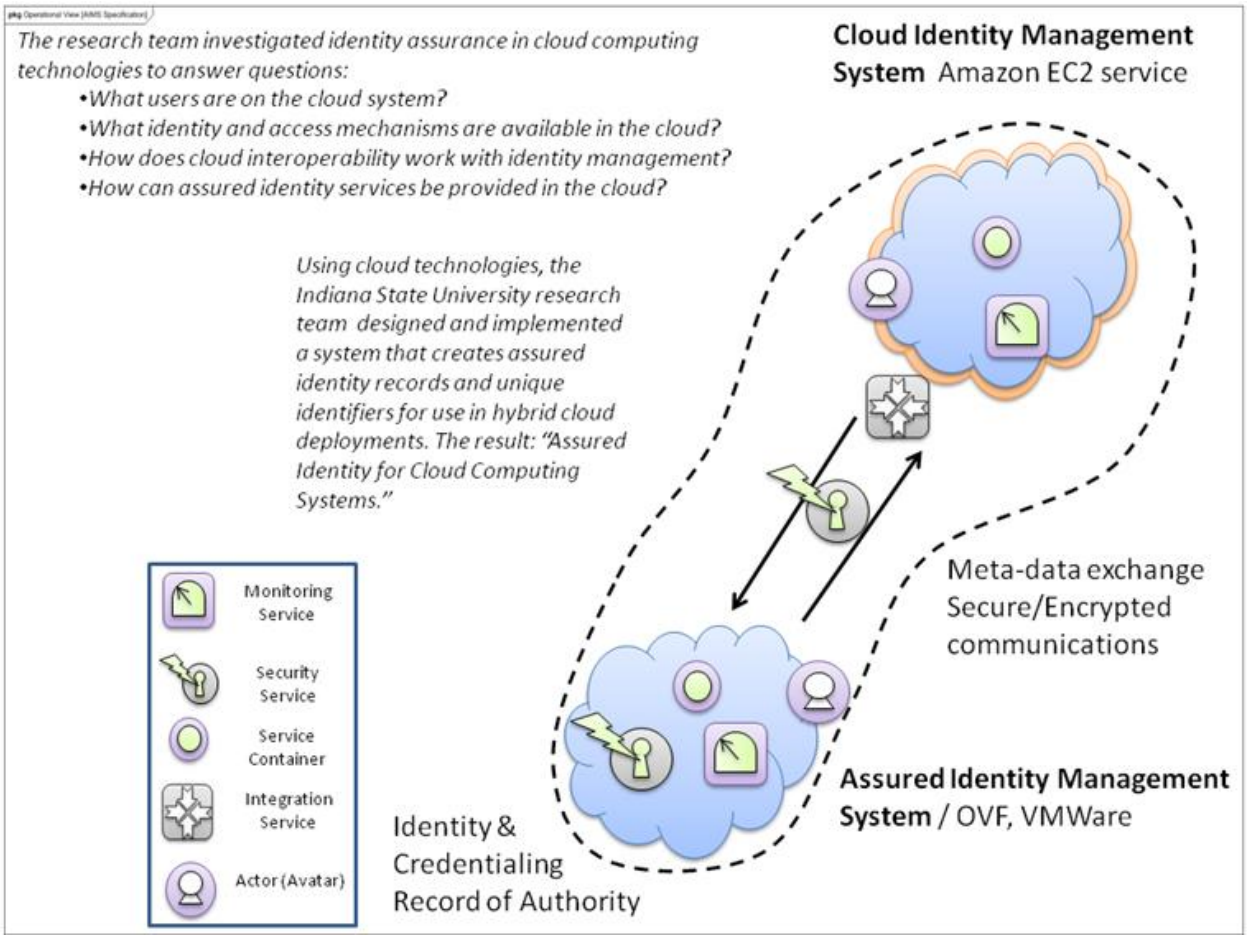

Figure 10.

OV-1 Operational Concept View, SysML Package

The OV-1 diagram shows two systems, the AIMS cloud on the bottom and the CIMS cloud on top with a system boundary represented by the dotted line. The diagram uses cloud notation that is currently evolving in the cloud research community. The notation shown here was proposed at the Open Group's (OMG) Cloudcamp meeting in Boston, MA on 21 July, 2010.

Key components within the clouds include service containers, monitoring services, and security services. The meta-data exchange and bi-directional communication links from AIMS to CIMS are represented by the four-arrow integration services box. Actors are shown in the purple circles and are intended to depict any end-user of AIMS or CIMS including customers and administrators. This is the only diagram from the DODAF reference model; the remaining artifacts in the study are SysML generated diagrams.

Building the SysML Model

The package diagram for the AIMS model shows which components were delivered. The diagram work products include requirements, behaviors, use cases, test cases, and a model library. Each of these components was expanded upon with specific work packages and products associated with each block (see Figure 11). The deliverables are represented by the diagrams or views found in this investigation.

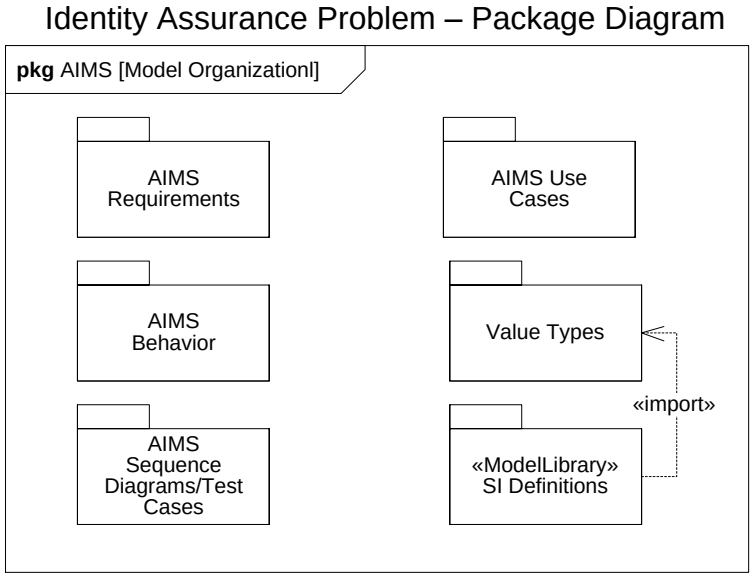

Figure 11.

Package Diagram for AIMS Organizational Model, SysML

The study conceptually designed a profile for the AIMS system. The idea of a SysML profile is a "template" that can be used for a specific domain. The profile offered helpful definitions and ideas that helped to create a common identity assurance model library. The AIMS model which contains attribute data and definitions (see Figure 12). The model feeds this information into the SysML profile.

Applying a Profile and Importing a Model Library

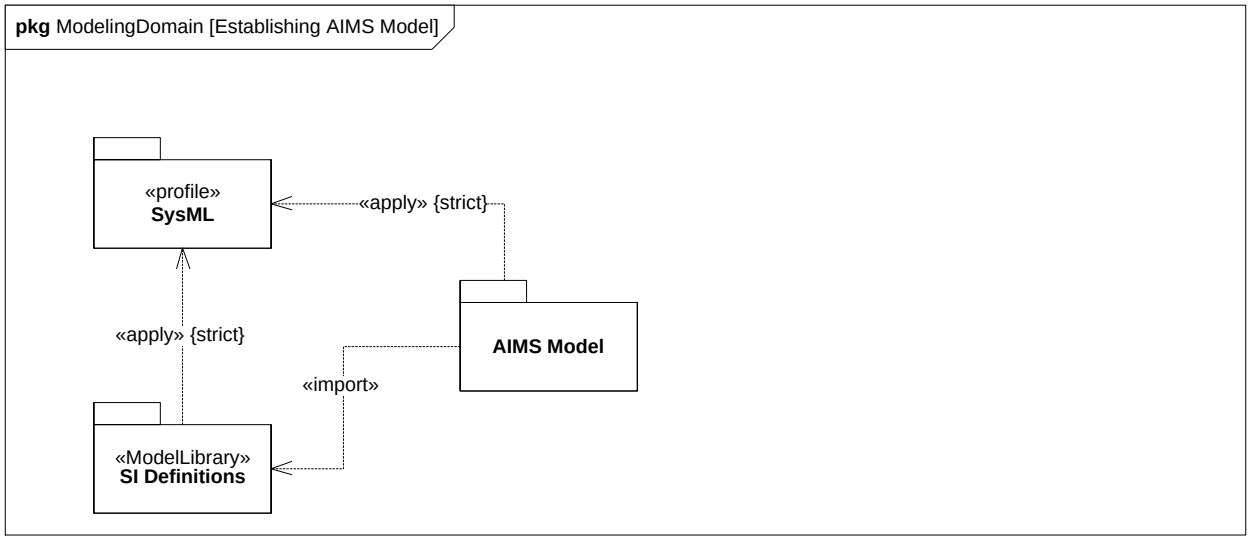

Figure 12.

Apply the Profile, SysML

The value types for this study focused on the communications protocols used to exchange data between the AIMS and CIMS components. A list of the protocol values included as part of the reference library is shown in Figure 13.

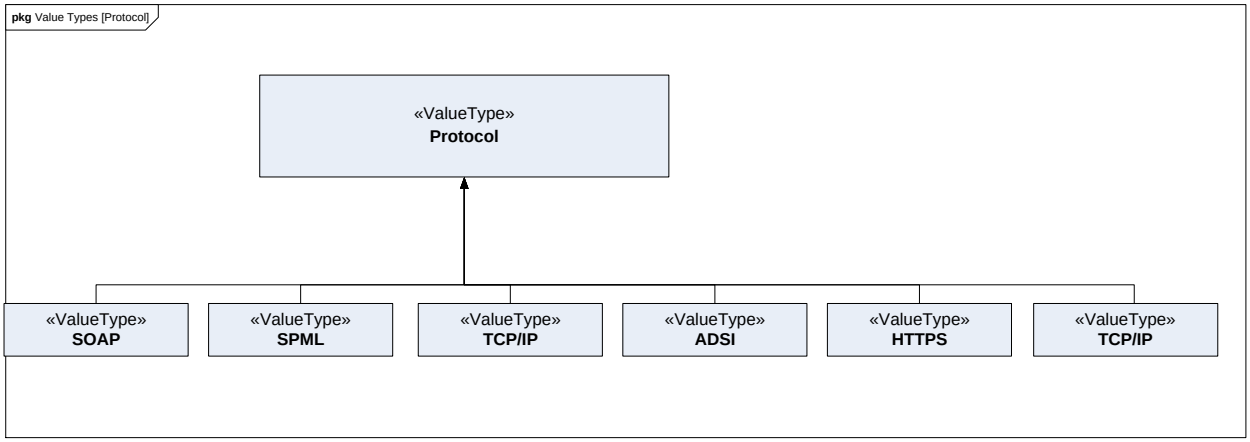

Figure 13.

Value Types Package Diagram, SysML

Systems Requirements

Requirements serve as the framework in which the design is built around. The research questions for this investigation are the source of requirements for the system. The requirements are structured in a manageable format with requirements identifier, description, and type of requirement. The various types of requirements include performance, functional, interface, computer software, data, and information security. They were gathered as constraints, features, and key goals for the system. The systems requirements are listed in Table 2.

Table 2.

System Requirements, SysML

Requirement ID	Requirement Description	Req Type
AIMS-REQ-1	The system shall support a network connection from CIMS to AIMS.	Systems Interface
AIMS-REQ-2	The network connection between CIMS and AIMS shall be mutually authenticated.	Systems Interface
AIMS-REQ-3	The network connection shall be encrypted with at least 128-bit encryption.	InfoSec

Table 2 *(continued)*

Requirement ID	Requirement Description	Req Type
AIMS-REQ-4	The network connection between the CIMS and AIMS shall be implemented using HTTPS.	InfoSec
AIMS-REQ-5	The system shall conform to the Service Provisioning Markup Language (SPML) 2.0 specification	Computer Software
AIMS-REQ-6	The system implementation shall utilize Sun Java System Identity Manager Version 8.1 (or higher).	Computer Software
AIMS-REQ-7	The system shall have a graphical user interface for the management and administration of identity records.	Functional
AIMS-REQ-8	The system shall create and maintain an atomic, unique value for each identity record	Functional
AIMS-REQ-9	The system shall link an identity record in AIMS to a corresponding record in CIMS	Functional
AIMS-REQ-10	The system shall process pipe-delimited flat files.	Functional
AIMS-REQ-11	The system shall have the ability to receive HR sponsored assurance data	Data
AIMS-REQ-12	The system shall provide a bi-directional interface between AIMS and CIMS.	Systems Interface
AIMS-REQ-13	The system shall support the 'Configurator' role for identity records management	Functional
AIMS-REQ-14	The system shall update interfacing subsystems within 5 minutes.	Performance
AIMS-REQ-15	The system shall maintain the current state of the subject associated with an identity record.	Functional
AIMS-REQ-16	The system shall provide a link to the state of resources linked to an identity record.	Functional
AIMS-REQ-17	The system shall have the capability to exchange meta-data attributes between CIMS and AIMS	Functional
AIMS-REQ-18	The system shall connect to Active Directory.	Functional
AIMS-REQ-19	The system shall implement active sync functionality	Functional
AIMS-REQ-20	The system shall implement reconciliation functionality	Functional
AIMS-REQ-21	The system shall implement initial seeding functionality	Functional
AIMS-REQ-22	The system shall manage logical access accounts (computer/information assets). Manage refers to create, remove, expire, archive, disable, re-enable, update, review.	Functional
AIMS-REQ-23	The system shall interface with Active Directory using the Active Directory System Interface (ADSI) interface specification.	Systems Interface
AIMS-REQ-24	The system shall log identity record provisions	InfoSec
AIMS-REQ-25	The system shall log identity record updates	InfoSec
AIMS-REQ-26	The system shall log identity record deprovisions	InfoSec
AIMS-REQ-27	The system shall log identity record access revocations	InfoSec

Table 2 *(continued)*

Requirement ID	Requirement Description	Req Type
AIMS-REQ-28	The system shall log change in levels of assurance	InfoSec
AIMS-REQ-29	The system shall log secure identification of users	InfoSec
AIMS-REQ-30	The systems share provide capability for single-sign-on	InfoSec

The investigation identified 30 requirements from the following categories: functional, interface, information security, performance, data, and software. Some of the requirements were derived from the COTS (computer off the shelf) products such as Sun Identity Manager while others, such as the use of Microsoft Active Directory, were included due to their common presence in enterprise computing environments. The SysML requirements diagram is displayed in Figure 14. It graphically displays the requirements categories using the *«requirement»* designation. Each requirement is shown under one of the six categories which it is related to. For example, the audit and logging requirements can be found under the information security identifier. CIMS to AIMS, HR to AIMS, and Active Directory to AIMS requirements are located under the systems interface requirements category.

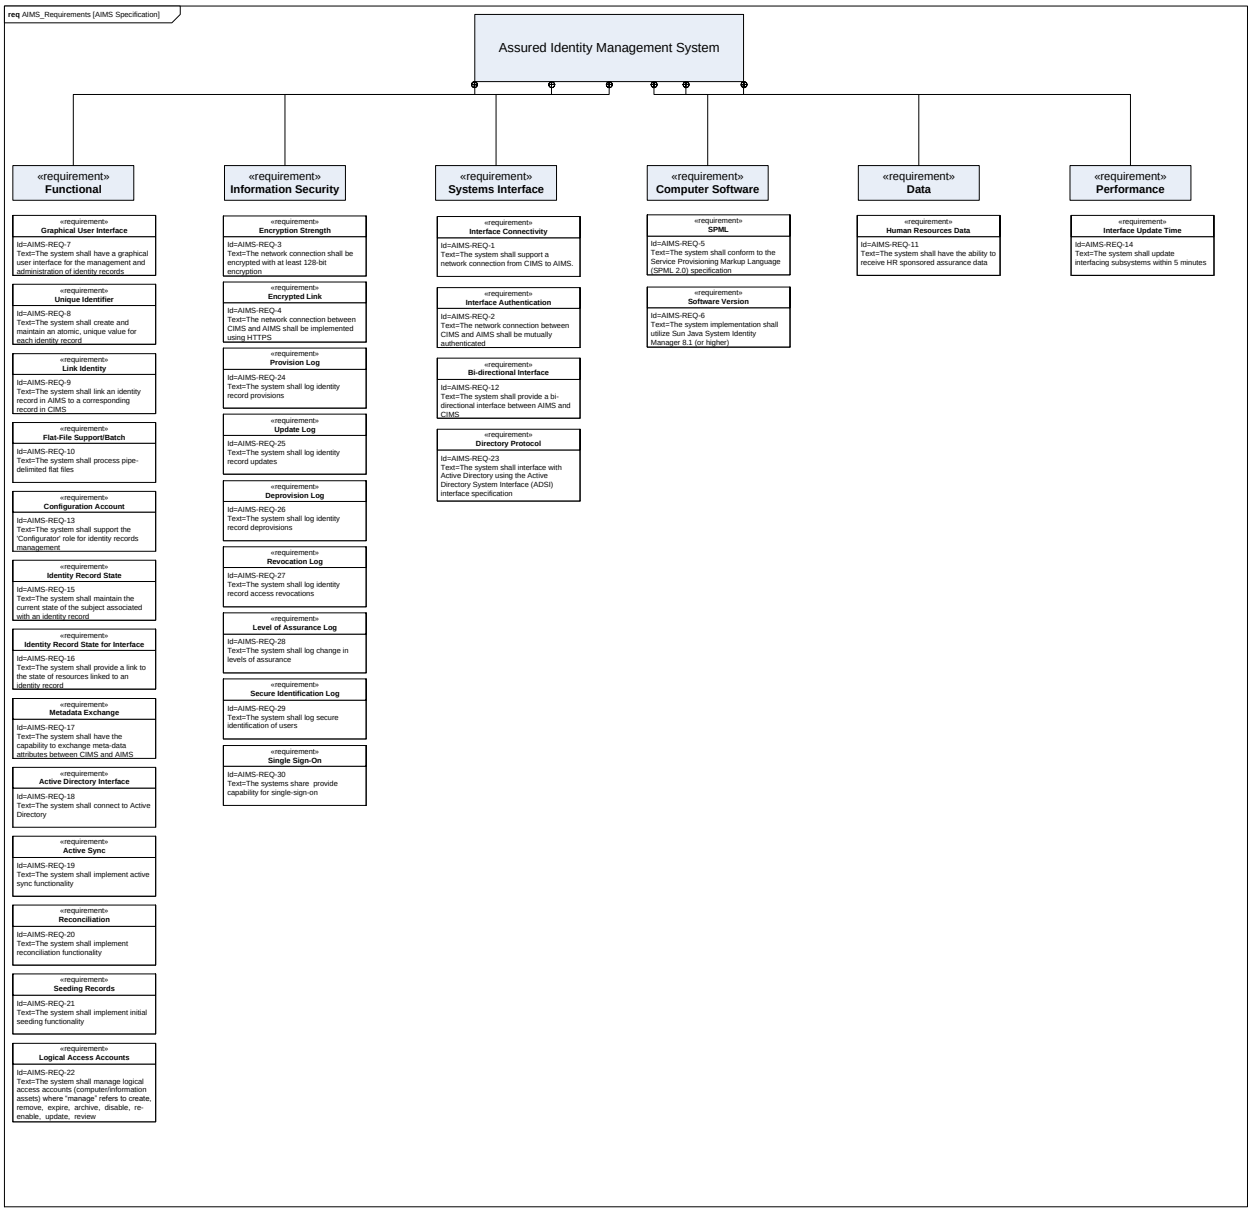

Figure 14.

Requirements Specification, SysML

Reference Architecture

The principal research investigator developed a reference architecture that was used to design and build the components and systems interfaces for the investigation. The reference

architecture served as a blueprint or a guide to follow through the design and implementation process. Coupled with architecture conventions and guiding principles, the reference architecture was a valuable tool that helped ensure consistency throughout the investigation. The reference architecture depicted in Figure 15 shows the various layers of authentication, authorization, and identity access interfaces. These layers loosely follow the cloud services model. The application layer consists of AIMS and CIMS applications, generally considered a management and services type of offering rather than an "end-user" application; these are basically systems management tools for identity and access records. The middleware application programming interfaces (APIs) allow connectivity to the platform layer and operating systems. These APIs serve as the conduit between the application and the operating system and directory services layers. The hardware layers included the physical network interface cards (NIC), disk storage, central processing unit (CPU), memory, and service bus for communication. The final layer is the directory services layer that included account records, authentication credentials, and related attributes such as roles and groups.

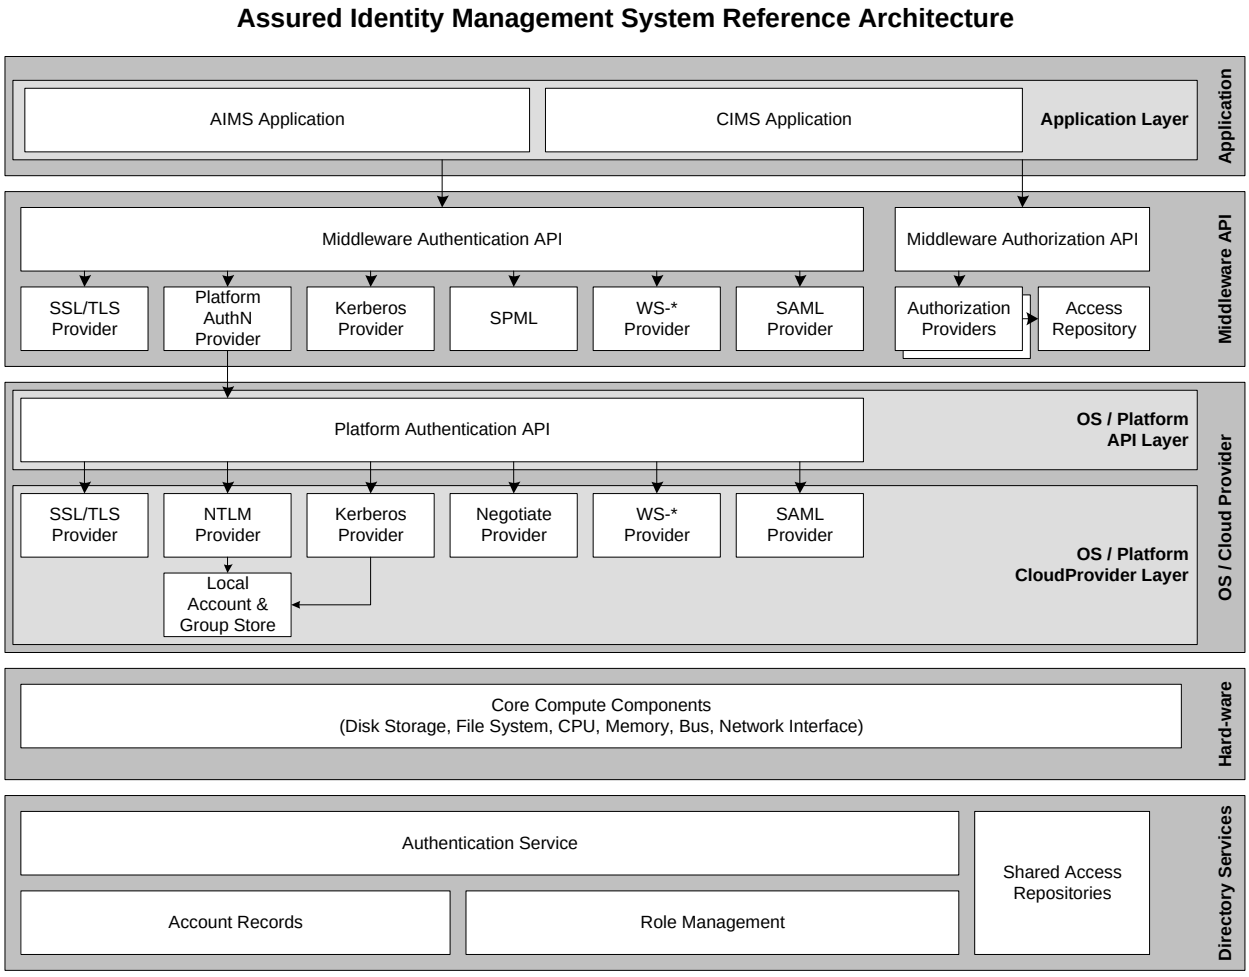

Figure 15.

Reference Architecture

Architecture Conventions

Several architecture conventions were used in the design of the research investigation. The first was to design a level of trust and identity assurance between systems to confidently establish the identity of all users, accounts, and records. This serves as the foundation and demonstrates the strength of AIMS in protecting cloud systems. Convergence of technologies within the infrastructure was a key architectural consideration. Virtualization technologies

continue to emerge as evident with the virtual disk and open virtual file system formats. The research project sought to converge identity and virtualization technologies where possible, primarily at the platform, infrastructure and operating systems layers. Consistent access enforcement was another architecture convention applied for administration across the applications. Establishment of trust between the Assured Identity Management System and the Cloud Identity Management System can only be established with consistent access and trust mechanisms in place.

The investigation was defined as a system of systems based on design criteria, architecture conventions, and the configuration of the subsystems. Maier (2000) defines systems of systems by the following:

1. *Operational independence of system elements: Each system element fulfills valid purposes in its own right. If the system is disassembled the separate elements continue to function.*
2. *Managerial independence of system elements: Limited centralized control. System elements are, at least in part, managed for their own purposes*
3. *Evolutionary development: the system arises over time*
4. *Geographic distribution: System elements can readily exchange only information (there are counterexamples)*

Maier continues to describe the systems as independently owned and operated systems that have common rules and a united purpose, not unlike a franchise system. The AIMS and CIMS components maintained operational independence; however the research project established a level of trust for identity and credentialing data between the two subsystems. If the trust was dissolved or the identity ceased to flow from AIMS to CIMS and back, the two systems would

operate independently with the potential that the cloud based system would not have assurance that identities have been fully vetted. In other words, the CIMS would continue to operate at a lower level of assurance as identity data would not be provided securely from AIMS. Both AIMS and CIMS were located in geographically disparate locations. AIMS was an on-premise, co-located with the HR system, however the CIMS was hosted by a cloud-provider.

The systems interfaces were defined by three types including external, inter, and intra systems. The external interface was between the AIMS subsystem and the CIMS subsystem. The two subsystems are independent from each other, operated and managed differently, but share the common need for assured identity data management. The inter subsystem interface was between the HR system and AIMS where critical data is passed from the human resources system within the same organization. Intra subsystems are defined as those internal to a specific subsystem such as the application server to database connection within the AIMS application to Active Directory system.

AIMS assigned a unique identifier that represents a persistent fixture to each record stored in the system. This unique identifier was used to correlate records between AIMS and CIMS and further prevent non-repudiation. The CIMS system used the unique identifier during synchronization of attributes between the two systems. The unique identifier also helped manage authoritative sources of data and reduce duplication of records and identity data. AIMS created the unique accountID attribute for each record based on the UUID version 4 standard. The Cloud Identity Management System used this value as the unique identifier for the same entity in its system. Rather than have each subsystem create its own unique identifiers and cross-reference the attributes, the AIMS UUID value is authoritative and was used for all cloud systems and other integrated subsystems.

The Java version 4 UUID instantiation is built on a scheme relying on random numbers. This algorithm sets the version number as well as two reserved bits. All other bits are set using a random or pseudorandom data source. Version 4 UUIDs have the form xxxxxxxx-xxxx-4xxx-yxxx-xxxxxxxxxxxx with any hexadecimal digits for x but only one of 8, 9, A, or B for y. e.g. f47ac10b-58cc-4372-a567-0e02b2c3d479.

Near real-time change notifications represent another architecture convention used in this research investigation. Upon notification of events, changes in attributes, or account status, AIMS immediately notified CIMS of these changes. This is done for expediency as well as a recommended security practice. Account creation and attribute updates occurred in a near real time fashion to accurately reflect the state of the account or record. Transactions in queue may cause delay in business processes or have the potential to incur multiple requests for the same action. As accounts were deprovisioned, removed, and revoked, the system updated accordingly to reduce risk of exposure, data leakage, and unnecessary access to sensitive data. In a business or organizational setting, policy and audit practices dictate how long to archive the accounts and records. For the purposes of this investigation, revoked accounts and records remained in the system indefinitely.

Infrastructure

For the investigation, commercial off the shelf (COTS) components including virtual appliances were designed and implemented. The system consists of the following infrastructure components:

1. Assured Identity Management System – Ubuntu 8.1, Sun (Oracle) Identity Manager 8.1, Tomcat application server.

2. Cloud Identity Management System – Sun Identity Manager 8.1, Tomcat application server.
3. Active Directory – Microsoft Windows 2003sp1, Active Directory Target System.
4. Amazon Web Services EC2 - emulator host to simulate cloud.

The AIMS system was based on the Virtual Machine Disk Format v1.1, a specification developed by VMWare, Inc to describe and document virtual machine storage. The virtual machine consisted of six vmdk disk files and a vmx configuration file that contained descriptive information about the system. These virtual disk images were managed by the principal investigator and secured on the host file system.

The implementation of the Assured Identity Management System (AIMS) consisted of several layers of virtualization. AIMS utilized the "virtual appliance model" where complex development environments are constructed and ready for a quick deployment. Virtual appliances are pre-configured virtual machines that have a foundation of application settings, operating system settings, and in some cases database builds ready for deployment. Virtual appliances are helpful in demonstration, test, and proof of concept, and for short-term implementation routines.

The emerging Open Virtualization Format specification introduced by the Distributed Management Task Force in 2010 was used in addition to the VMDK file system configuration. The OVF format allowed for a portable, secure, and efficient standards-compliant packaging format for virtual machines. OVF supported validation of the virtual machine and contains extensible metadata descriptors related to the virtual machine convenient for distribution of the appliances. Given the open nature of this research investigation, the OVF was used to allow the widest distribution of virtual hosts and cloud service providers as well as the greatest flexibility.

In the interest of repeatability, the project may be repeated using different providers or hypervisor technologies using the pre-configured appliances and configuration components.

The VMDK disk files were converted to OVF format using the OVF Tool version 2.0.1 offered by VMWare. The OVF tool is available to convert .vmx, .vmdk, .vmsd, .vmxf, and .nvram files to OVF as well as import and export them for distribution. The command line utility used the following syntax:

ovftool C:\Users\sysadmin\AIMS-vm.vmx C:\Users\sysadmin\AIMS-vm.ovf

Validation of the OVF Tool output was confirmed using the schemaValidate command:

ovftool --schemaValidate package.ovf

In both cases the conversion from the virtual disk to the open virtualization format was successful. The OVF format demonstrated the portability and compatibility of the AIMS virtual application configuration.

The guest operating system was Ubuntu 8.1, a popular and well-known Linux distribution. Ubuntu was chosen for two main reasons. The first is ease of use for the principal investigator. Ubuntu 8.1 is a stable Linux release that is easily installed, configured, and "application friendly" in the sense that many open source and third party tools are available, compliant, and compatible with the operating system. The second reason for choosing Ubuntu is the cloud-ready feature of Ubuntu. The operating system can be bundled up as a virtual machine and deployed into the cloud using the Amazon EC2 cloud hosting service. EC2 is a Xen-based open source hypervisor technology offered through Amazon Web Services (AWS). The virtual appliances used in the configuration are compatible with multiple virtual hosts and cloud computing service providers. The portability of AIMS demonstrated a key characteristic of cloud systems.

Microsoft's Active Directory is a commonly-used COTS product and source of consolidated directory information. Given its deep market penetration in the enterprise market segment, AIMS utilized Active Directory to initialize the creation of identity records for the sample application. After initialization, the provisioned account record, the existing records were synchronized using scheduled tasks in the AIMS application. AIMS created and managed an identity record for all lifecycle events in the Active Directory.

AIMS-Active Directory synchronization involved three distinct processes. These processes represented the initialization of identity records through seeding, the continued synchronization between the subsystems, and finally the reconciliation process to identify and resolve potential data conflicts.

Seeding: These processes were used to load existing enterprise account data into AIMS. Seeding data loads relied on AIMS's reconciliation function.

Active Synchronization: Commonly referred to as 'Active Sync'; Active Synchronization consists of the processes used for synchronizing changes from an authoritative resource such as human resources, into AIMS on an on-going hour-to-hour basis.

Reconciliation: Reconciliation was used to periodically compare resource accounts in AIMS with the accounts actually present in the resources. Reconciliation correlates account data and performs a comparison of full account records.

These three processes were used to move account and identity data among subsystems in accordance with use cases.

Sun Identity Manager version 8.1 was chosen as the provisioning and account management workflow tool for several reasons. Industry research company Gartner Group has consistently placed the Identity Manager product in its leadership quadrant for identity

management and access tools. Although similar tools in the leadership quadrant include Oracle's Waveset application and Microsoft's Forefront Identity Lifecycle Manager, the principal investigator chose Sun's platform due to familiarity with the product and its SPML2 capabilities.

Sun made the Identity Manager product readily available for development, proof-of-concept, and demonstration purposes. Sun offered full product versions and access to documentation libraries for its identity middleware product line to consumers, researchers, and engineers. Since Oracle's 2010 acquisition of Sun Microsystems Corporation, the new owner has continued to provide software downloads for free, under the Developer License that allows full use of product versions at no charge while developing and prototyping applications, and for self-educational purposes.

The principal investigator was familiar with the workflow functionality of the Identity Manager. Based on the XPRESS language, customized workflows were created for account creation in AIMS. XPRESS uses the Netbeans IDE 6.1 integrated development environment. The workflows started as an extensible markup language XML template and are modified to meet the requirements of the AIMS use cases. A set of common workflows out of the box were used for the identity lifecycle events.

To summarize the technologies used in the study, Table 3 organizes the technical components of the system categorized by operation systems, interfaces, virtual disk format, programming languages, and applications platform.

Table 3.

Inventory of Software Components, Interfaces, Languages, and Applications Used

Operating Systems	Interfaces	Virtualization Format	Languages	Applications
Windows 7	SOAP	VMDK	SPML2	Sun Identity Manager
Ubuntu 8.1	HTTPS	VMX	Xpress	NetBeans 6.1 IDE
Solaris 10/Zones		OVF	XML	ArchStudio 4
Amazon EC2			SysML	OVF Tool

Systems Interfaces

AIMS to CIMS Interface

The AIMS to CIMS connection represents one of the core interfaces in the investigation. This is where the trust fabric for identity assurance was sustained through the cloud. AIMS served as the record of authority and exchanged identity record information and updates with CIMS. The bi-directional interface was constructed through an encrypted tunnel with near-real time exchange of information.

AIMS to Active Directory Interface

Active Directory is an enterprise directory based on the Lightweight Directory Access Protocol (LDAP) and Active Directory Service Interface (ADSI). In the study, Active Directory was a foundation service that applications use to query account record data and directory objects. User and group objects were stored here as well as computing resources, service accounts, and network policies. AIMS communicated to the Active Directory instance using the AIMS adapter in a bi-directional fashion. AIMS remained the record of authority for identity data thereby taking precedence for data attributes. The connection was proxied through a Windows based gateway server required for AIMS to Active Directory connections. The gateway was a COTS

constraint, necessary due to the .NET COM libraries available only on the Windows operating system. AIMS controlled all communications with Active Directory and error handling was managed by AIMS in the form of reattempting failed connections periodically. The Active Directory schema fields are listed in Table 4.

Table 4.

Active Directory Fields

Field Name	Type	Length	Format	AIMS Field	Comments
sAMAccountName	String	0 - 256	Unicode String	ACCOUNT_ID	SAMAccount Name
HRemplD	String	No min or max defined	Unicode String	SUBJECT_ID	Human Resources Employee Identification Number
givenName	String	1 – 64	Unicode String	FIRST_NAME	First name of employee
middleName	String	0 – 64	Unicode String	MIDDLE_NAME	Middle name of employee
sn	String	1 – 64	Unicode String	LAST_NAME	Last name/Surname of employee
mail	String	0 – 256	Unicode String	EMAIL	Corresponds to RFC-2822 email address
uSNChanged	String	No min or max defined	Large Integer/ Interval	USNCHANGED	Incremented counter used by active sync to detect account creations and modifications
LevelofAssurance	String	1 – 1024	Unicode String	LOA_VALUE	Level of Assurance designation
userPrincipleName	String	Max 1024	Unicode String	USER_PRINCIPLE_NAME	User Principle Number
DN	String		Distinguished Name	DN	Distinguished Name

Table 4 *(continued)*

Field Name	Type	Length	Format	AIMS Field	Comments
CN	String	1 – 64	Unicode String	CN	Common Name
objectGuid	String	32	Octet String	OBJECTGUID	AIMS unique, persistent identifier

HR to AIMS Interface

The Human Resources (HR) to AIMS interface was a flat file interface. The attributes in the Table 5 represent the Human Resources system flat file (see Table 5). These attributes were used to establish the identity record in AIMS and provide credential data. The flat file was sent daily from the HR system and reflects changes in the subject's relationship with the organization. AIMS compared the data to see if any attributes have changed and took appropriate actions to update the changes.

Table 5.

Human Resources Data Fields

Field Name	Type	Length	Format	AIMS Field	Comments
NATIONAL_ID	String	20	Mixed case	NATIONAL_ID	National ID; must be alphanumeric
EMPL_ID	String	11	Num	SUBJECT_ID	HR Employee Identification
NAME_PREFIX	String	4	Mixed case	NAME_PREFIX	Miss, Mr., Mrs., Ms., Dr.
FIRST_NAME	String	30	Mixed case	FIRST_NAME	
MIDDLE_NAME	String	30	Mixed case	MIDDLE_NAME	
LAST_NAME	String	30	Mixed case	LAST_NAME	

Table 5 *(continued)*

Field Name	Type	Length	Format	AIMS Field	Comments
NAME_SUFFIX	String	15	Mixed case	NAME_SUFFIX	Jr., Sr., II, III, etc.
BUSINESS_DIV	String	5	Mixed case	BUSINESS_DIV	Business Division
WORK_ADDRESS	String	35	Mixed case	WORK_ADDRESS	Work Address Line 1
WORK_CITY	String	30	Mixed case	WORK_CITY	Work City
WORK_STATE	String	30	Mixed case	WORK_STATE	Work State
WORK_ZIP	String	12	Mixed case	WORK_ZIP	Work Zip Code
WORK_COUNTRY	String	2	Mixed case	WORK_COUNTRY	Work Country
DEPT_ID	String	10	Mixed case	DEPT_ID	Department Code
DEPTID_DESC	String	30	Mixed case	DEPTID_DESC	Department Code Description
MGR_ID	String	8		MGR_ID	Manager's Employee Identifier

Table 5 *(continued)*

Field Name	Type	Length	Format	AIMS Field	Comments
JOB_ID	String	6	Mixed case	JOB_ID	Job Description Identifier
JOB_DESC	String	30	Mixed case	JOB_DESC	Description of JOB_ID
WORK_PHONE	String	24	Mixed case	WORK_PHONE	Work phone number as provided by employee
sAMAccountName	String	0 - 256	Unicode String		ACCOUNT_ID (if available)
EMAIL	String	70	Mixed case	EMAIL	RFC-822 address (email standard)
LOA_VALUE	String	1 – 1024	Unicode String	LOA_VALUE	Level of Assurance designation

Languages Used

A number of programming languages were used to build the environment on which AIMS and CIMS run. The languages were workflow processes and extensible markup languages. Generally, markup languages are customized for a specific need by using extensible attributes and definitions. For the purposes of this study, SysML, XML, SPML2, and XPRESS languages were used.

SysML is a graphical modeling language and notation tool based on the UML2 profile commonly used in the systems engineering discipline. SysML was designed as a visual

modeling language to provide semantics and notation by helping systems engineers and systems architects to identify and solve problems. SysML can be used with multiple design methodologies including agile and model based design approaches. The model and system artifacts are tool independent, meaning that any SysML-compliant tool can be used to visualize, analyze, verify, validate, and verify the design. For this investigation, three tools were used to generate SysML: ArchStudio4, Topcased, and Microsoft Visio.

XPRESS is an XML-based functional language used by Sun's Identity Manager product to collect, transform, and present identity related data. The AIMS workflow logic and forms were based on XPRESS. The AIMS administration pages presented in the user interface (UI) were supported by XPRESS syntax embedded in XML tags. The attributes collected on the form were transformed and stored in the AIMS database repository.

Services Provisioning Markup Language version 2.0 (SPML2) was used by AIMS to broker provisioning requests. The AIMS SPML2 listener service listened for any incoming requests that contain expressions to process. The expression types were based on the use case scenarios: create, update, read, delete, and search. The SPML2 packages were contained in a Simple Object Access Protocol (SOAP) protocol envelope. The SOAP envelope contained schema information, attribute, and header information such as where to send the SPML2 message. Each message contained a Provisioning Service Object (PSO) considered a target and a corresponding PsoID (PSO Identifier) for each object. To summarize, SPML2 code was the payload and SOAP was the payload carrier that delivers the message to an SPML2 listener service. Once the message was delivered, AIMS unpacked the payload and processed the request.

Design Tools

ArchStudio4 was used to design the initial visual concept. The principal investigator deployed ArchStudio into the Eclipse 3.6.1 platform, but soon found limitations in terms of support for the tool. The investigator had trouble locating practitioners with experience to assist in the use of ArchStudio; many were unfamiliar or had never heard of the application. The Object Modeling Group (OMG) maintains a list of vendors that support SysML modeling tools including IBM Rhapsody, Papyrus, Sparx Systems' Enterprise Architect, Microsoft's Visio with SysML templates, and No Magic's MagicDraw SysML. The investigator found a combination of Visio templates and a new open source tool called Topcased to be a working combination.

Visio is Microsoft's visualization and modeling tool. Network and systems engineers may recognize its capabilities for process diagrams, network diagrams, ability to depict interfaces, connections, and data flows. Based on the OMG SysML specification, Pavel Hruby, with help from Lockheed Martin Corporation's Sandy Friedenthal, created a set of stencils for use in Visio. The stencils are based on the SysML 1.0 standards and are easily imported into Visio as a collection of objects. The stencils include activity diagrams, block diagrams, packages, parametrics, profiles, requirements, sequence diagrams, state diagrams, and use cases.

Topcased is an open source modeling and design tool managing out of France. The tool was bundled with Eclipse 3.5 (Galileo) in a single download. Whereas some versions of Eclipse, Java, and the processor architecture type (32-bit versus 64-bit) caused compatibility problems, the Topcased packaged was easily installed and configured on the Windows 7 64-bit modeling workstation. Topcased had readily available tutorials, advanced model exemplars, and a quick-start guide to get the investigator up to speed in short order. Figure 16 shows a screenshot of the integrated development environment (IDE) within Eclipse with the Topcased user interface.

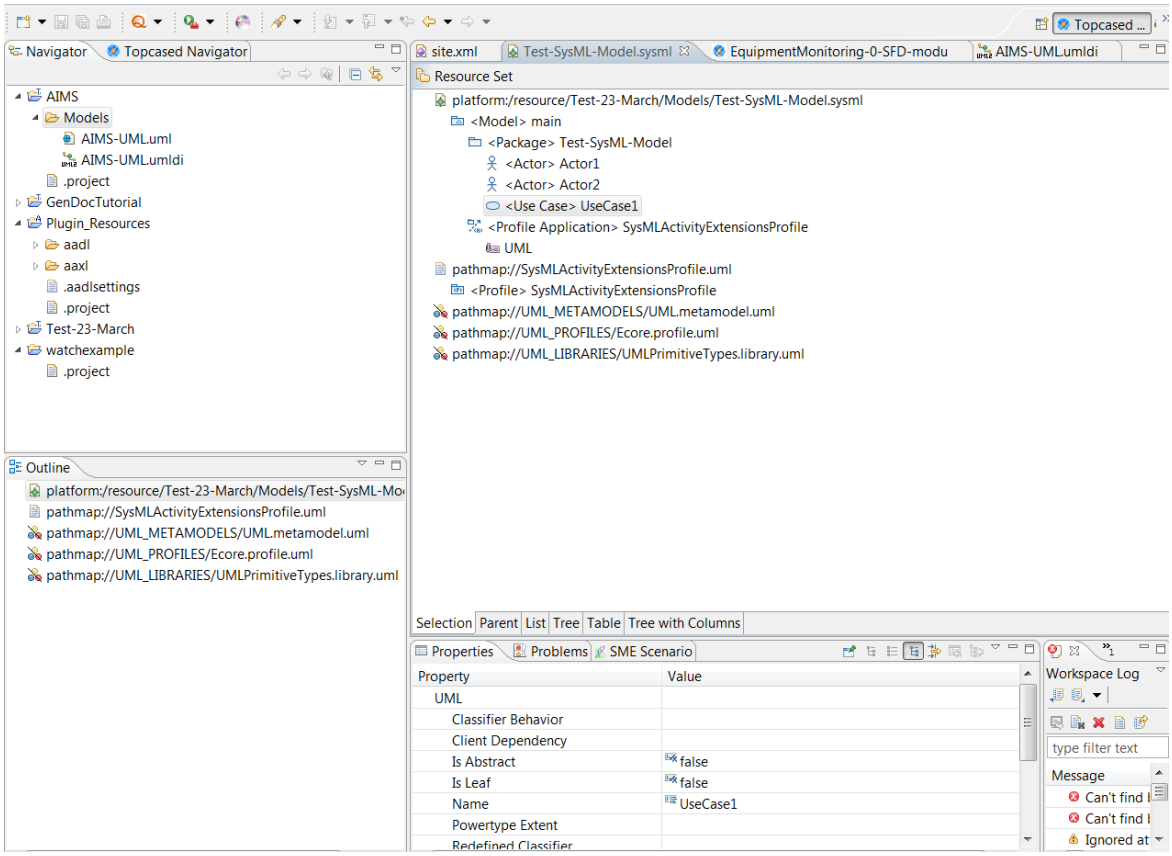

Figure 16.

Topcased User Interface Deployed in Eclipse 3.5

In the end, all three modeling tools were used to build the use cases and analyze the relationships between actors and components. ArchStudio 4 uniquely offered the xADL 2.0 architecture description language. Visio offered a comfort level in terms of ease of use and familiarity by the investigator. Visio also included special stencils created and reviewed by leading researchers in the SysML community. Finally, Topcased was the most mature of the three modeling tools in terms of features and support. Examples, tutorials, and a pre-configured installation made Topcased easy to work in.

Use Case Design

In order to answer the research questions posed in this study, a series of ten fully dressed use case scenarios were developed. The use cases were tailored to meet the study's objectives. They are based on feedback, collaboration, and related process work from a variety of process teams, working groups, and standards bodies. The use case template is unique compared to previous use case formats; prior work did not contain cloud specific references. One of the investigations contributions to the body of knowledge, these new use case templates allow designers to understand the cloud deployment model and cloud service model with greater clarity. One of the design goals is that these use cases represent an extensible, reusable format thereby allowing practitioners and researchers to build upon the core foundation of work in this investigation.

The use case design included a number of attributes and corresponding values for each scenario. The field names and descriptions are listed in Table 6. The use cases were fully-dressed with detailed information relating to a number of technical and business functions. One important customization for this investigation is the addition of fields for Cloud Service Model and Cloud Deployment Model. These two items have never been added to a fully-dressed use case, however they are vital components when dealing with cloud systems. The service model and deployment model helped to frame the context of the scenario. For example, if the service model is within the Platform-as-a-Service, the analyst reading the use case immediately has a frame of reference, i.e. the process will be conducted within the platform layer. Similarly, the deployment model designation allows the analyst to understand how the environment has been implemented, specifically with connectivity, security, and data exchange among cloud environments. The remaining fields defined actors, triggers, alternate courses, policy impacts,

constraints, and open issues. Details on each field varied with the amount of information available.

Table 6.

Use Case Description

Use Case Field Name	Description
Use Case Name	Title of the use case
Use Case Id	Number of use case, for reference purposes
Priority	Priority for testing or program management usage
Source	Program name, source of requirement, for tracking purposes
Primary Business Actor	Primary business actor who will be performing activity
Primary System Actor	Primary system actor, technical or analyst subject
Other Participating Actors	Related actors who may have minor role or role other than primary in the scenario
Other Interested Stakeholders	Stakeholders with interest in the scenario; business unit, program sponsor, customer, advocate, shareholder, etc
Cloud Deployment Model	One of four deployment models used in the scenario
Cloud Service Model	Defined service models with ability to define custom, such as Identity-as-a-Service, Storage-as-a-Service, etc
Description	Summary description of the scenario, used to set the stage for the following steps
Pre-Condition	Pre-requisites necessary to be done prior to executing the use case
Trigger	The action or process that begins the use case scenario; Can be automated or manual
Typical Course Of Events	Logical flow of events with actor actions and system response as designed
Alternate Courses	Alternate course defined as another path through the use case if an action or response in the typical course fails
Conclusion	Summarizes when use case concludes
Post-Condition	Summarizes the expected outcome as a result of executing the use case
Business Rules	Identifies related business rules, include job roles, system status, or business functions associated with the use case

Table 6 *(continued)*

Use Case Field Name	Description
Policy Impacts	Identifies business, technical, security policies affected by the scenario
Implementation Constraints And Specifications Assumptions	Known constraints or assumptions related to the scenario or systems components are documented here
Open Issues	Open issues such as to-be-determined items, discrepancies, or deficiencies captured
Notes/Use Case Diagram	Follow-up notes, diagrams, schematics, and similar information that may help analysts with the use case

Sources for Use Cases

A number of sources were used for the structure and format of the use cases for this investigation. The NIST definitions were adopted for the taxonomy and lexicon by the study including delivery models, deployment models, and characteristics. Research by Mell and Grance from NIST is emerging as a commonly referenced source in the literature. Despite the evolving nature of the literature, many of the use case committees, working groups, and standards bodies defer to the NIST definition as a launch point.

Standards-focused organizations that are currently working on cloud computing use cases include:

- The Cloud Computing Use Case Discussion Group has produced 4 versions of the Cloud Computing Uses Cases White Paper (2010).
- NIST Standards Acceleration Jumpstarting Adoption of Cloud Computing (SAJACC)
- Kantara Initiative Consumer Identity and Federation Interoperability Groups
- Oasis Cloud-ID TC
- [SNIA] "Cloud Storage Use Cases", Storage Network Industry Association, Version 0.5 rev 0, June 8, 2009.

The investigation reviewed over 90 use cases in draft format, working format, pending approval, and approved from the five standards-focused technical organizations. Two of the organizations were focused exclusively on identity management within cloud computing. The use cases found in these groups were generally related to lifecycle management, provisioning of access, auditing, accessibility, and identity federation. The third group included use case development for all areas of cloud computing including access provisioning, performance bursting, virtual machine management, and similar infrastructure scenarios. The fourth group built use cases under the broad umbrella of cloud security, which includes identity management as well as non-identity management security use cases. The final group was focused exclusively on storage within the cloud. None of the use cases in the last group included identity management, lifecycle management, nor provisioning. Most of the use cases found in the final group dealt with data retention, data backup, data management and data policy issues.

The data was collected by contacting cloud computing focused working groups within standards bodies. In each case, the use cases and group work products were posted online in a collaborative fashion. The data was easily accessible and openly shared with committee members and the general public. In one case the group used Google discussion forums to collaborate; another used a knowledge-based application called Confluence to share documents and provide a wiki tool.

The input was consolidated into a matrix for analysis. Of the 90 use cases collected, 34 were found to be related to the core functions of identity management: lifecycle management, provisioning, access, auditing, authorization, and authentication. Further analysis found that of the 34 use cases, none were in a fully dressed use case format. By design, most of the use cases, were business case or concept level. It should be noted a number of these use cases were in

working mode and draft format. The use cases in this study built upon the work of previous researchers and their working artifacts.

Taxonomy

The SysML taxonomy was used to build and model the use cases. The taxonomy carries over many of the structures found in the Unified Modeling Language (UML) and builds upon existing work in the UML specification. The requirements diagram is one of the new features introduced with SysML (see Figure 17). Other diagrams including the activity diagram were modified and adopted with the new SysML standard. Figure 17 also shows the taxonomy for SysML used in this study.

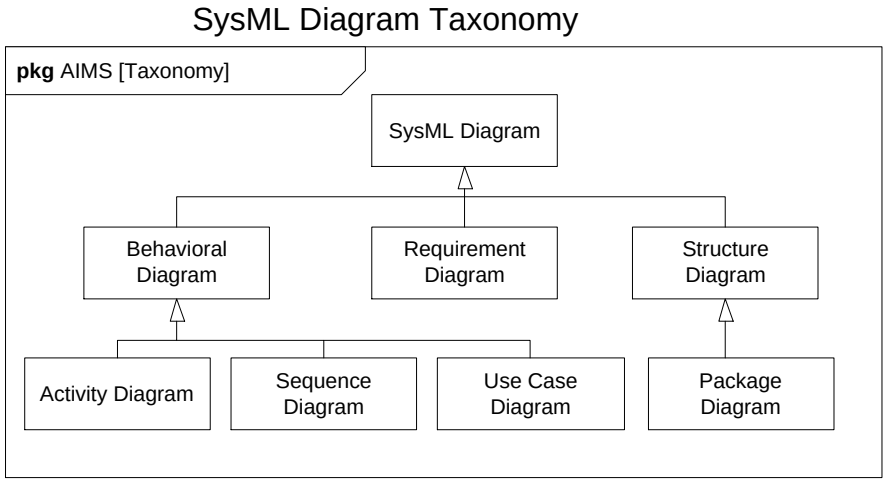

Figure 17.

AIMS SysML Taxonomy

This investigation used logical structures given the emphasis on virtualization. The work products delivered for this study include the high level conceptual operational view (OV-1) and a package diagram under the structure diagram type (bottom right box in the figure above). For behavioral diagrams, this investigation produced activity diagrams, sequence diagrams, use cases

diagrams, and package diagrams. The introduction of the cloud platform presents a change in the behavior of how computing resources are used; specifically in the identity assurance realm with secure transmission of account and identity data between the clouds.

SysML diagrams are easily identified by the header label with information contained in the pentagon shaped shape found in the top left corner of the artifact. Within the pentagon object there are five descriptors; the first three are required and the last two are optional. The first component of the label written in bold font indicates the diagram kind such as activity, sequence, block diagram, etc (see Figure 18). The next descriptor is the model element type such as activity, block, or other. The third component is the model element name. The final two optional descriptors are the diagram name and the diagram usage. The frame sets a boundary for the diagram and all content is contained within the structure frame.

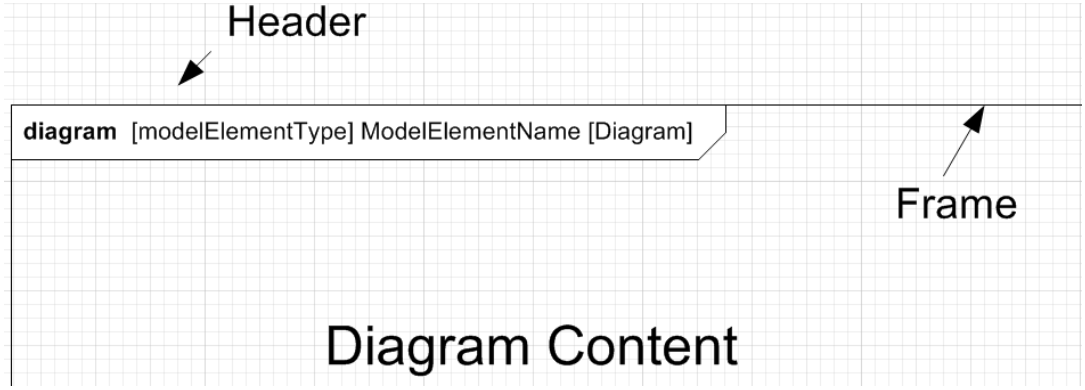

Figure 18.

SysML Header Information

Breaking down the header syntax further, some examples of each attribute are provided. In the book *A practical guide to SysML: Systems Model Language* authors Friedenthal et al.

define the construct of the header. Based on these definitions, the following values were used in the study:

The *diagram kind* with abbreviations include:

- Activity diagram – **act**
- Package diagram **– pkg**
- Requirement diagram **– req**
- Use case diagram **- uc**

The *model element type* references used in the naming convention include:

- Activity diagram – used to show control operations or scenarios
- Package diagram – could be a work package, model, library, profile, or a view
- Requirement diagram – shows requirements and relationships within system
- Use case diagram – could be a package, model, or library with respect to a scenario

The *model element name* reference is customizable and typically used to increase the clarity of the artifact while preventing potential confusion with similar work packages in the model.

These remaining two attributes are *diagram name* and *diagram usage*. While optional, these two provide a description of the diagram and how it should be used. For example, later in the study, sequence diagrams were used as test cases to validate requirements using the following convention: "**sd** Secure Identification of User [Test]". The diagram usage descriptor indicates how the particular sequence diagram was intended to be used: as a test case.

Use Cases

Use cases were described with the SysML designation **uc**. The operational use case view depicts the ten scenarios covered in this investigation. The view shows how the actors interacted with the systems, organized by vertical swim lanes representing each actor. Each use case was

labeled in an oval and a number of the use cases are dependent on others give that certain data objects or actions are included, denoted by the «*include*» syntax (see Figure 19).

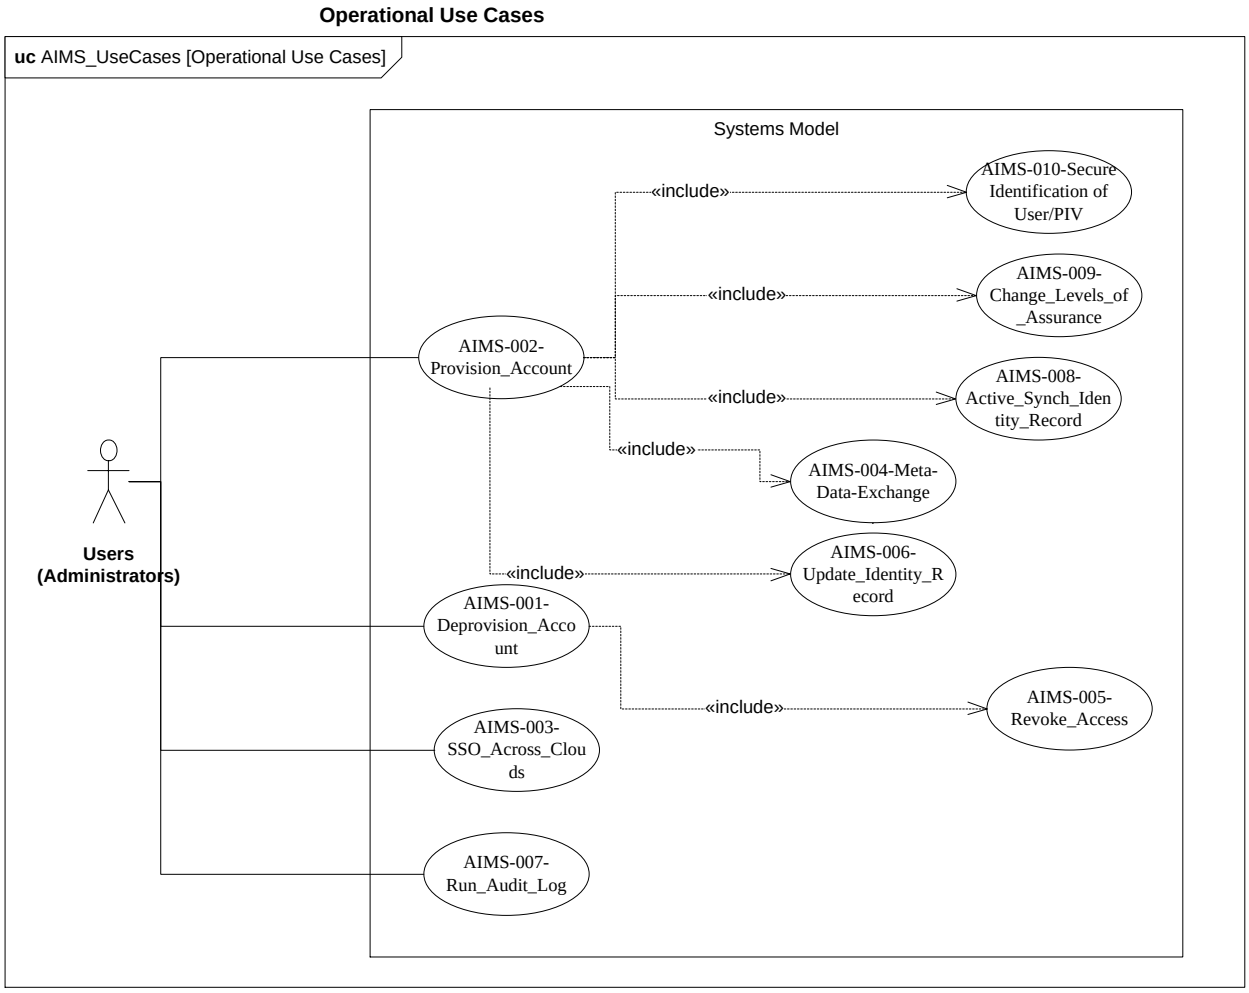

Figure 19.

Operational Use Case View, SysML

The use cases created for this study have a foundation in the 34 identity management and assurance related use cases started by industry working groups. As a result, ten scenarios were chosen based on common functions between tools capability with the SPML2 specification from OASIS and the SPML2WavesetOperations JAVA class provided by Sun Identity Manager, and

regularly used operations for managing identities. The use cases are listed in Table 7: Use Case Scenarios.

Table 7.

Use Case Scenarios

Use Case Title	Use Case Number	Use Case Description
AIMS-001-Deprovision_Account	AIMS-001	Deprovision Account
AIMS-002-Provision_Account	AIMS-002	Provision Account
AIMS-003-SSO_Across_Clouds	AIMS-003	Single Sign On to Cloud Environment
AIMS-004-Meta-Data-Exchange	AIMS-004	Exchange of Meta-Data attributes
AIMS-005-Revoke_Access	AIMS-005	Revocation of access
AIMS-006-Update_Identity_Record	AIMS-006	Update existing identity record
AIMS-007-Run_Audit_Log	AIMS-007	Audit log
AIMS-008-Active_Synch_Identity_Record	AIMS-008	Synchronization of Identity Record
AIMS-009-Change_Levels_of_Assurance	AIMS-009	Change Levels of Assurance
AIMS-010-Secure Identification of User/PIV	AIMS-010	Secure Identification of End-User

The fully dressed use cases are listed below. The cloud deployment model and the cloud service model are special features for the purposes of this investigation. The "system response" column under the typical course of events primarily focuses on AIMS. Given that AIMS was the record of authority for identity information, the end users and administrators primarily use the AIMS user interface to manage account lifecycle events. For the purposes of this investigation, the technical reference model was the primary focus; therefore assumptions were made for business and policy impacts.

ASSURED IDENTITY MANAGEMENT SYSTEM
USE CASE PACKAGE AIMS-001

Author (s): **Daniels** **Date:** **07-February-2011**
Version: **1.0**

USE CASE NAME:	Deprovision Account	**USE CASE TYPE & LEVEL**
USE CASE ID:	AIMS-001, version 1	Business:
PRIORITY:	High	System/Solution:
		Requirements ☑
SOURCE:	*Research Project: Assured Identity for Cloud Computing (2011, Daniels)*	Analysis
PRIMARY BUSINESS ACTOR:	AIMS	Design
		Fully Dressed ☑
PRIMARY SYSTEM ACTOR:	AIMS Technical Team	
OTHER PARTICIPATING ACTORS:	Cloud Identity System	
OTHER INTERESTED STAKEHOLDERS:		
CLOUD DEPLOYMENT MODEL:	☐ Public ☐ Private ☐ Community ☑ Hybrid	
CLOUD SERVICE MODEL:	☐ Software-as-a-Service (SaaS) ☑ Platform-as-a-Service (PaaS) ☐ Infrastructure-as-a-Service (IaaS) ☑ Other (example: Identity-as-a-Service)	
DESCRIPTION:	The use case is initiated when AIMS receives a deprovisioning request upon a notification from human resources, industrial security, or other authorized party. This workflow is triggered based on a request by the AIMS Administrator.	
PRE-CONDITION:	AIMS system must be able to issue a Deprovision function with appropriate authorization. The deprovision action must be issued to the downstream integrated Cloud Identity System. A record of the resource being enrolled must exist in the AIMS database for auditing	
TRIGGER:	The use case is initiated when AIMS receives a deprovisioning request upon a notification from human resources, industrial security, or other authorized party. This workflow is triggered based on a request by the	

USE CASE NAME:	Deprovision Account	USE CASE TYPE & LEVEL
	AIMS administrator.via a manual process or when an account deprovisioning request has been requested.	
TYPICAL COURSE OF EVENTS:	**Actor Action**	**System Response**
	Step 1: The AIMS Administrator navigates to the appropriate location and navigates to a "Find User" form.	**Step 2**: AIMS displays a blank "Find User" form.
	Step 3: The AIMS administrator types in the Credentialed Entity's Employee ID, SamAccountName, Last Name, and/or First Name and hits Enter. **(Alt Step 3)**	**Step 4**: AIMS pre-populates the Name Prefix, First Name, Preferred First Name, Middle Name, Last Name, Name Suffix, Credentialed Entity Company, Credentialed Entity Type , EmailID, Owner SamAccountName, OwnerAccountID, user Principal Name, DN, CN, Export Control Designation, Credential Type, and Location Fields from Active Directory and the human resources database table. **(Alt Step 4a) (Alt Step 4b)**
	Step 5: The AIMS administrator requests to deprovision the appropriate Credentialed Entity record and hits the Save button.	**Step 6:** AIMS validates that all credentials associated with the credentialed entity have been cancelled.
		Step 7: AIMS sends deprovision packet to Cloud Identity System. **(Alt Step 7)**
	Step 8: Cloud Identity System sends acknowledgement.	**Step 9:** AIMS logs acknowledgement and use case concludes.
ALTERNATE COURSES:	**ALT Step 3:** The AIMS Administrator types in the Account ID (especially for Non-Employees, though it applies to both Employees and Non-Employees) and hits Enter.	
	Alt Step 4a: AIMS is unable to find Credentialed Entity data from human resources or Active Directory. Error message is displayed.	
	Alt Step 4b: If multiple results that meet the criteria are found, AIMS displays a list of all users that meet that criteria. AIMS administrator selects appropriate user from list.	
	Alt. Step 7: AIMS is unable to establish a connection with Cloud Identity System to send data. AIMS retries every 2 minutes for 40 tries. Following the first unsuccessful retry an error message is sent to an agreed upon distribution list.	

USE CASE NAME:	Deprovision Account	USE CASE TYPE & LEVEL
CONCLUSION:	The use case concludes when AIMS commits the de-provisioning updates to the provisioning database.	
POST-CONDITION:	AIMS sends a deprovisioning request to Cloud Identity System.	
BUSINESS RULES:	The AIMS administrator is responsible for approving all deprovisioning requests before the AIMS Technical Team deprovisions the requested resources. JOBSTATUS equal to “Deprovision_Requested”	
POLICY IMPACTS:	NA	
IMPLEMENTATION CONSTRAINTS AND SPECIFICATIONS		
ASSUMPTIONS:		
OPEN ISSUES:		
NOTES/USE CASE DIAGRAM:	See Figure 20	

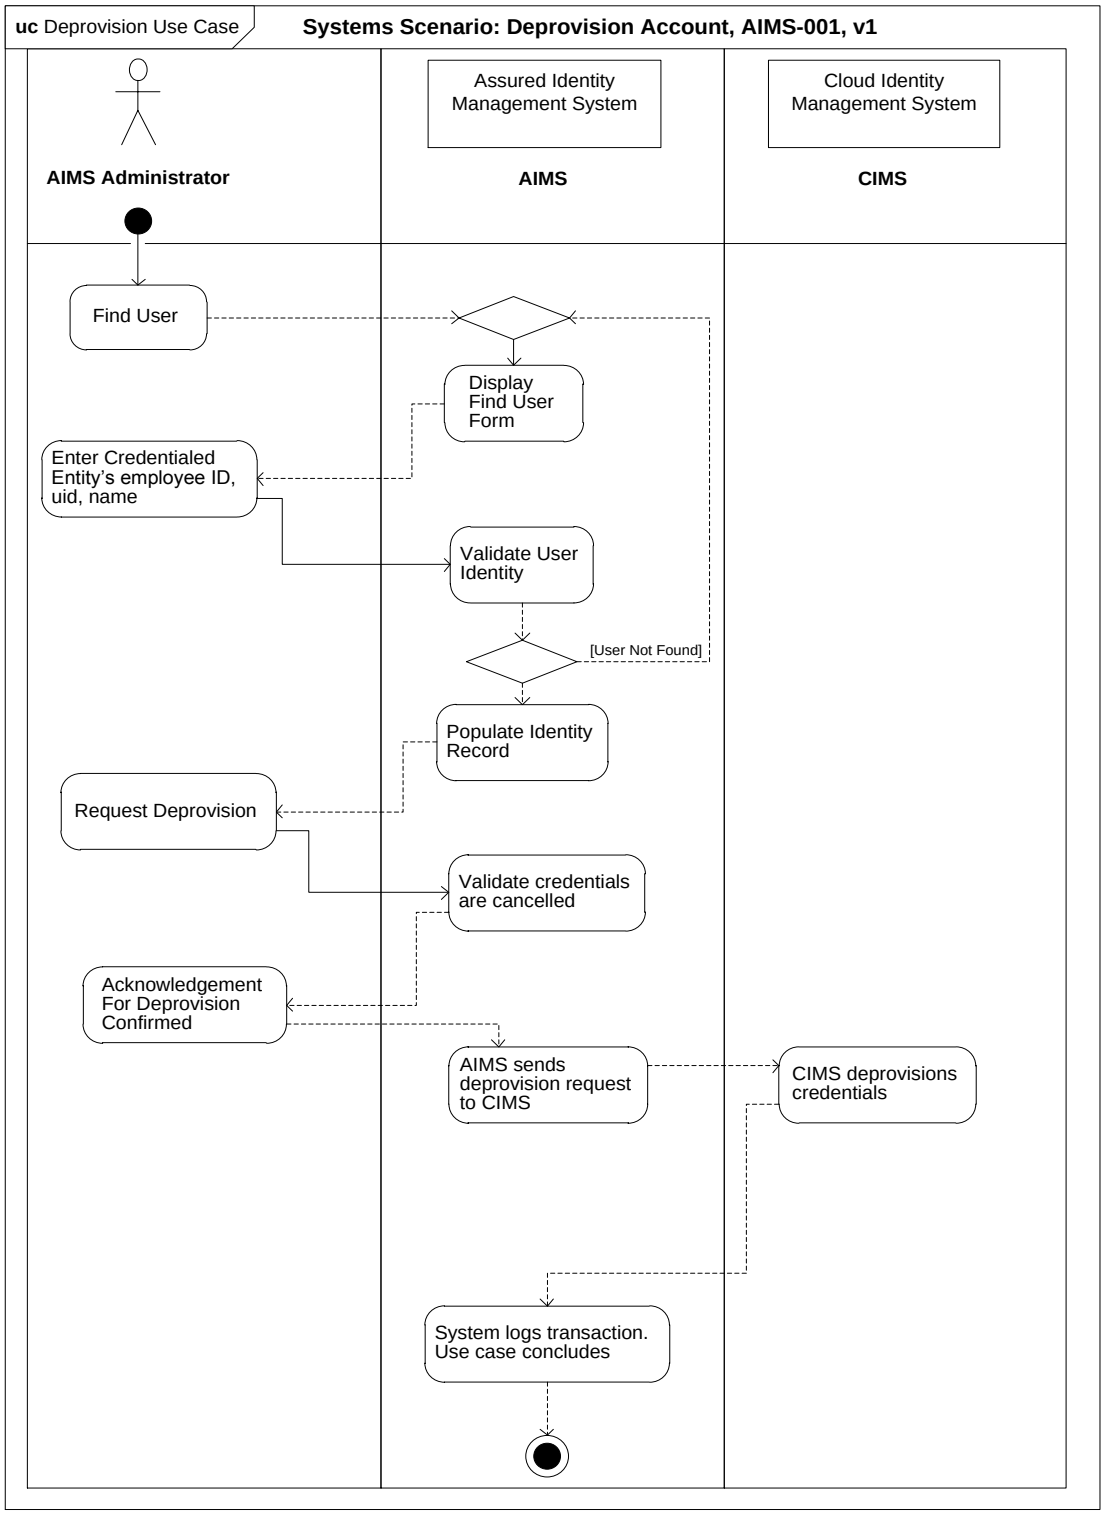

Figure 20.

Deprovisioning Use Case, AIMS-001

ASSURED IDENTITY MANAGEMENT SYSTEM
USE CASE PACKAGE AIMS-002

Author (s): **Daniels** **Date: 07-February-2011**
Version: 1.0

USE CASE NAME:	Provision Account	**USE CASE TYPE & LEVEL**
USE CASE ID:	AIMS-002, version 1	Business:
PRIORITY:	High	System/Solution:
		Requirements ☑
SOURCE:	*Research Project: Assured Identity for Cloud Computing (2011, Daniels)*	Analysis
PRIMARY BUSINESS ACTOR:	AIMS	Design
		Fully Dressed ☑
PRIMARY SYSTEM ACTOR:	AIMS Technical Team	
OTHER PARTICIPATING ACTORS:	Cloud Identity System	
OTHER INTERESTED STAKEHOLDERS:		
CLOUD DEPLOYMENT MODEL:	☐ Public ☐ Private ☐ Community ☑ Hybrid	
CLOUD SERVICE MODEL:	☐ Software-as-a-Service (SaaS) ☑ Platform-as-a-Service (PaaS) ☐ Infrastructure-as-a-Service (IaaS) ☑ Other (example: Identity-as-a-Service)	
DESCRIPTION:	The use case is initiated when AIMS receives a provisioning request upon a notification from human resources, industrial security, or other authorized party. This workflow is triggered based on a request by the AIMS Administrator.	
PRE-CONDITION:	AIMS system must be able to issue a provision function with appropriate authorization. The provision action must be issued to the downstream integrated Cloud Identity System.	
TRIGGER:	The use case is initiated when AIMS receives a provisioning request upon a notification from human resources, industrial security, or other authorized party. This workflow is triggered based on a request by the AIMS administrator.via a manual process or when an account	

	provisioning request has been submitted.	
TYPICAL COURSE OF EVENTS:	**Actor Action**	**System Response**
	Step 1: The AIMS Administrator verifies identity of user	**Step 2**: None
	Step 3: The AIMS administrator types in the Credentialed Entity's Employee ID, SamAccountName, Last Name, and/or First Name and hits Enter. **(Alt Step 3)**	**Step 4**: AIMS pre-populates the Name Prefix, First Name, Preferred First Name, Middle Name, Last Name, Name Suffix, Credentialed Entity Company, Credentialed Entity Type , EmailID, Owner SamAccountName, OwnerAccountID, user Principal Name, DN, CN, Export Control Designation, Credential Type, and Location Fields from Active Directory and the human resources database table. **(Alt Step 4a) (Alt Step 4b)**
	Step 5: The AIMS administrator requests to provision the appropriate Credentialed Entity identity record and hits the Save button.	**Step 6:** AIMS validates the credentialed entity has been created.
		Step 7: AIMS sends provision packet to Cloud Identity System. **(Alt Step 7)**
	Step 8: Cloud Identity System sends acknowledgement.	**Step 9:** AIMS logs acknowledgement and use case concludes.
ALTERNATE COURSES:	**Alt Step 3:** The AIMS Administrator types in the Account ID. If no identity record is pre-populated, the AIMS administrator has the option to manually create the record.	
	Alt Step 4a: AIMS is unable to find Credentialed Entity data from human resources or Active Directory. Error message is displayed.	
	Alt Step 4b: If multiple results that meet the criteria are found, AIMS displays a list of all users that meet that criterion. AIMS administrator selects appropriate user from list.	
	Alt. Step 7: AIMS is unable to establish a connection with Cloud Identity System to send data. AIMS retries every 2 minutes for 40 tries.	
CONCLUSION:	The use case concludes when AIMS commits the provisioning updates to the provisioning database and CIS acknowledges receipt of the new request.	

POST-CONDITION:	AIMS sends a provisioning request to Cloud Identity System.
BUSINESS RULES:	The AIMS administrator is responsible for approving all provisioning requests before the AIMS Technical Team provisions the requested resources. JOBSTATUS equal to “provision_Requested”
POLICY IMPACTS:	NA
IMPLEMENTATION CONSTRAINTS AND SPECIFICATIONS	
ASSUMPTIONS:	
OPEN ISSUES:	
NOTES/USE CASE DIAGRAM:	See Figure 21

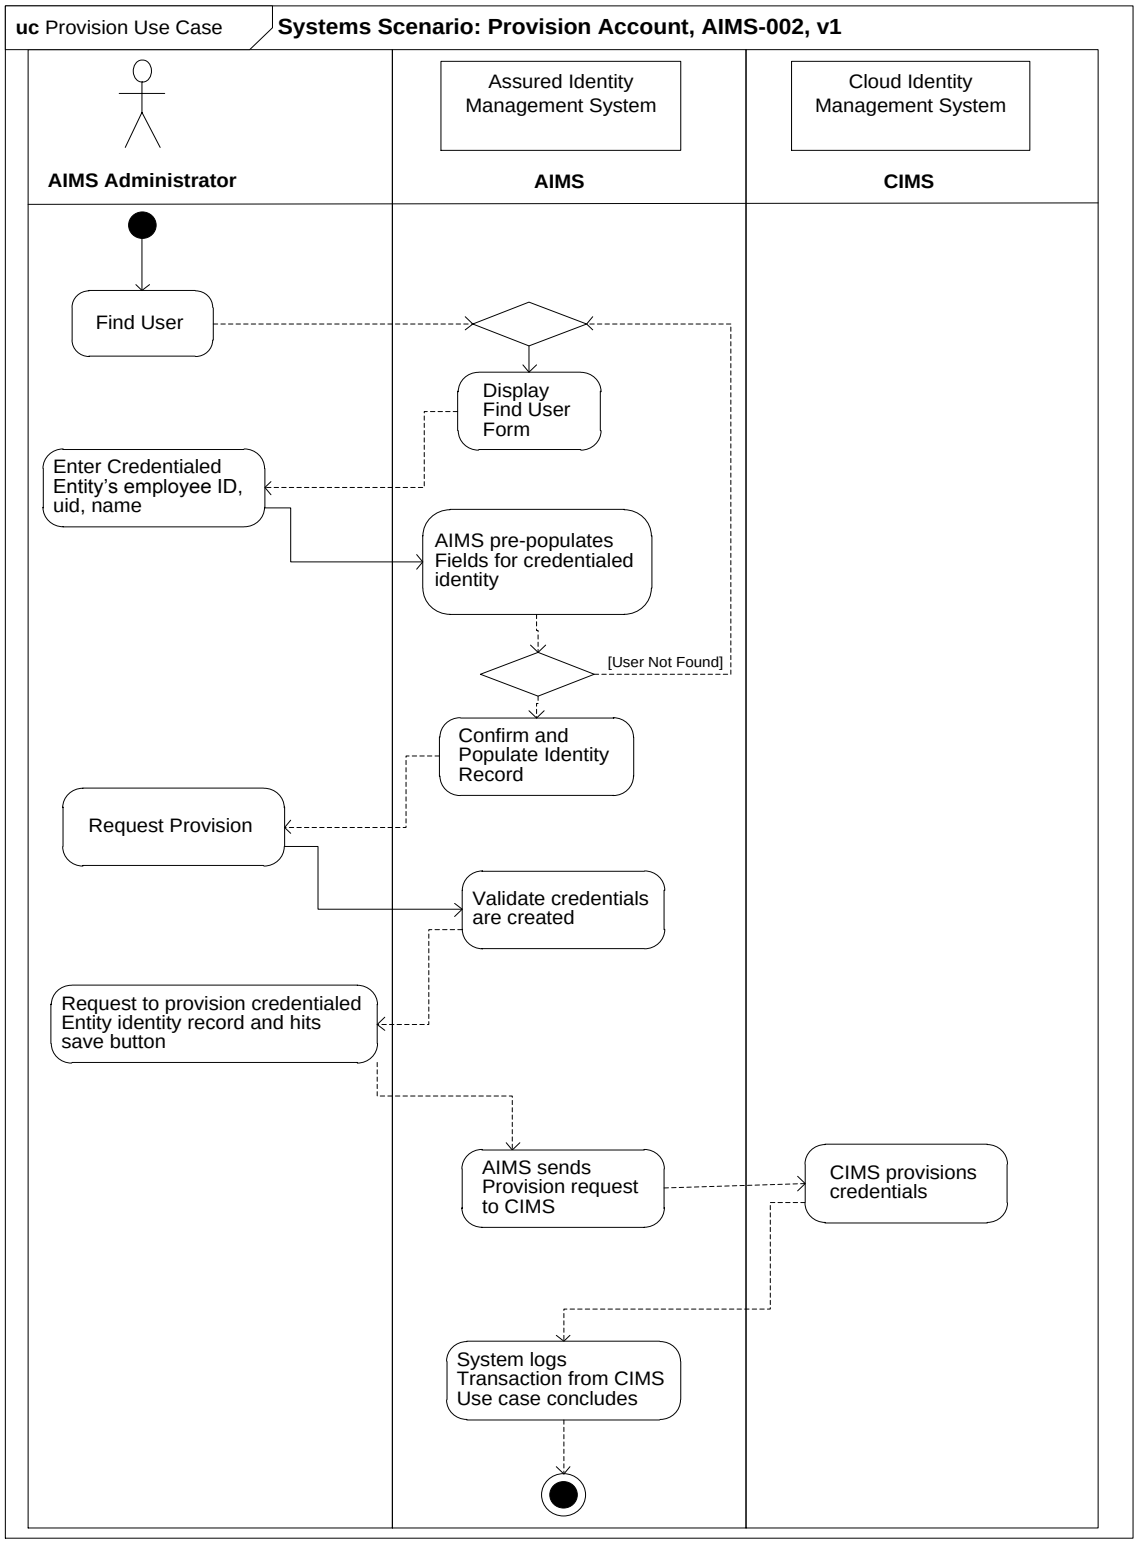

Figure 21.

Provisioning Use Case, AIMS-002

ASSURED IDENTITY MANAGEMENT SYSTEM
USE CASE PACKAGE AIMS-003

Author (s): **Daniels** **Date:** **07-February-2011**
Version: **1.0**

USE CASE NAME:	Single-Sign-On Across the Cloud	**USE CASE TYPE & LEVEL**
USE CASE ID:	AIMS-003, version 1	Business:
PRIORITY:	High	System/Solution: Requirements ☑
SOURCE:	*Research Project: Assured Identity for Cloud Computing (2011, Daniels)*	Analysis
PRIMARY BUSINESS ACTOR:	AIMS	Design Fully Dressed ☑
PRIMARY SYSTEM ACTOR:	AIMS Technical Team	
OTHER PARTICIPATING ACTORS:	Cloud Identity System	
OTHER INTERESTED STAKEHOLDERS:		
CLOUD DEPLOYMENT MODEL:	☐ Public ☐ Private ☐ Community ☑ Hybrid	
CLOUD SERVICE MODEL:	☐ Software-as-a-Service (SaaS) ☑ Platform-as-a-Service (PaaS) ☐ Infrastructure-as-a-Service (IaaS) ☑ Other (example: Identity-as-a-Service)	
DESCRIPTION:	The use case is initiated when user authenticates to AIMS, receives successful authorization to proceed. Within the same session, the user authenticates to CIS without being prompted to present credentials again.	
PRE-CONDITION:	User identity record must exist in AIMS and CIS identity repositories. User must be authorized in AIMS as well as CIS. Single-Sign-On occurs during a single session. AIMS and CIS use the same credential for authentication.	
TRIGGER:	The use case is initiated when user authenticates to AIMS and creates an authorized session.	
TYPICAL COURSE OF EVENTS:	**Actor Action**	**System Response**
	Step 1: The user authenticates to AIMS application.	**Step 2**: AIMS successfully authenticates user and creates session.

	Step 3: The user authenticates to CIS application.	**Step 4**: CIS successfully authenticates user and creates session.
	Step 5: Use case concludes.	
ALTERNATE COURSES:	N/A	
CONCLUSION:	The use case concludes after the user successfully authenticates to the CIS application, thereby confirming single-sign-on functionality.	
POST-CONDITION:	None	
BUSINESS RULES:	Single-sign-on functionality is commonly requested among enterprise applications	
POLICY IMPACTS:	NA	
IMPLEMENTATION CONSTRAINTS AND SPECIFICATIONS	Single-sign-on functionality is dependent on similar levels of assurance. The user identity within CIS and AIMS must be vetted at the same level to ensure single-sign-on.	
ASSUMPTIONS:		
OPEN ISSUES:		
NOTES/USE CASE DIAGRAM:	See Figure 22	

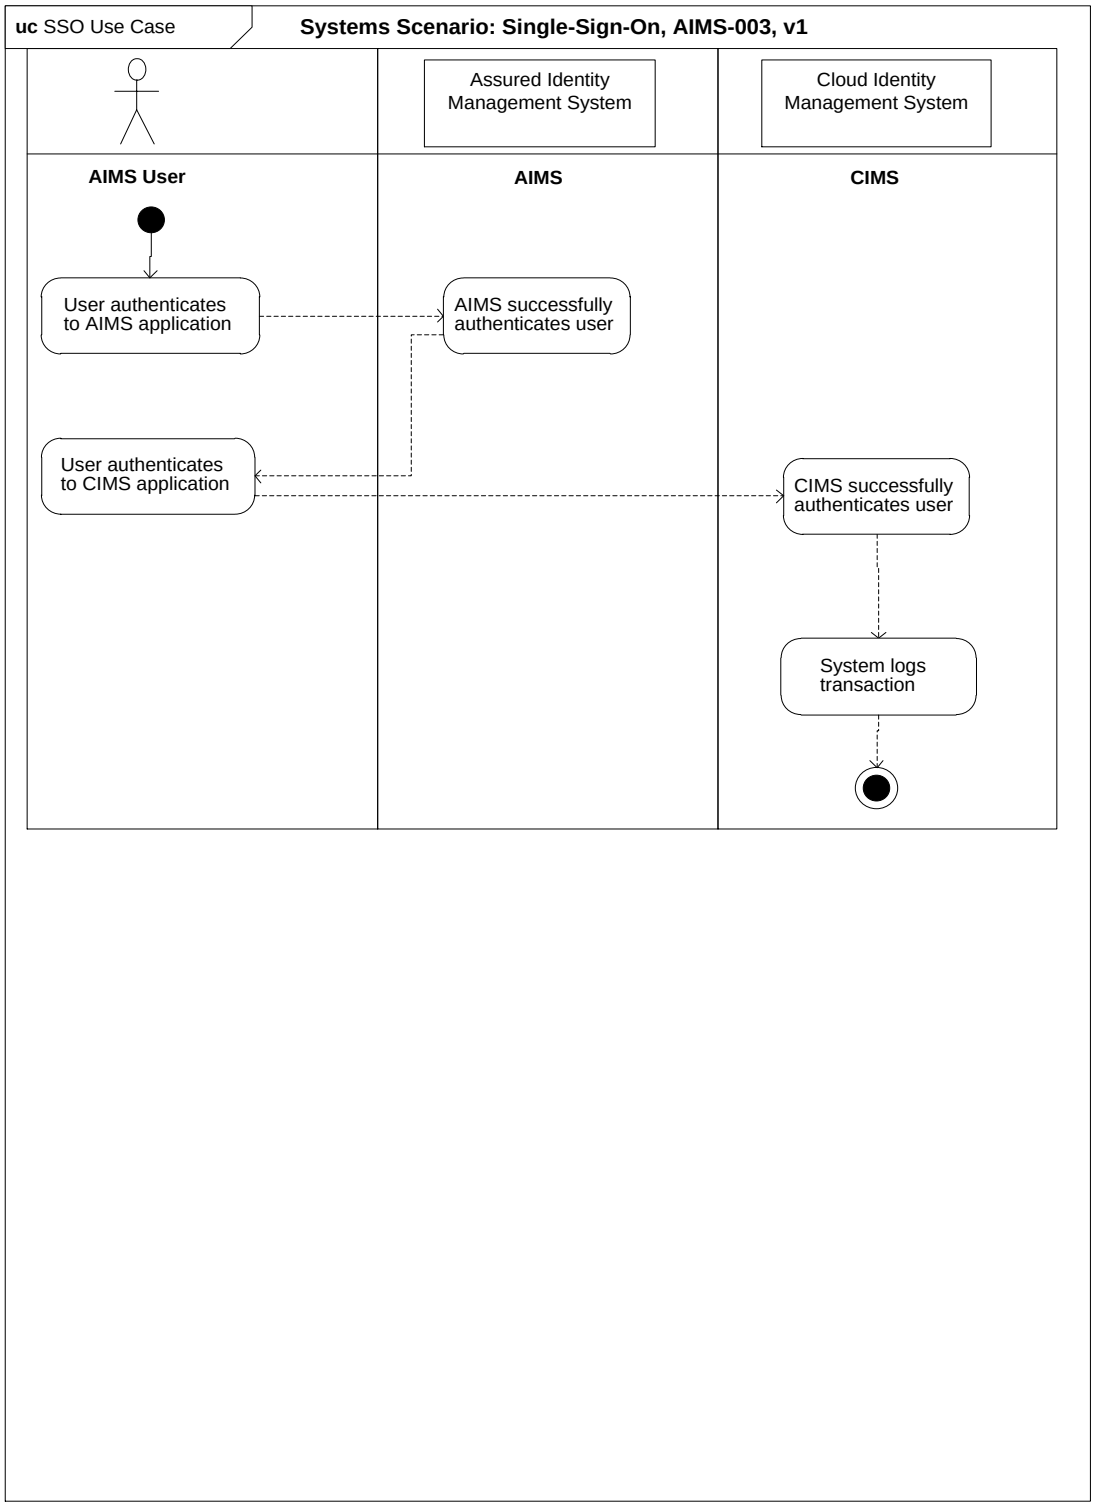

Figure 22.

Single Sign-On Use Case, AIMS-003

ASSURED IDENTITY MANAGEMENT SYSTEM
USE CASE PACKAGE AIMS-004

Author (s): **Daniels** **Date:** **07-February-2011**
Version: **1.0**

USE CASE NAME:	Meta Data Exchange	**USE CASE TYPE & LEVEL**
USE CASE ID:	AIMS-004, version 1	Business:
PRIORITY:	High	System/Solution: Requirements ☑
SOURCE:	*Research Project: Assured Identity for Cloud Computing (2011, Daniels)*	Analysis
PRIMARY BUSINESS ACTOR:	AIMS	Design Fully Dressed ☑
PRIMARY SYSTEM ACTOR:	AIMS Technical Team	
OTHER PARTICIPATING ACTORS:	Cloud Identity System	
OTHER INTERESTED STAKEHOLDERS:		
CLOUD DEPLOYMENT MODEL:	☐ Public ☐ Private ☐ Community ☑ Hybrid	
CLOUD SERVICE MODEL:	☐ Software-as-a-Service (SaaS) ☑ Platform-as-a-Service (PaaS) ☐ Infrastructure-as-a-Service (IaaS) ☑ Other (example: Identity-as-a-Service)	
DESCRIPTION:	The use case is initiated when AIMS receives a provisioning request upon a notification from human resources, industrial security, or other authorized party. The meta data exchange occurs when AIMS provisions to the user to CIS and provides the data required. This workflow is triggered based on a request by the AIMS Administrator.	
PRE-CONDITION:	AIMS system must be able to issue a provision function with appropriate authorization. The provision action must be issued to the downstream integrated Cloud Identity System.	
TRIGGER:	The use case is initiated when AIMS receives a provisioning request upon a notification from human resources, industrial security, or other	

	authorized party. This workflow is triggered based on a request by the AIMS administrator.via a manual process or when an account provisioning request has been submitted.	
TYPICAL COURSE OF EVENTS:	**Actor Action**	**System Response**
	Step 1: The AIMS Administrator verifies identity of user	**Step 2**: None
	Step 3: The AIMS administrator types in the Credentialed Entity's Employee ID, SamAccountName, Last Name, and/or First Name and hits Enter. **(Alt Step 3)**	**Step 4**: AIMS pre-populates the Name Prefix, First Name, Preferred First Name, Middle Name, Last Name, Name Suffix, Credentialed Entity Company, Credentialed Entity Type , EmailID, Owner SamAccountName, OwnerAccountID, user Principal Name, DN, CN, Export Control Designation, Credential Type, and Location Fields from Active Directory and the human resources database table. **(Alt Step 4a) (Alt Step 4b)**
	Step 5: The AIMS administrator requests to provision the appropriate Credentialed Entity identity record and hits the Save button.	**Step 6:** AIMS validates the credentialed entity has been created.
		Step 7: AIMS sends provision packet to Cloud Identity System. **(Alt Step 7)**
	Step 8: Cloud Identity System sends acknowledgement.	**Step 9:** AIMS logs acknowledgement, data attributes confirmed by viewing log. Use case concludes.
ALTERNATE COURSES:	**Alt Step 3:** The AIMS Administrator types in the Account ID. If no identity record is pre-populated, the AIMS administrator has the option to manually create the record.	
	Alt Step 4a: AIMS is unable to find Credentialed Entity data from human resources or Active Directory. Error message is displayed.	
	Alt Step 4b: If multiple results that meet the criteria are found, AIMS displays a list of all users that meet that criterion. AIMS administrator selects appropriate user from list.	
	Alt. Step 7: AIMS is unable to establish a connection with Cloud Identity System to send data. AIMS retries every 2 minutes for 40 tries.	
CONCLUSION:	The use case concludes when CIS confirms the meta-data is passed from	

	AIMS to CIS. Log entry will confirm meta-data successfully exchanged.
POST-CONDITION:	AIMS sends a provisioning request to CIS.
BUSINESS RULES:	The AIMS administrator is responsible for approving all provisioning requests before the AIMS Technical Team provisions the requested resources. CIS data must match AIMS data attributes AIMS is the record of the authority
POLICY IMPACTS:	Sensitive data must be securely stored in CIS
IMPLEMENTATION CONSTRAINTS AND SPECIFICATIONS	Meta-data exchange rules must be created in the systems design. This will determine which attributes are sent to CIS
ASSUMPTIONS:	The use case uses the provisioning part of the identity life cycle to fulfill the meta-data exchange functionality.
OPEN ISSUES:	
NOTES/USE CASE DIAGRAM:	See Figure 23

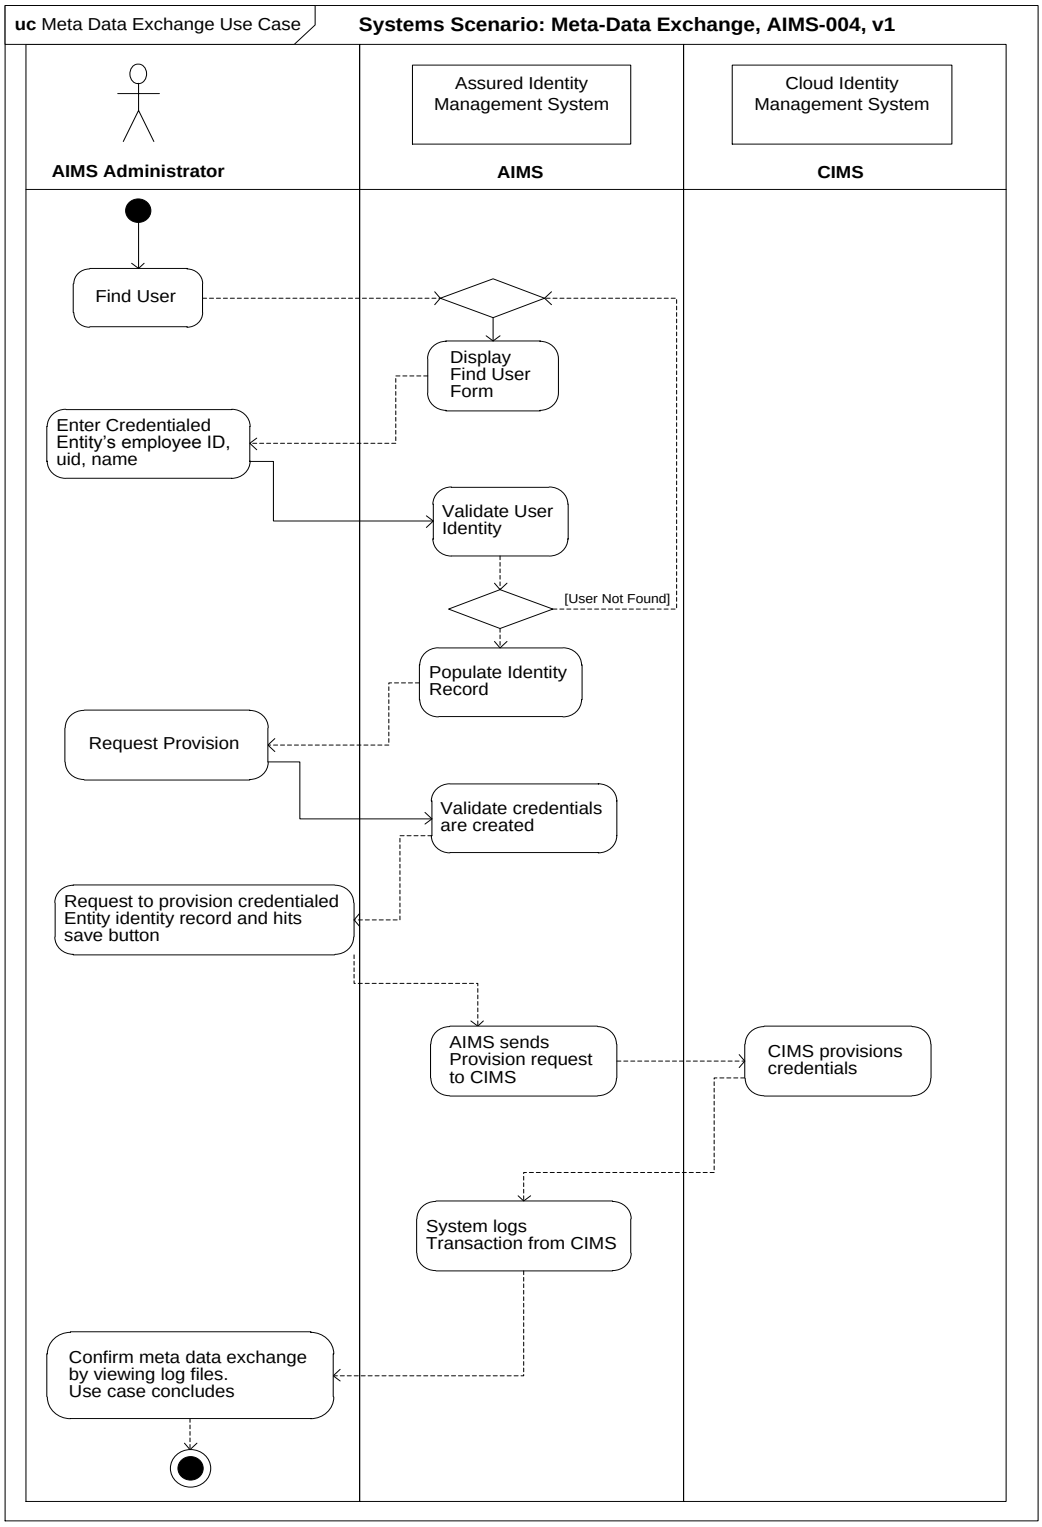

Figure 23.

Meta Data Exchange Use Case, AIMS-004

ASSURED IDENTITY MANAGEMENT SYSTEM
USE CASE PACKAGE AIMS-005

Author (s): **Daniels** **Date:** **07-February-2011**
Version: **1.1**

USE CASE NAME:	Revocation of Access	**USE CASE TYPE & LEVEL**
USE CASE ID:	AIMS-005, version 1	Business:
PRIORITY:	High	System/Solution:
		Requirements ☑
SOURCE:	*Research Project: Assured Identity for Cloud Computing (2011, Daniels)*	Analysis
PRIMARY BUSINESS ACTOR:	AIMS Technical Team	Design
		Fully Dressed ☑
PRIMARY SYSTEM ACTOR:	AIMS System	
OTHER PARTICIPATING ACTORS:	Cloud Identity System	
OTHER INTERESTED STAKEHOLDERS:		
CLOUD DEPLOYMENT MODEL:	☐ Public ☐ Private ☐ Community ☑ Hybrid	
CLOUD SERVICE MODEL:	☐ Software-as-a-Service (SaaS) ☑ Platform-as-a-Service (PaaS) ☐ Infrastructure-as-a-Service (IaaS) ☑ Other (example: Identity-as-a-Service)	
DESCRIPTION:	The use case is initiated when AIMS receives a revoke access request upon a notification from human resources, industrial security, or other authorized party. This workflow is triggered based on a revocation request by the AIMS Administrator.	
PRE-CONDITION:	AIMS system must be able to revoke access for an entity with appropriate authorization. The revocation action must be applicable to the downstream integrated Cloud Identity System.	

TRIGGER:	The use case is initiated when AIMS receives a revocation request upon a notification from human resources, physical security, information security, application owner, or other authorized party. This workflow is triggered based on a revocation request by the AIMS administrator	
TYPICAL COURSE OF EVENTS:	**Actor Action**	**System Response**
	Step 1: The AIMS Administrator navigates to the appropriate location and navigates to a "Find User" form.	**Step 2**: AIMS displays a blank "Find User" form.
	Step 3: The AIMS administrator types in the credentialed entity's Employee ID, SamAccountName, Last Name, and/or First Name and hits Enter.	**Step 4**: AIMS displays the entity's record.
	Step 5: The AIMS administrator requests to revoke access for the appropriate account record and hits the Save button.	**Step 6:** AIMS validates that record associated with the credentialed entity has been revoked.
	Step 7: The AIMS administrator reviews the audit log to confirm access has been revoked.	**Step 8:** AIMS displays confirmation of access revocation
	Step 9: Use case concludes	
ALTERNATE COURSES:	**None**	
CONCLUSION:	The use case concludes when Cloud Identity System confirms account record and access have been revoked.	
POST-CONDITION:	None	
BUSINESS RULES:	Revocation of access due to termination, firing, suspension, or similar action AIMS is the record of authority.	
POLICY IMPACTS:	NA	
IMPLEMENTATION		

CONSTRAINTS AND SPECIFICATIONS	
ASSUMPTIONS:	
OPEN ISSUES:	
NOTES/USE CASE DIAGRAM:	See Figure 24

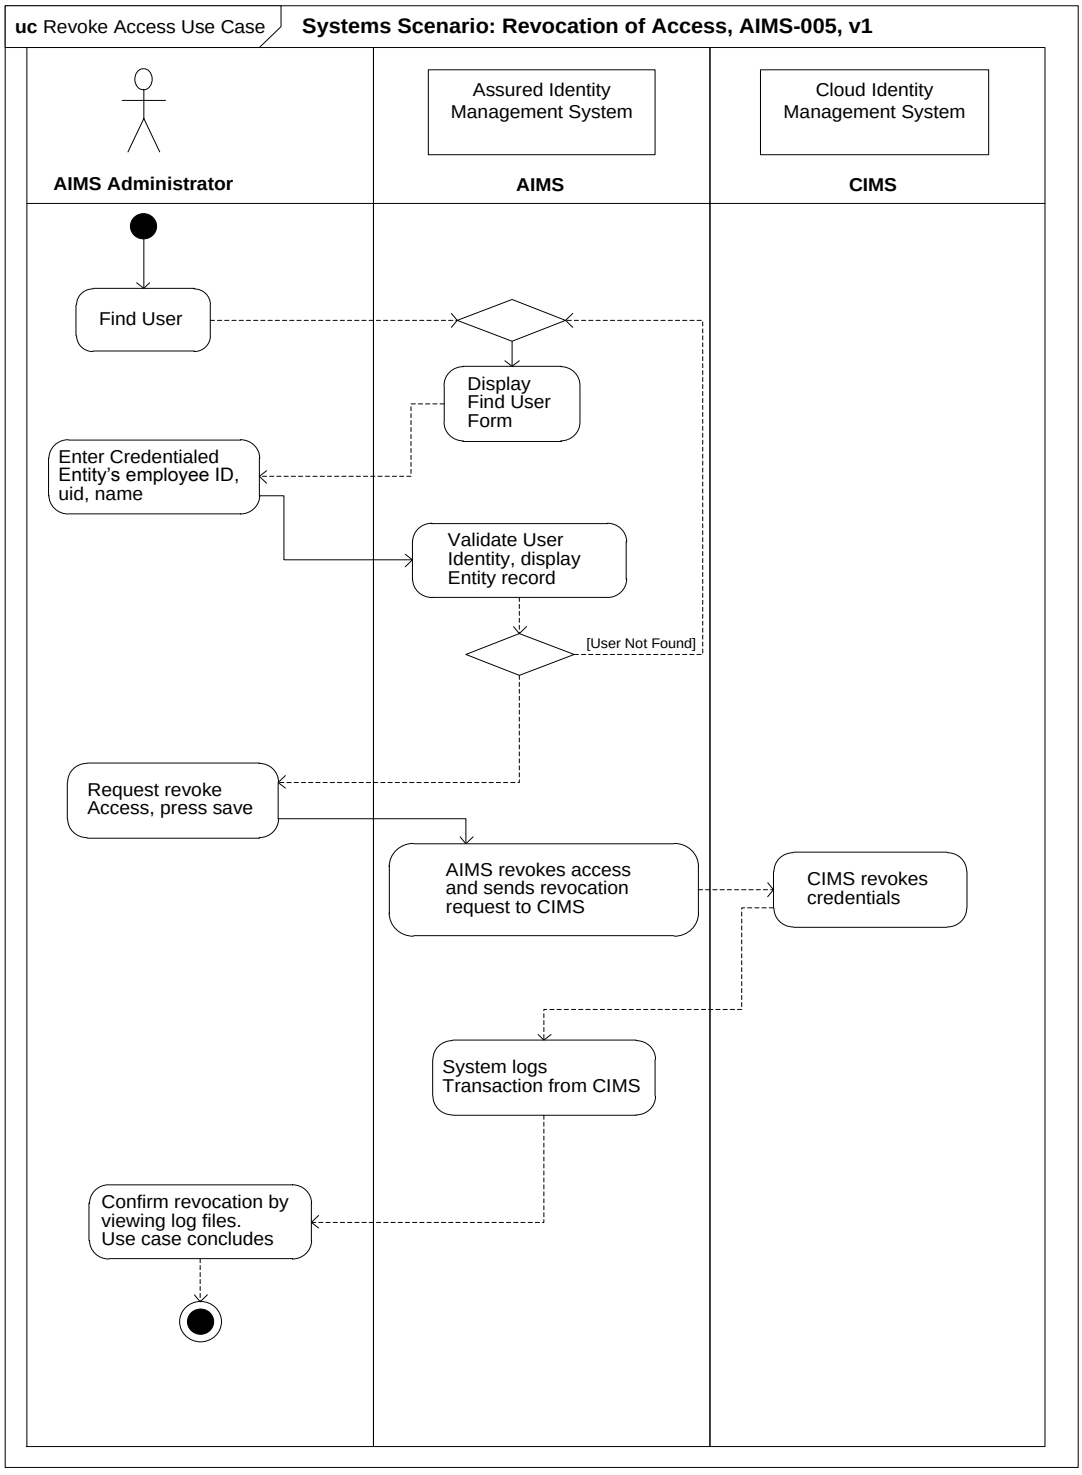

Figure 24.

Revocation of Access Use Case, AIMS-005

ASSURED IDENTITY MANAGEMENT SYSTEM
USE CASE PACKAGE AIMS-006

Author (s): **Daniels** **Date:** **07-February-2011**
Version: **1.1**

USE CASE NAME:	Update Identity Record	**USE CASE TYPE & LEVEL** Business:
USE CASE ID:	AIMS-006, version 1	
PRIORITY:	High	System/Solution: Requirements ☑
SOURCE:	*Research Project: Assured Identity for Cloud Computing (2011, Daniels)*	Analysis
PRIMARY BUSINESS ACTOR:	AIMS Technical Team	Design Fully Dressed ☑
PRIMARY SYSTEM ACTOR:	AIMS System, End User	
OTHER PARTICIPATING ACTORS:	Cloud Identity System	
OTHER INTERESTED STAKEHOLDERS:		
CLOUD DEPLOYMENT MODEL:	☐ Public ☐ Private ☐ Community ☑ Hybrid	
CLOUD SERVICE MODEL:	☐ Software-as-a-Service (SaaS) ☑ Platform-as-a-Service (PaaS) ☐ Infrastructure-as-a-Service (IaaS) ☑ Other (example: Identity-as-a-Service)	
DESCRIPTION:	The use case is initiated when AIMS receives a request to update an existing record from human resources, individual, or other authorized party. This workflow is triggered based on a revocation request by the AIMS Administrator.	
PRE-CONDITION:	AIMS system must be able to update information for an entity with appropriate authorization. The update action must be applicable to the integrated Cloud Identity System.	
TRIGGER:	The use case is initiated when AIMS receives an update request upon a notification from human resources, individual, or other authorized party. This workflow is triggered based on a revocation request by the AIMS	

	administrator	
TYPICAL COURSE OF EVENTS:	**Actor Action**	**System Response**
	Step 1: The AIMS Administrator navigates to the appropriate location and navigates to a "Find User" form.	**Step 2**: AIMS displays a blank "Find User" form.
	Step 3: The AIMS administrator types in the credentialed entity's Employee ID, SamAccountName, Last Name, and/or First Name and hits Enter.	**Step 4**: AIMS displays the entity's record.
	Step 5: The AIMS administrator updates appropriate data fields for the account record and hits the Save button.	**Step 6:** AIMS validates that record associated with the credentialed entity has been updated.
	Step 7: The AIMS administrator reviews the audit log to confirm record has been updated.	
	Step 8: Use case concludes	
ALTERNATE COURSES:	**None**	
CONCLUSION:	The use case concludes when AIMS administrator confirms account record has been updated information	
POST-CONDITION:	None	
BUSINESS RULES:	Name change, organizational change might drive a record to be updated.	
POLICY IMPACTS:	NA	
IMPLEMENTATION CONSTRAINTS AND SPECIFICATIONS		
ASSUMPTIONS:		
OPEN ISSUES:		
NOTES/USE CASE	See Figure 25	

DIAGRAM:	

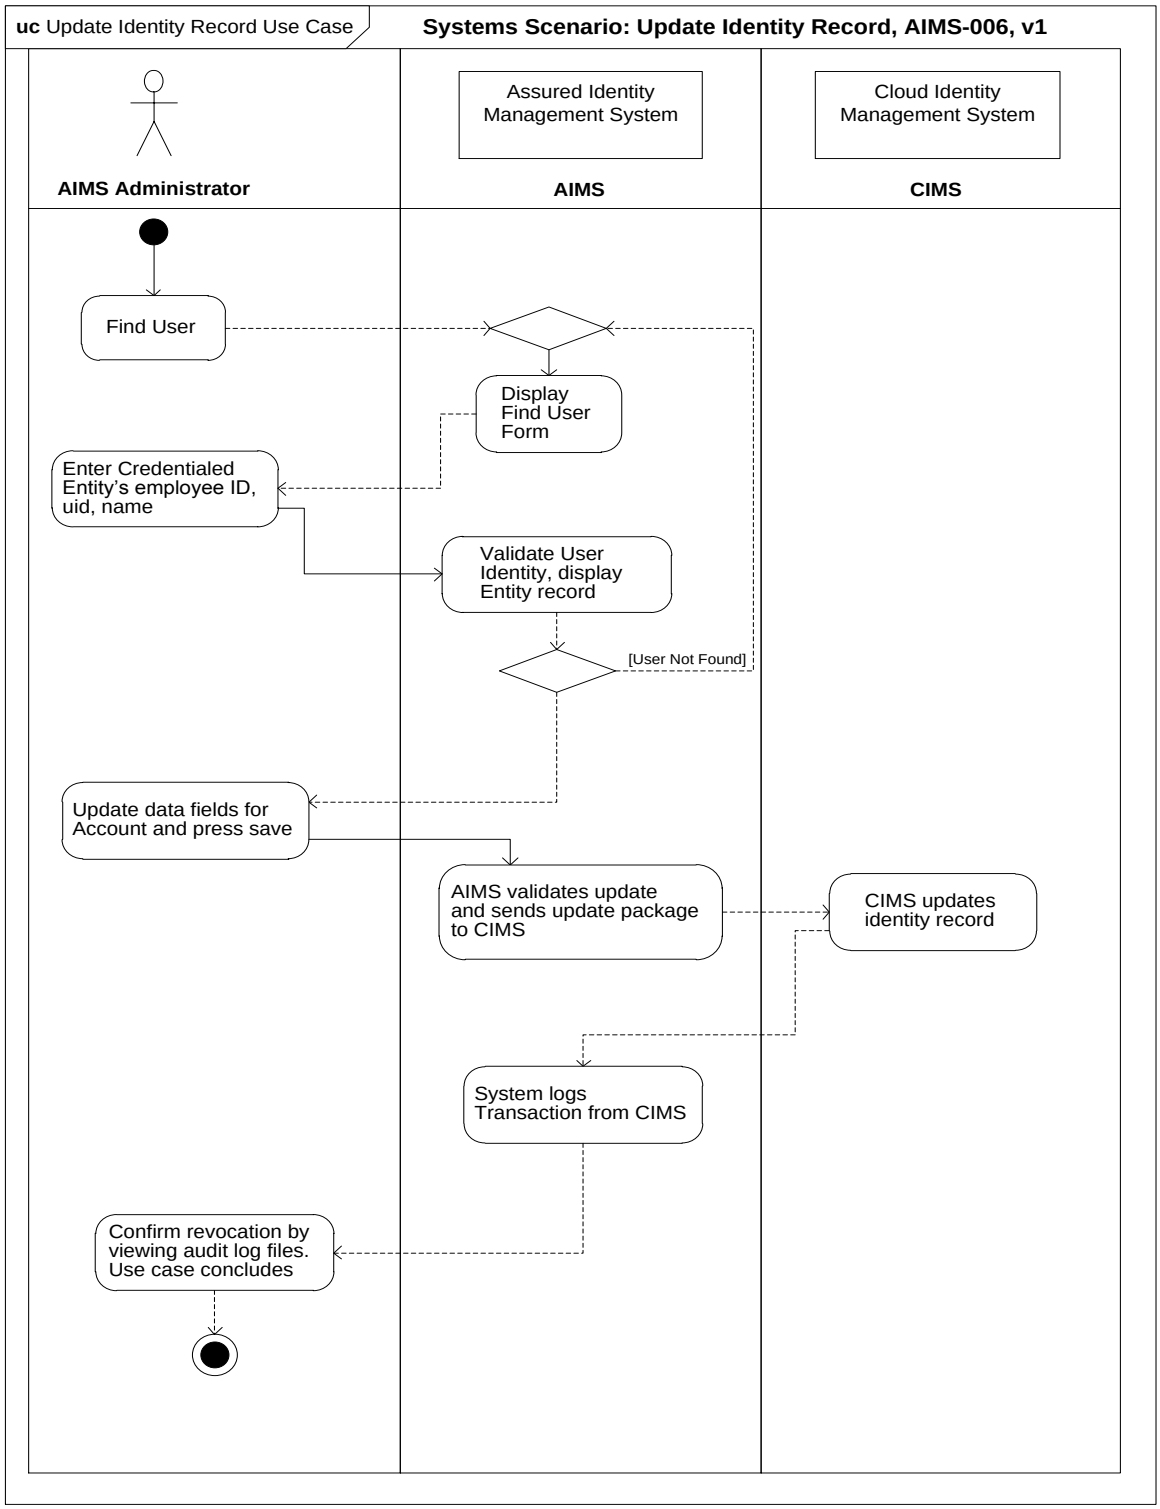

Figure 25.

Update Identity Record Use Case, AIMS-006

ASSURED IDENTITY MANAGEMENT SYSTEM
USE CASE PACKAGE AIMS-007

Author (s): **Daniels** **Date:** **07-February-2011**
Version: **1.1**

USE CASE NAME:	Run Audit Log	**USE CASE TYPE & LEVEL** Business:
USE CASE ID:	AIMS-007, version 1	
PRIORITY:	High	System/Solution: Requirements ☑
SOURCE:	*Research Project: Assured Identity for Cloud Computing (2011, Daniels)*	Analysis
PRIMARY BUSINESS ACTOR:	AIMS Technical Team	Design Fully Dressed ☑
PRIMARY SYSTEM ACTOR:	AIMS System	
OTHER PARTICIPATING ACTORS:	Cloud Identity System	
OTHER INTERESTED STAKEHOLDERS:		
CLOUD DEPLOYMENT MODEL:	☐ Public ☐ Private ☐ Community ☑ Hybrid	
CLOUD SERVICE MODEL:	☐ Software-as-a-Service (SaaS) ☑ Platform-as-a-Service (PaaS) ☐ Infrastructure-as-a-Service (IaaS) ☑ Other (example: Identity-as-a-Service)	
DESCRIPTION:	The use case is initiated when AIMS Administrator checks the audit log.	
PRE-CONDITION:	AIMS system must be able to synchronize data between AIMS and the Cloud Identity System	
TRIGGER:	The use case is initiated when AIMS administrator accesses audit information for a user record including identity record updates, provisioning, and deprovisioning This workflow is triggered based on audit information access by the AIMS administrator	
TYPICAL COURSE OF EVENTS:	**Actor Action**	**System Response**
	Step 1: The AIMS Administrator navigates to the audit tab.	**Step 2**: AIMS displays audit information form

	Step 3: The AIMS administrator types in the credentialed entity's Employee ID, SamAccountName, Last Name, and/or First Name and hits Enter.	**Step 4**: AIMS displays the entity's audit record.
	Step 5: The AIMS administrator confirms actions for provisioning, deprovisioning, access revocation, and updates	**Step 6:** AIMS validates that record associated with the request has valid audit information.
	Step 7: The AIMS administrator reviews the audit log to confirm record has been updated.	
	Step 8: Use case concludes	
ALTERNATE COURSES:	**None**	
CONCLUSION:	The use case concludes when AIMS administrator confirms account record audit information is accurate	
POST-CONDITION:	None	
BUSINESS RULES:	Auditing, accounting, and compliance	
POLICY IMPACTS:	NA	
IMPLEMENTATION CONSTRAINTS AND SPECIFICATIONS		
ASSUMPTIONS:		
OPEN ISSUES:		
NOTES/USE CASE DIAGRAM:	See Figure 26	

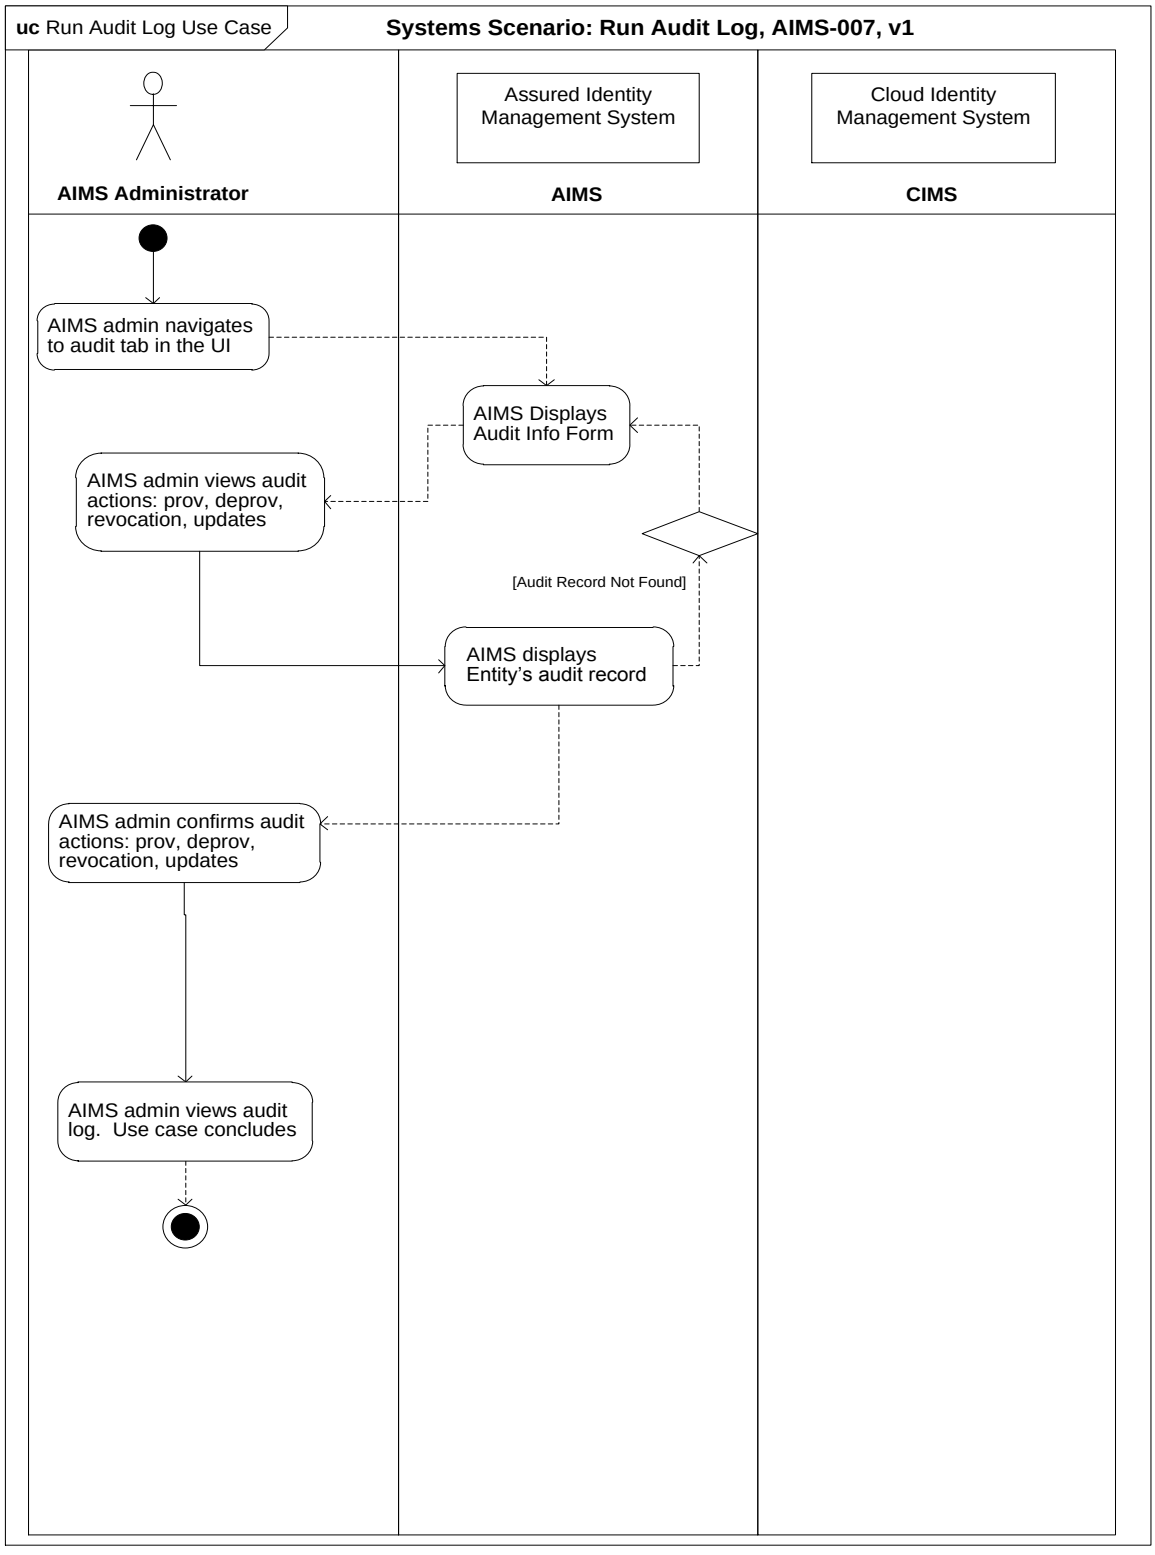

Figure 26.

Run Audit Log Use Case, AIMS-007

ASSURED IDENTITY MANAGEMENT SYSTEM
USE CASE PACKAGE AIMS-008

Author (s): **Daniels** **Date:** **07-February-2011**
Version: **1.1**

USE CASE NAME:	Synchronization of Identity Record	**USE CASE TYPE & LEVEL**	
USE CASE ID:	AIMS-008, version 1	Business:	
PRIORITY:	High	System/Solution:	
		Requirements	☑
SOURCE:	*Research Project: Assured Identity for Cloud Computing (2011, Daniels)*	Analysis	
PRIMARY BUSINESS ACTOR:	AIMS Technical Team	Design	
		Fully Dressed	☑
PRIMARY SYSTEM ACTOR:	AIMS System		
OTHER PARTICIPATING ACTORS:	Cloud Identity System		
OTHER INTERESTED STAKEHOLDERS:			
CLOUD DEPLOYMENT MODEL:	☐ Public ☐ Private ☐ Community ☑ Hybrid		
CLOUD SERVICE MODEL:	☐ Software-as-a-Service (SaaS) ☑ Platform-as-a-Service (PaaS) ☐ Infrastructure-as-a-Service (IaaS) ☑ Other (example: Identity-as-a-Service)		
DESCRIPTION:	The use case is initiated when AIMS Administrator checks the audit log.		
PRE-CONDITION:	AIMS system must be able to synchronize data between AIMS and the Cloud Identity System		
TRIGGER:	The use case is initiated when AIMS administrator accesses audit information for a user record including identity record updates, provisioning, and deprovisioning This workflow is triggered based on audit information access by the AIMS administrator		
TYPICAL COURSE OF EVENTS:	**Actor Action**	**System Response**	
	Step 1: The AIMS	**Step 2**: AIMS displays resources	

	Administrator navigates to the resources tab.	
	Step 3: The AIMS administrator clicks synchronization with Cloud Identity System	**Step 4**: AIMS runs synchronization job.
	Step 5: The AIMS administrator confirms synchronization job accessing application logs	**Step 6:** AIMS validates the synchronization job is complete
	Step 7: Use case concludes	
ALTERNATE COURSES:	None	
CONCLUSION:	The use case concludes when AIMS administrator confirms synchronization between AIMS and the Cloud Identity System	
POST-CONDITION:	None	
BUSINESS RULES:	Routine updates, data availability	
POLICY IMPACTS:	NA	
IMPLEMENTATION CONSTRAINTS AND SPECIFICATIONS		
ASSUMPTIONS:		
OPEN ISSUES:		
NOTES/USE CASE DIAGRAM:	See Figure 27	

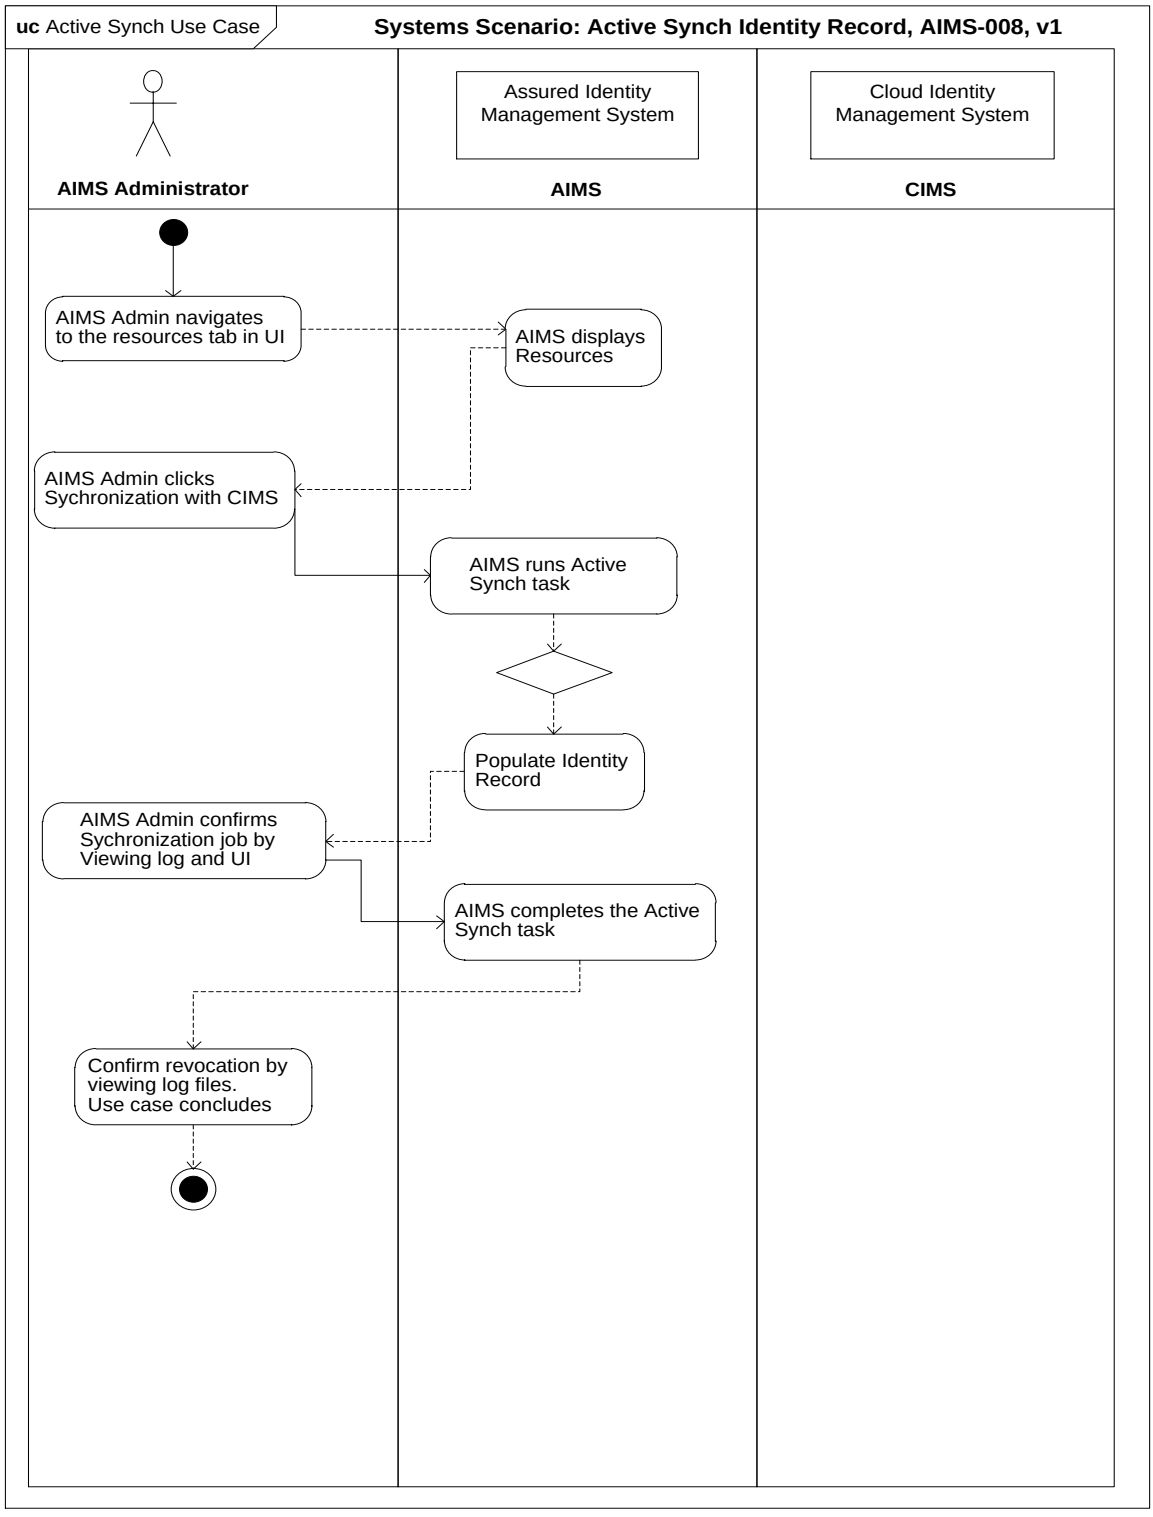

Figure 27.

Active Sync Use Case, AIMS-008

ASSURED IDENTITY MANAGEMENT SYSTEM
USE CASE PACKAGE AIMS-009

Author (s): **Daniels** **Date:** **07-February-2011**
Version: **1.1**

USE CASE NAME:	Change Level of Assurance	**USE CASE TYPE & LEVEL**
USE CASE ID:	AIMS-009, version 1	Business:
PRIORITY:	High	System/Solution:
		Requirements ☑
SOURCE:	*Research Project: Assured Identity for Cloud Computing (2011, Daniels)*	Analysis
PRIMARY BUSINESS ACTOR:	AIMS Technical Team	Design
		Fully Dressed ☑
PRIMARY SYSTEM ACTOR:	AIMS System, End User	
OTHER PARTICIPATING ACTORS:	Cloud Identity System	
OTHER INTERESTED STAKEHOLDERS:		
CLOUD DEPLOYMENT MODEL:	☐ Public ☐ Private ☐ Community ☑ Hybrid	
CLOUD SERVICE MODEL:	☐ Software-as-a-Service (SaaS) ☑ Platform-as-a-Service (PaaS) ☐ Infrastructure-as-a-Service (IaaS) ☑ Other (example: Identity-as-a-Service)	
DESCRIPTION:	The use case is initiated when AIMS receives a request to change level of assurance for an existing record from human resources, information security, individual, or other authorized party. This workflow is triggered based on a revocation request by the AIMS Administrator.	
PRE-CONDITION:	AIMS system must be able to change the level of assurance for an entity with appropriate authorization. The update action must be applicable to the integrated Cloud Identity System.	
TRIGGER:	The use case is initiated when AIMS receives a change level of assurance request upon a notification from human resources, information security,	

	individual, or other authorized party. This workflow is triggered based on a change level of assurance request by the AIMS administrator	
TYPICAL COURSE OF EVENTS:	**Actor Action**	**System Response**
	Step 1: The AIMS Administrator navigates to the appropriate location and navigates to a "Find User" form.	**Step 2**: AIMS displays a blank "Find User" form.
	Step 3: The AIMS administrator types in the credentialed entity's Employee ID, SamAccountName, Last Name, and/or First Name and hits Enter.	**Step 4**: AIMS displays the entity's record.
	Step 5: The AIMS administrator updates appropriate data field for level of assurance for the account record and hits the Save button.	**Step 6:** AIMS validates that record associated with the credentialed entity has been updated with a new level of assurance.
	Step 7: The AIMS administrator reviews the audit log to confirm record has been updated.	
	Step 8: Use case concludes	
ALTERNATE COURSES:	**None**	
CONCLUSION:	The use case concludes when AIMS administrator confirms account record has been changed to reflect the new level of assurance.	
POST-CONDITION:	None	
BUSINESS RULES:	Organizational change, project assignment, or other event might drive a record to change the associated level of assurance.	
POLICY IMPACTS:	NA	
IMPLEMENTATION CONSTRAINTS AND SPECIFICATIONS		

ASSUMPTIONS:	
OPEN ISSUES:	
NOTES/USE CASE DIAGRAM:	See Figure 28

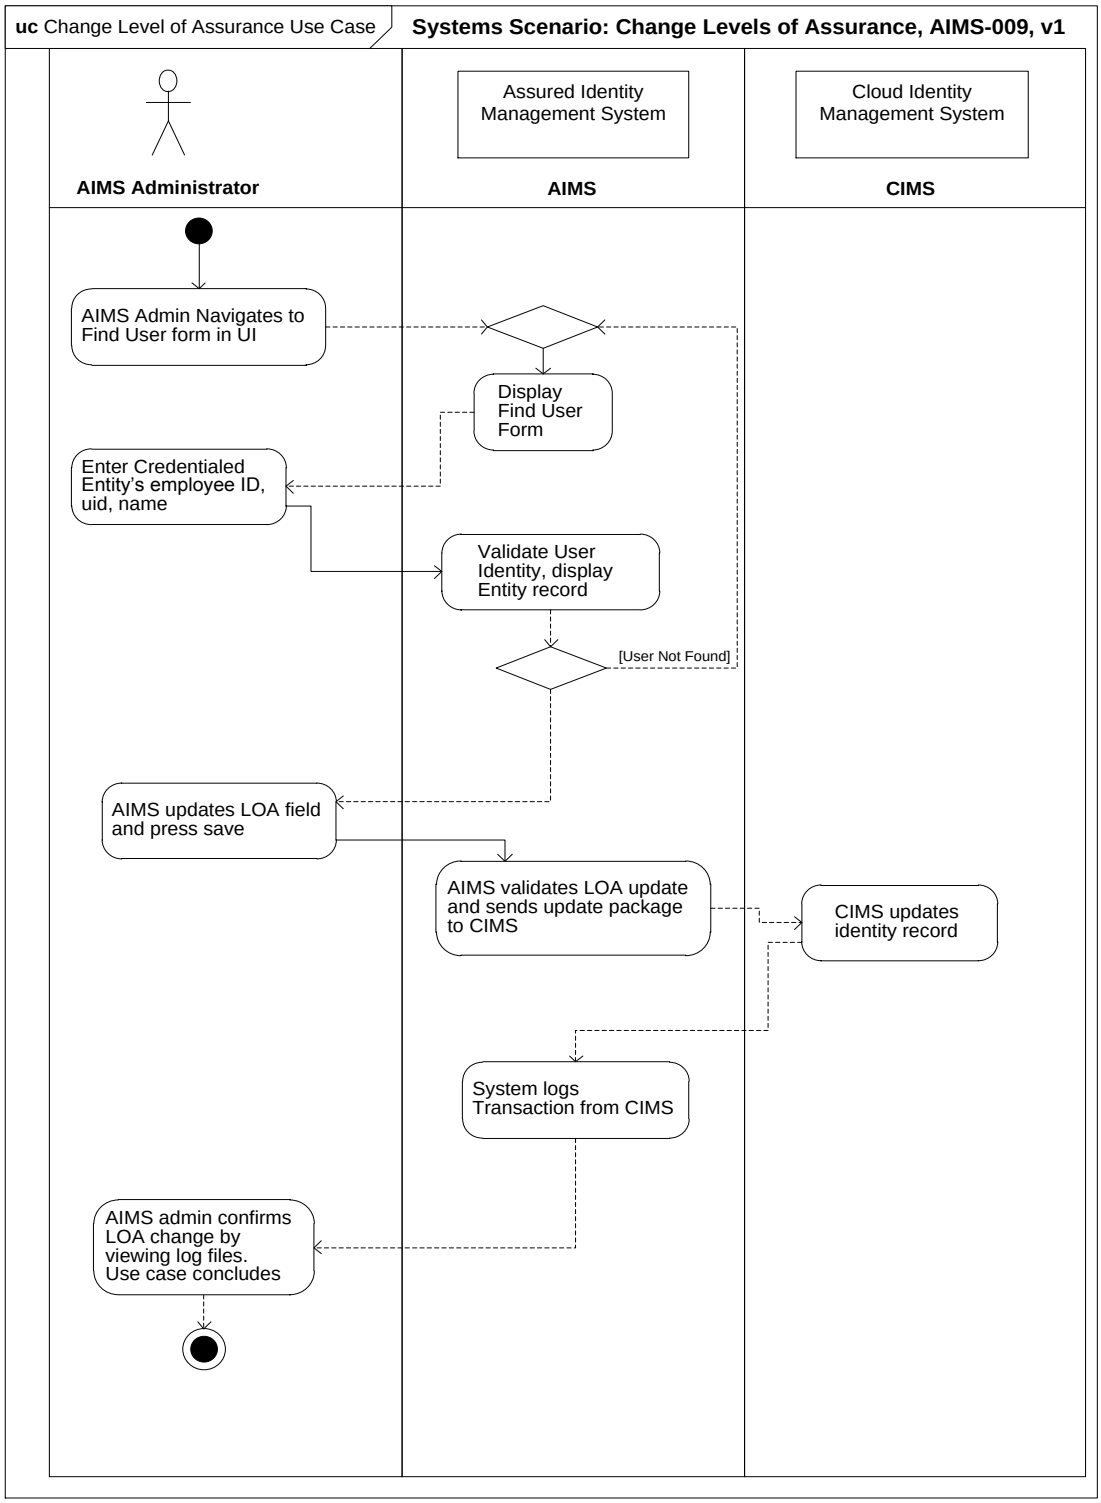

Figure 28.

Change Level of Assurance Use Case, AIMS-009

ASSURED IDENTITY MANAGEMENT SYSTEM
USE CASE PACKAGE AIMS-010

Author (s): **Daniels** **Date:** **07-February-2011**
Version: **1.1**

USE CASE NAME:	Secure Identification of End-User	**USE CASE TYPE & LEVEL** Business: System/Solution: Requirements ☑ Analysis Design Fully Dressed ☑
USE CASE ID:	AIMS-010, version 1	
PRIORITY:	High	
SOURCE:	*Research Project: Assured Identity for Cloud Computing (2011, Daniels)*	
PRIMARY BUSINESS ACTOR:	AIMS Technical Team	
PRIMARY SYSTEM ACTOR:	AIMS System, End User	
OTHER PARTICIPATING ACTORS:		
OTHER INTERESTED STAKEHOLDERS:		
CLOUD DEPLOYMENT MODEL:	☐ Public ☐ Private ☐ Community ☑ Hybrid	
CLOUD SERVICE MODEL:	☐ Software-as-a-Service (SaaS) ☑ Platform-as-a-Service (PaaS) ☐ Infrastructure-as-a-Service (IaaS) ☑ Other (example: Identity-as-a-Service)	
DESCRIPTION:	The use case is initiated when AIMS receives a request to identify a new or existing user record from human resources, information security, individual, or other authorized party. This workflow is triggered based on an identification request by the AIMS Administrator.	
PRE-CONDITION:	AIMS system must be able to confirm the identity for an entity with appropriate authorization. The update action must be applicable to the integrated Cloud Identity System.	
TRIGGER:	The use case is initiated when AIMS receives a receives a request to identify a new or existing user record from human resources, information security, individual, or other authorized party. This workflow is triggered	

	based on an identification request by the AIMS administrator	
TYPICAL COURSE OF EVENTS:	**Actor Action**	**System Response**
	Step 1: The AIMS Administrator navigates to the appropriate location and navigates to a "Find User" form.	**Step 2**: AIMS displays a blank "Find User" form.
	Step 3: The AIMS administrator types in the credentialed entity's Employee ID, SamAccountName, Last Name, and/or First Name and hits Enter.	**Step 4**: AIMS displays the entity's record.
	Step 5: The AIMS administrator confirms the identification of the end user.	**Step 6:** Use case concludes.
ALTERNATE COURSES:	**None**	
CONCLUSION:	The use case concludes when AIMS administrator confirms the identity of the end-user against the existing AIMS record.	
POST-CONDITION:	None	
BUSINESS RULES:	Identity confirmation may be required for new project assignments, travel, or similar organizational change, and/or audit functions.	
POLICY IMPACTS:	NA	
IMPLEMENTATION CONSTRAINTS AND SPECIFICATIONS		
ASSUMPTIONS:		
OPEN ISSUES:		
NOTES/USE CASE DIAGRAM:	See Figure 29	

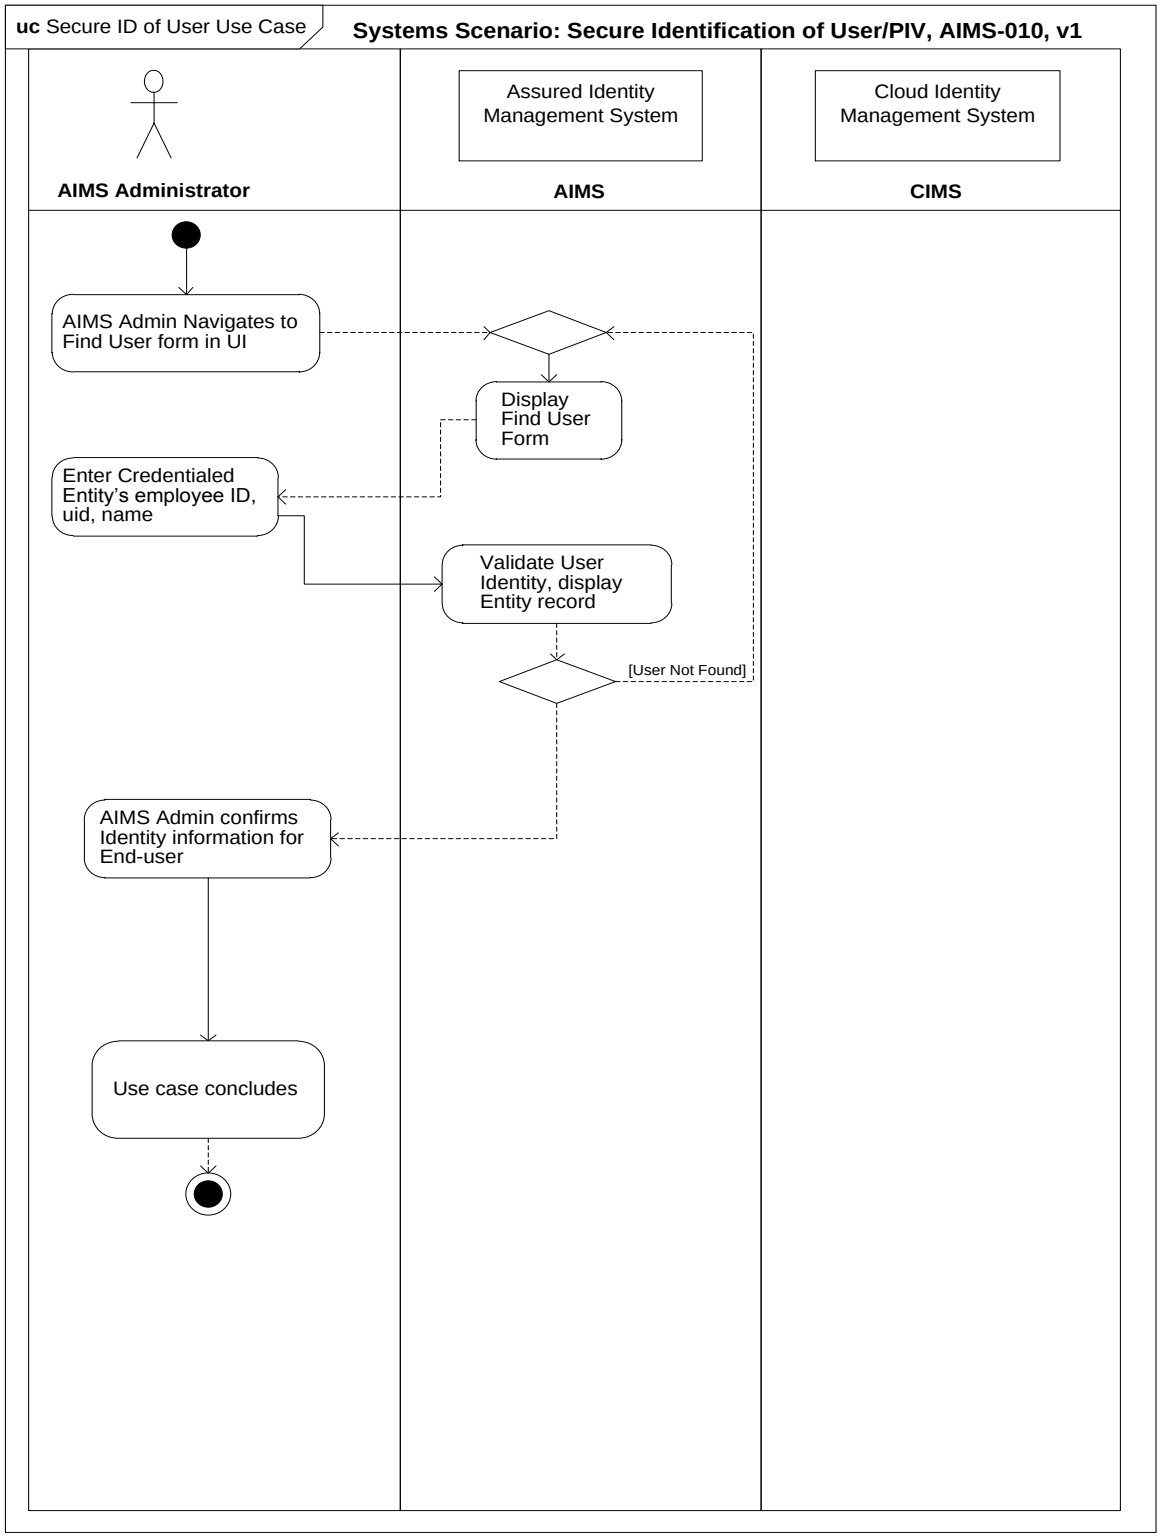

Figure 29.

Secure Identification of User, AIMS-010

Use Case Validation

Use case validation was performed by testing the scenarios in a sequence diagram format. The SysML standard describes sequence diagrams as a way to visually depict operations and process interaction between components. The investigation created sequence diagrams for the purpose of testing functionality as well as the visual benefit.

Designated by an "sd" label, the sequence diagrams depict a series of vertical "swim lanes" owned by an actor from the use case scenarios. In most scenarios, the actors were the AIMS Administrator, AIMS system, and CIMS system. Horizontal lines denoted the action being taken on a particular step of the test. The gray boxes placed vertically on the diagram between process lines indicate an action or "processing" takes place between the actors. The sequence diagrams were executed according to the descriptors on the left side of the diagram. In conjunction with use cases, the sequence diagrams were a particularly helpful tool in determining what actions should happen at each step as well as the final expect result at the conclusion of the scenario. A summary of the sequence diagrams is displayed Table 8. Ten sequences were tested, aligning with the ten use case scenarios.

Table 8.

Sequence Diagram Test Cases

Sequence Diagram Title	Sequence Test Number	Sequence Diagram Description
TC-001-Deprovision_Account	TC-001	Deprovision Account
TC-002-Provision_Account	TC-002	Provision Account
TC-003-SSO_Across_Clouds	TC-003	Single Sign On to Cloud Environment
TC-004-Meta-Data-Exchange	TC-004	Exchange of Meta-Data attributes
TC-005-Revoke_Access	TC-005	Revocation of access
TC-006-Update_Identity_Record	TC-006	Update existing identity record
TC-007-Run_Audit_Log	TC-007	Audit log
TC-008-Active_Synch_Identity_Record	TC-008	Synchronization of Identity Record
TC-009-Change_Levels_of_Assurance	TC-009	Change Levels of Assurance
TC-010-Secure Identification of User/PIV	TC-010	Secure Identification of End-User

Sequence Diagrams

Each of the ten sequence diagram test scenarios are depicted in SysML form (see figures 30-39):

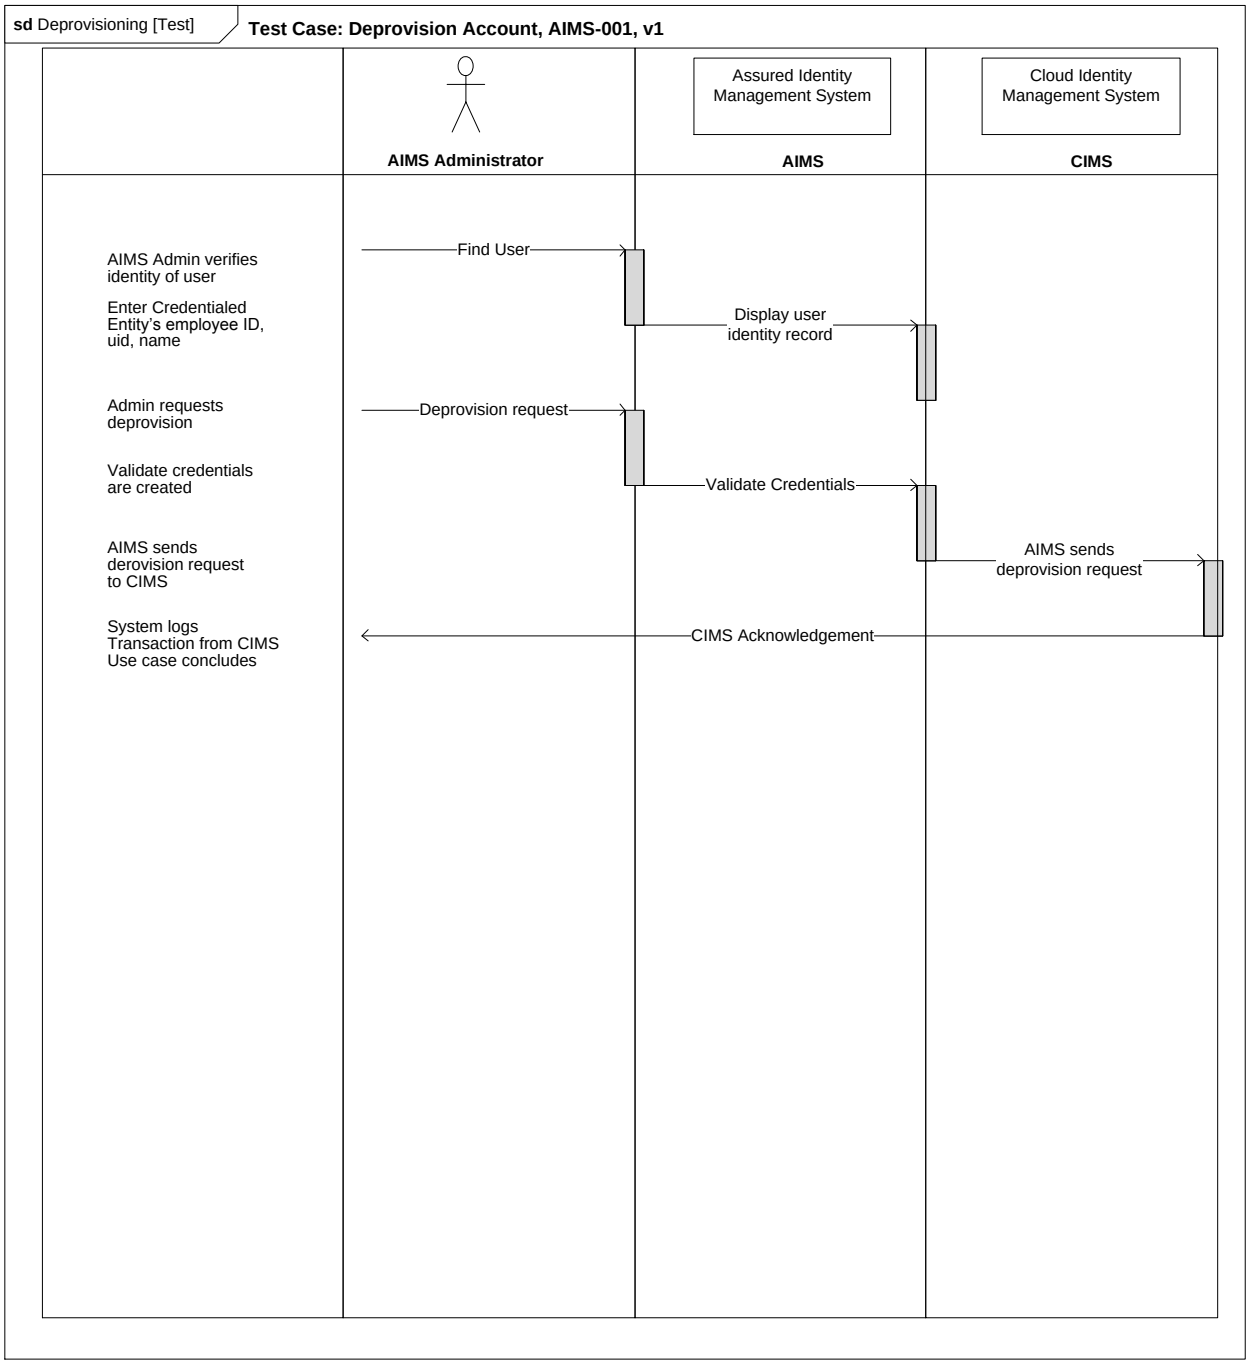

Figure 30.

Deprovision Sequence Diagram TC-001, SysML

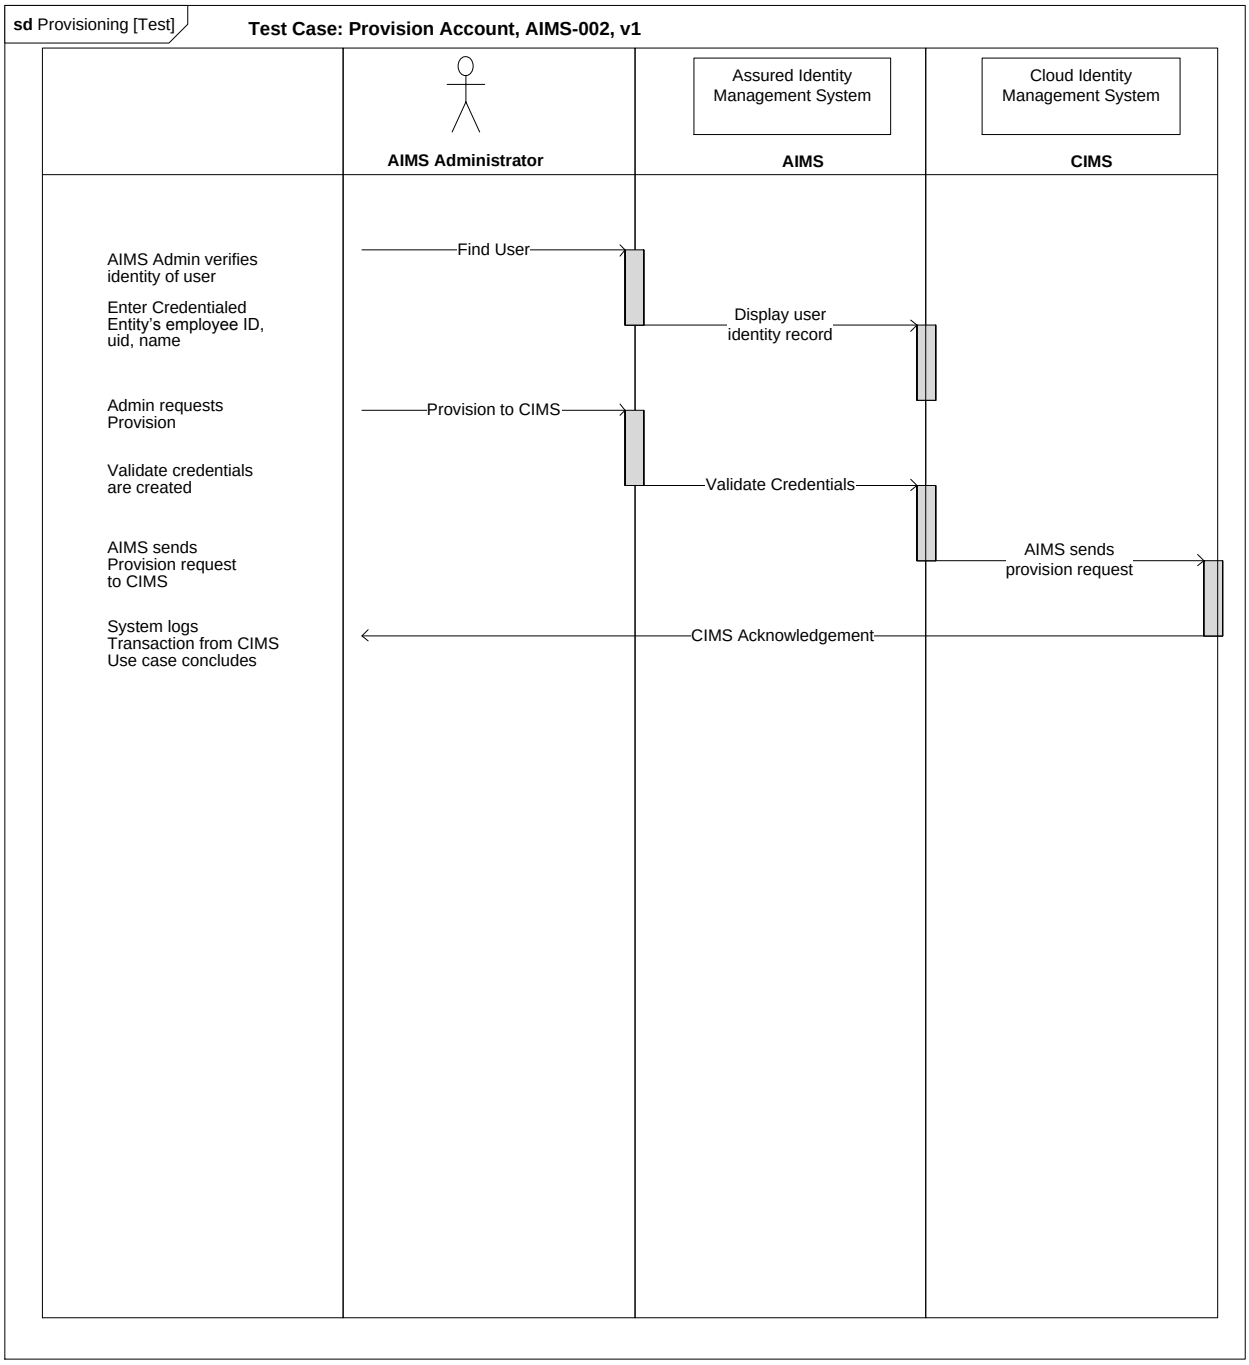

Figure 31.

Provision Sequence Diagram TC-002, SysML

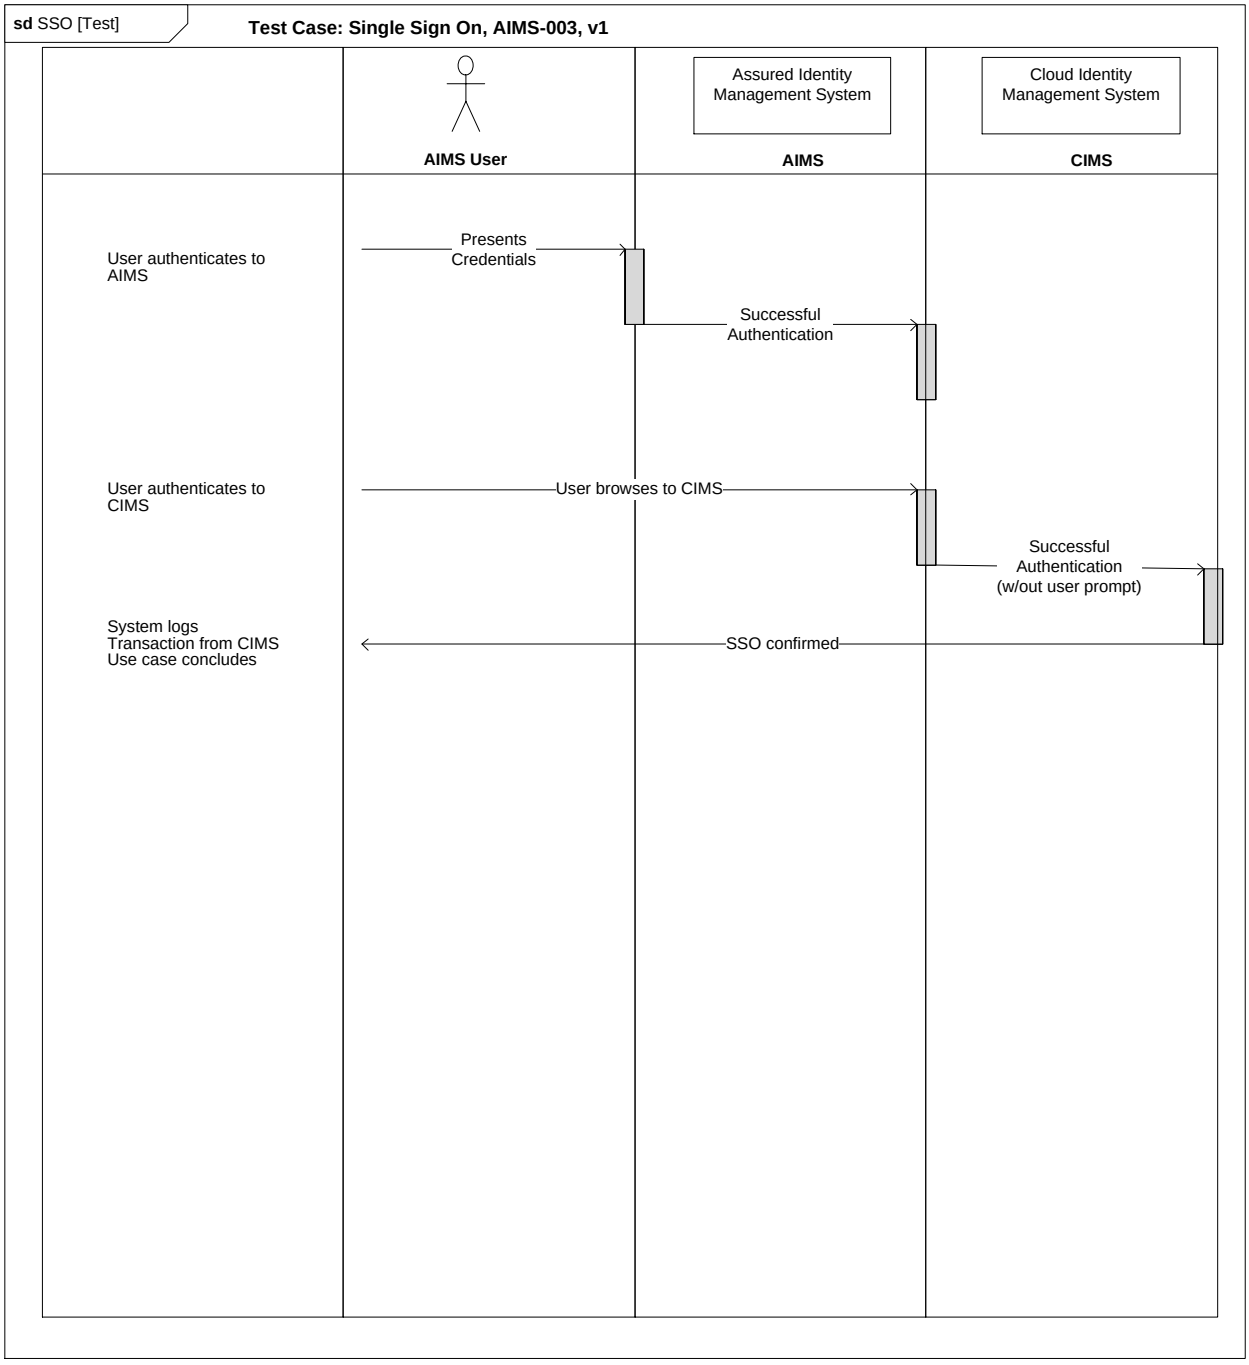

Figure 32.

Single Sign-On (SSO) Sequence Diagram TC-003, SysML

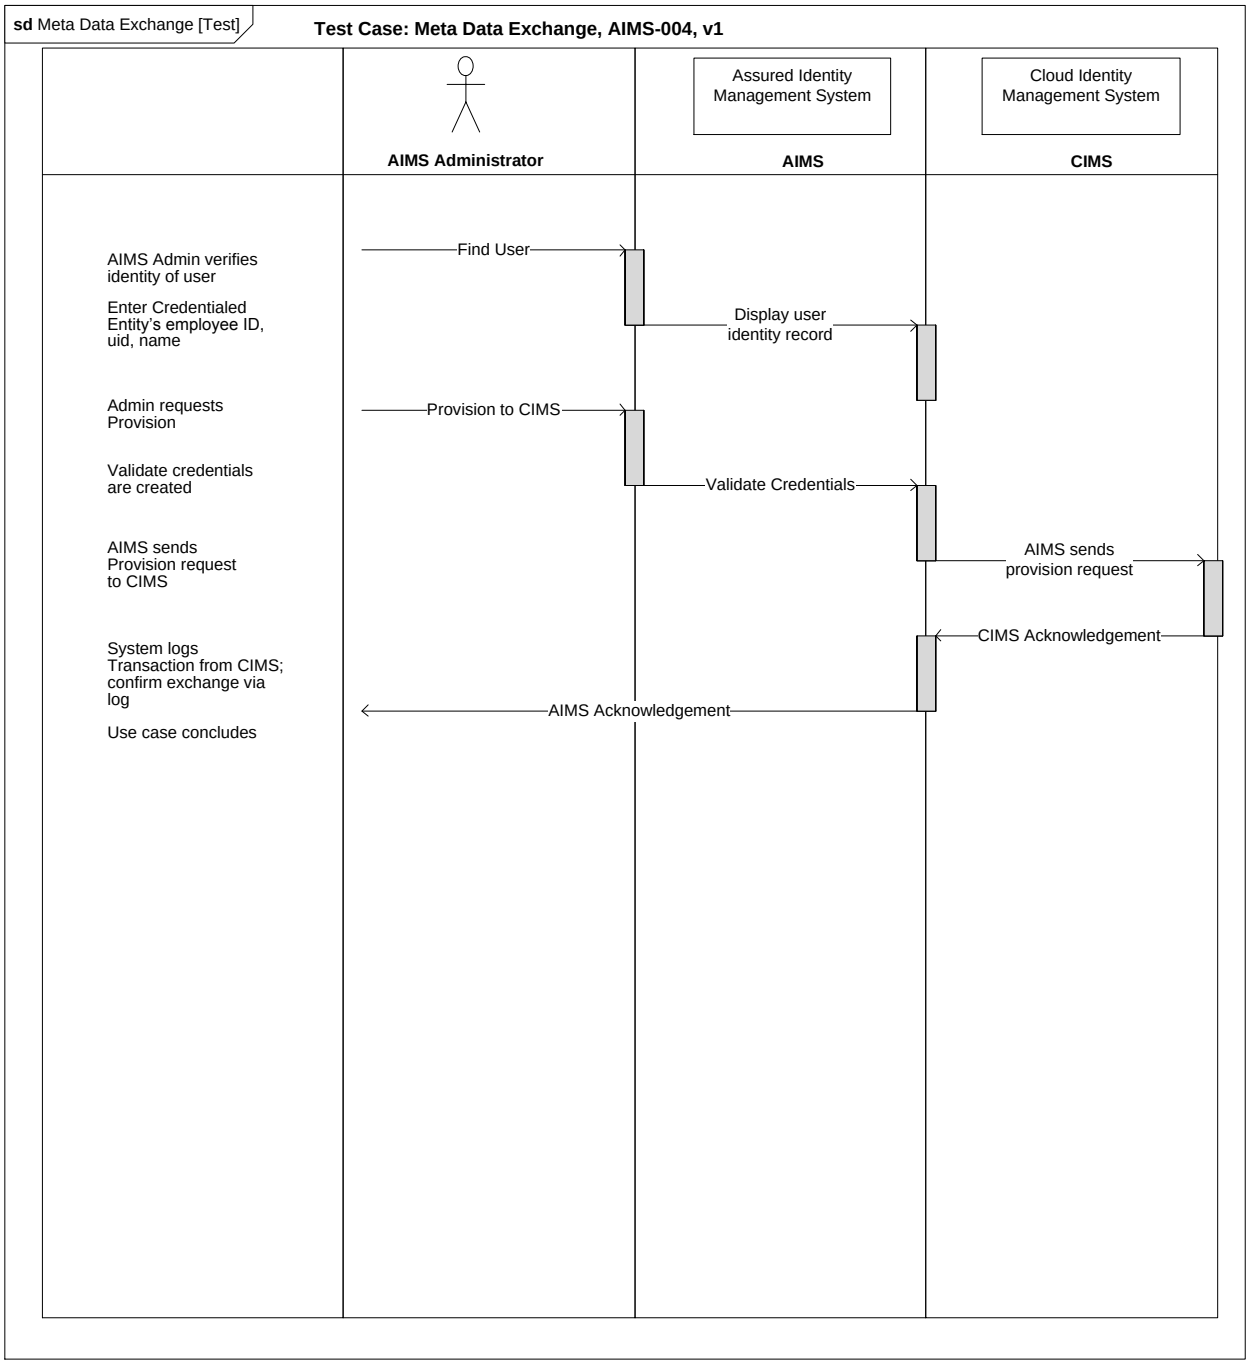

Figure 33.

Meta-Data Exchange Sequence Diagram, TC-004

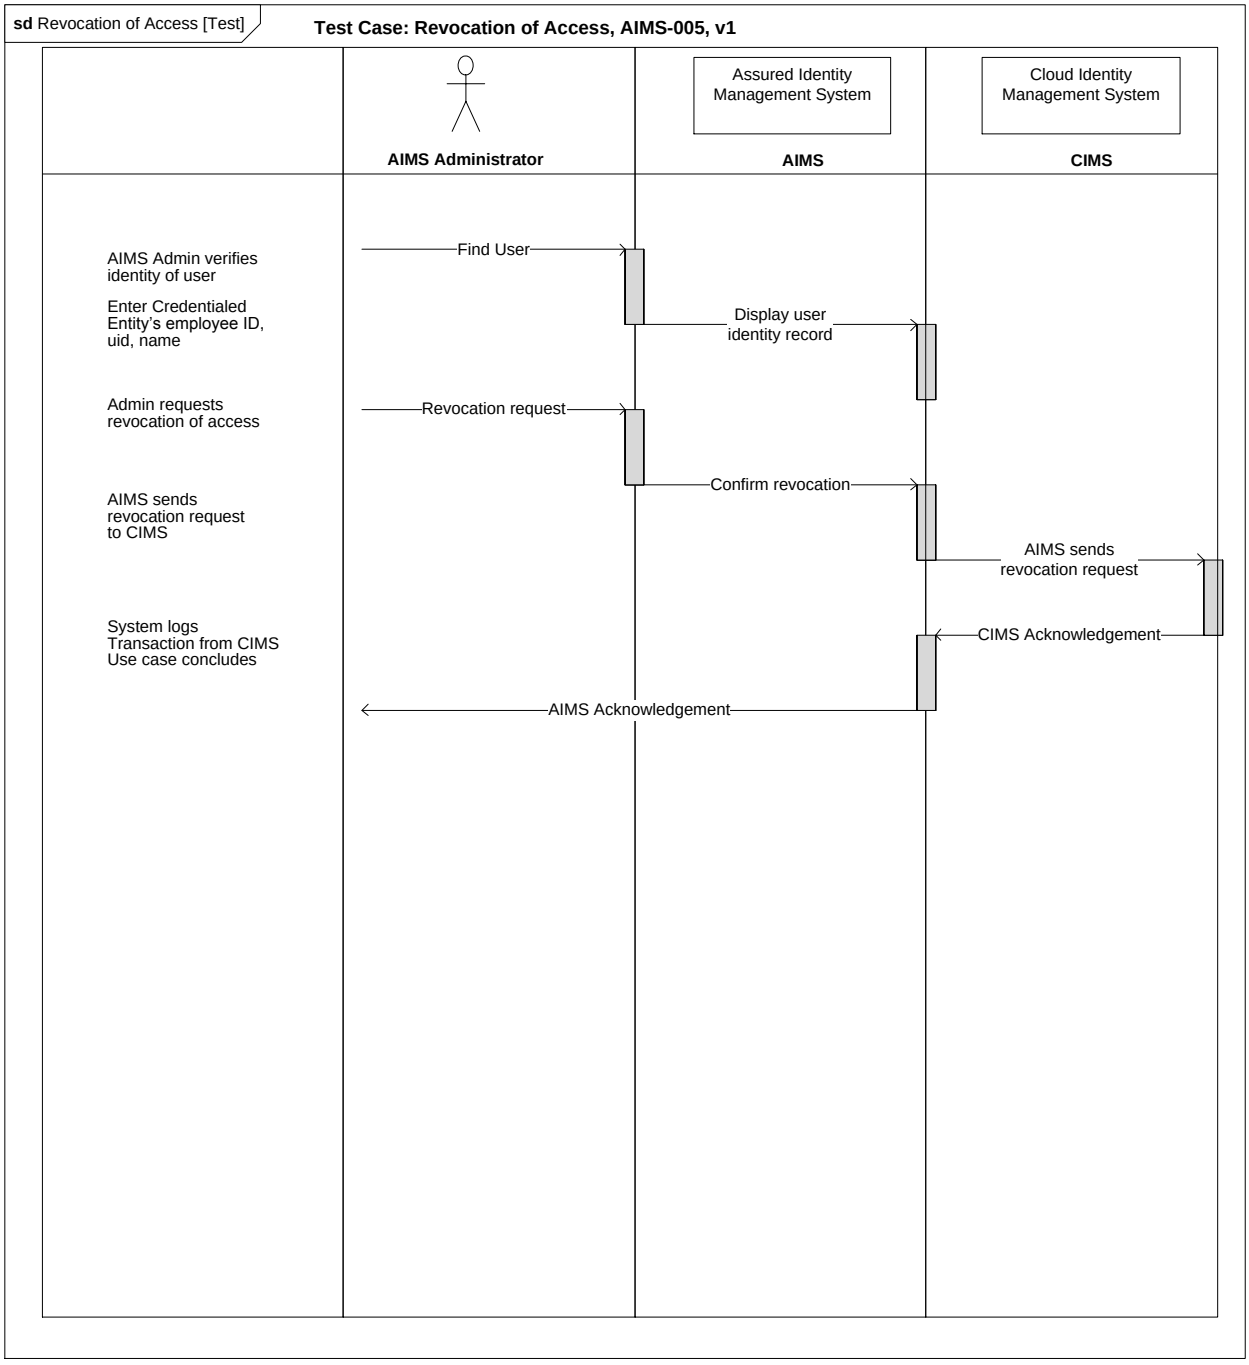

Figure 34.

Revocation of Access Sequence Diagram TC-005, SysML

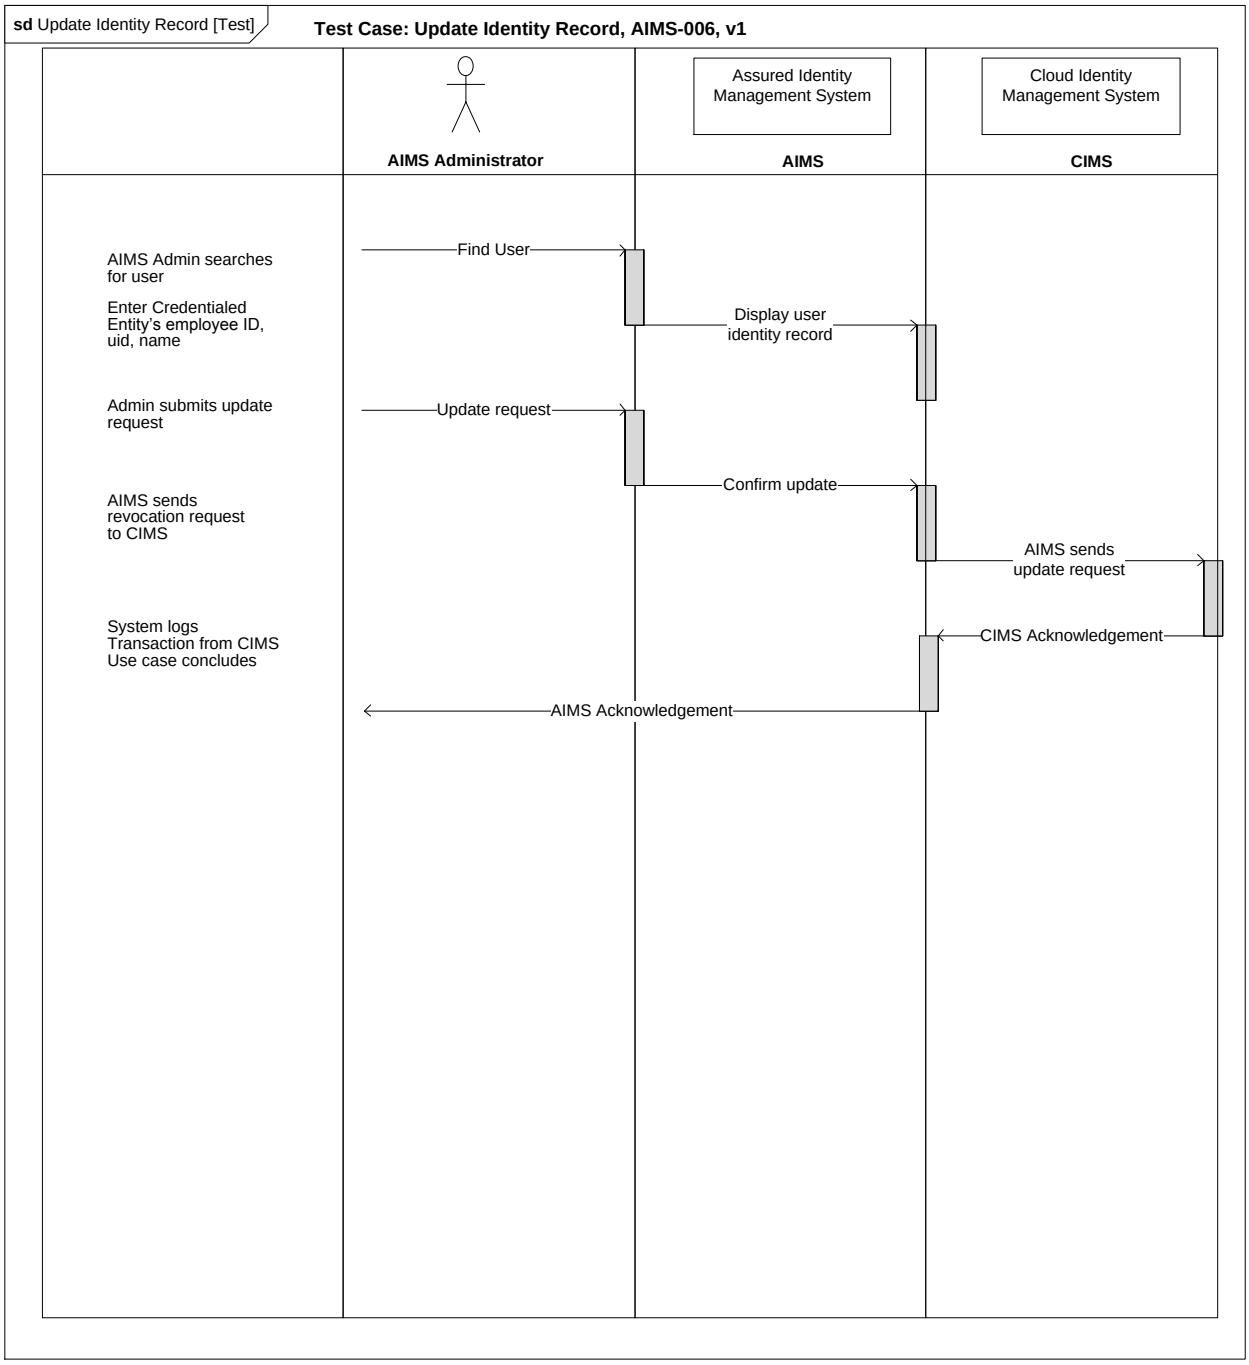

Figure 35.

Update Identity Record Sequence Diagram TC-006, SysML

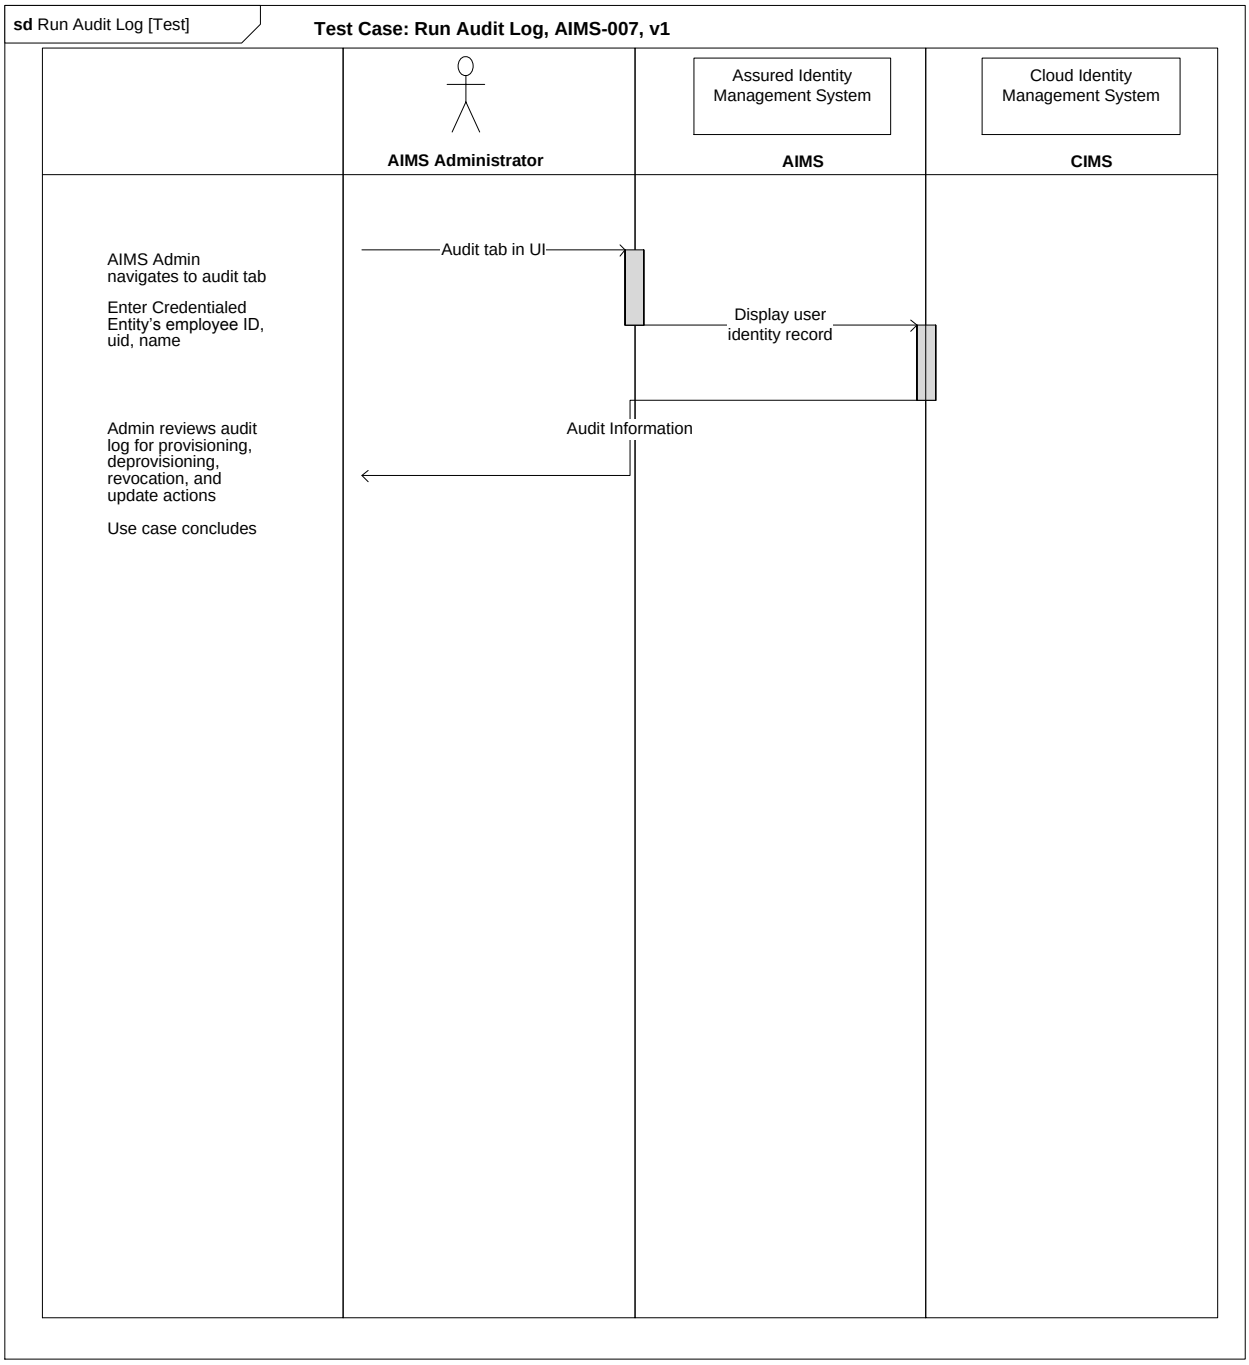

Figure 36.

Run Audit Log Sequence Diagram TC-007, SysML

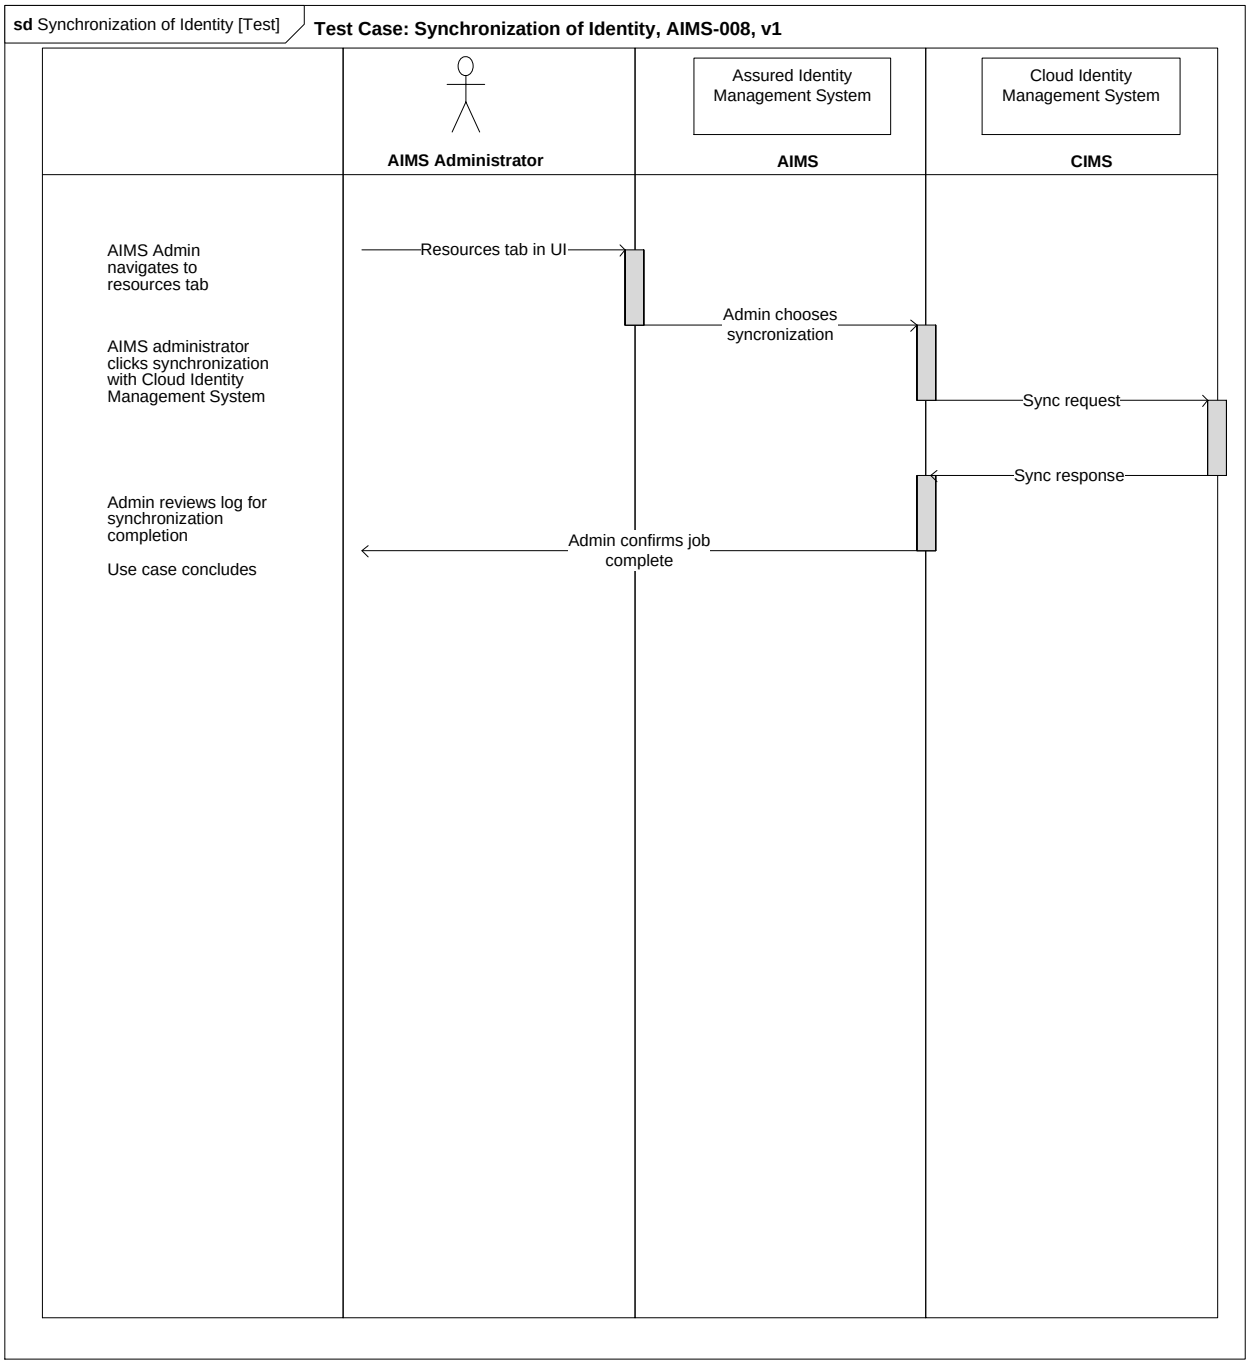

Figure 37.

Synchronization of Identity Sequence Diagram TC-008, SysML

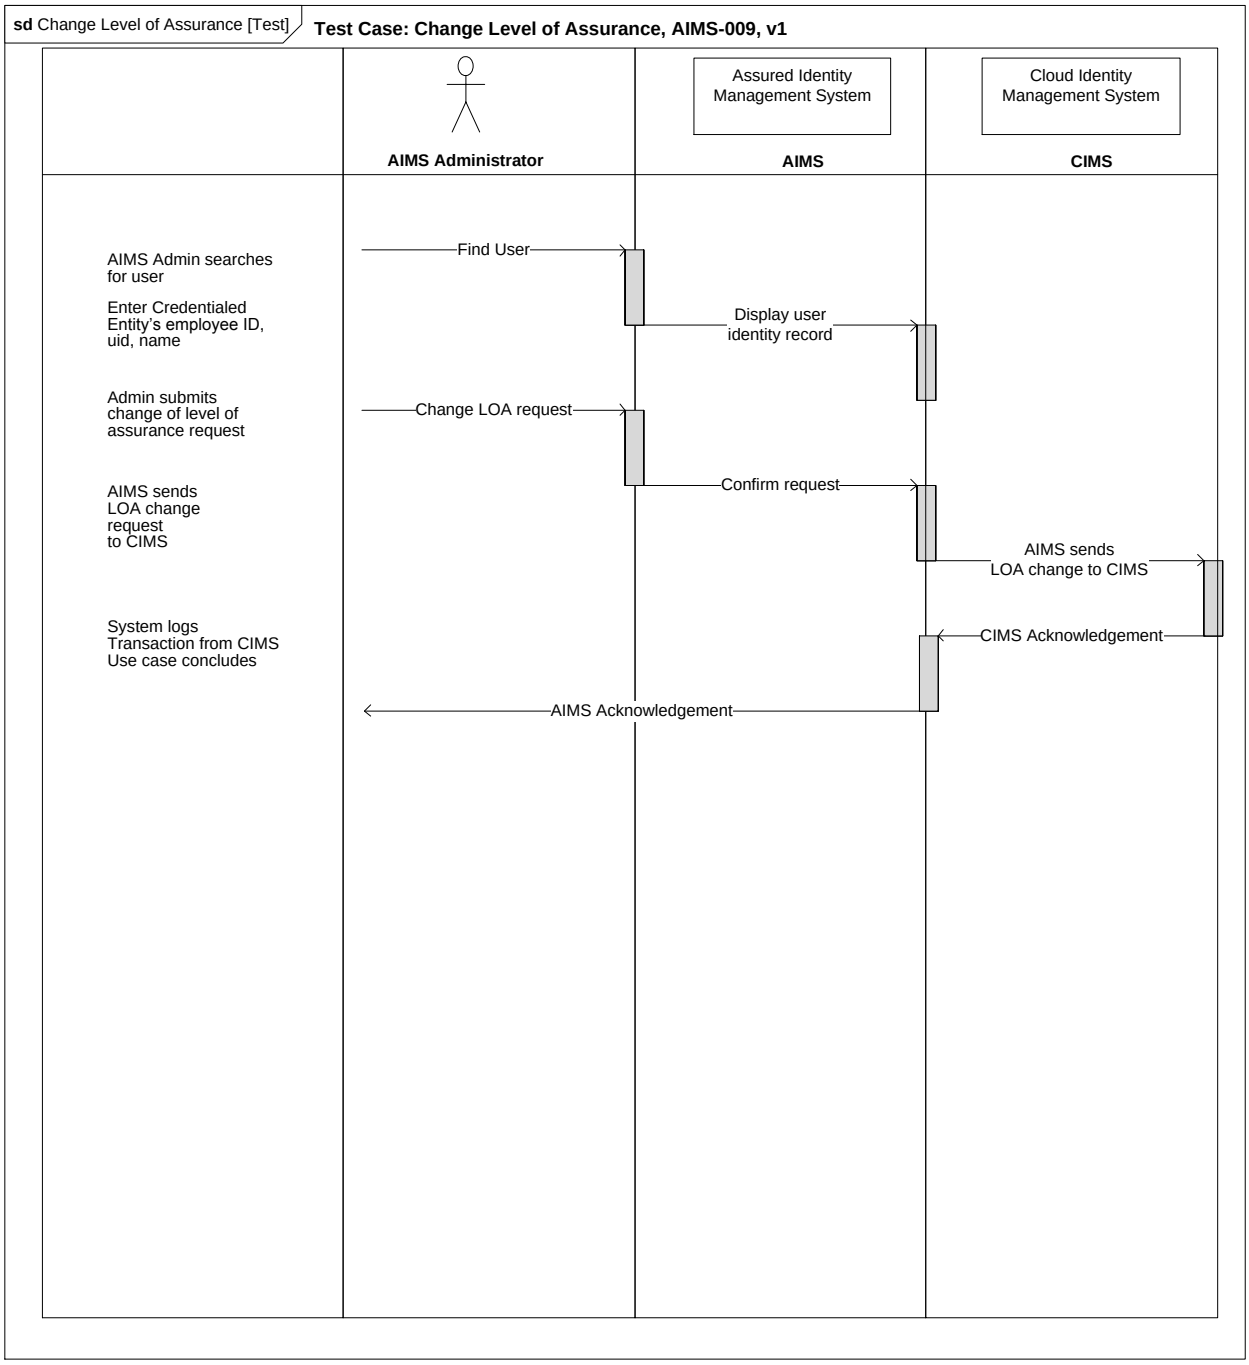

Figure 38.

Change Level of Assurance Sequence Diagram TC-009, SysML

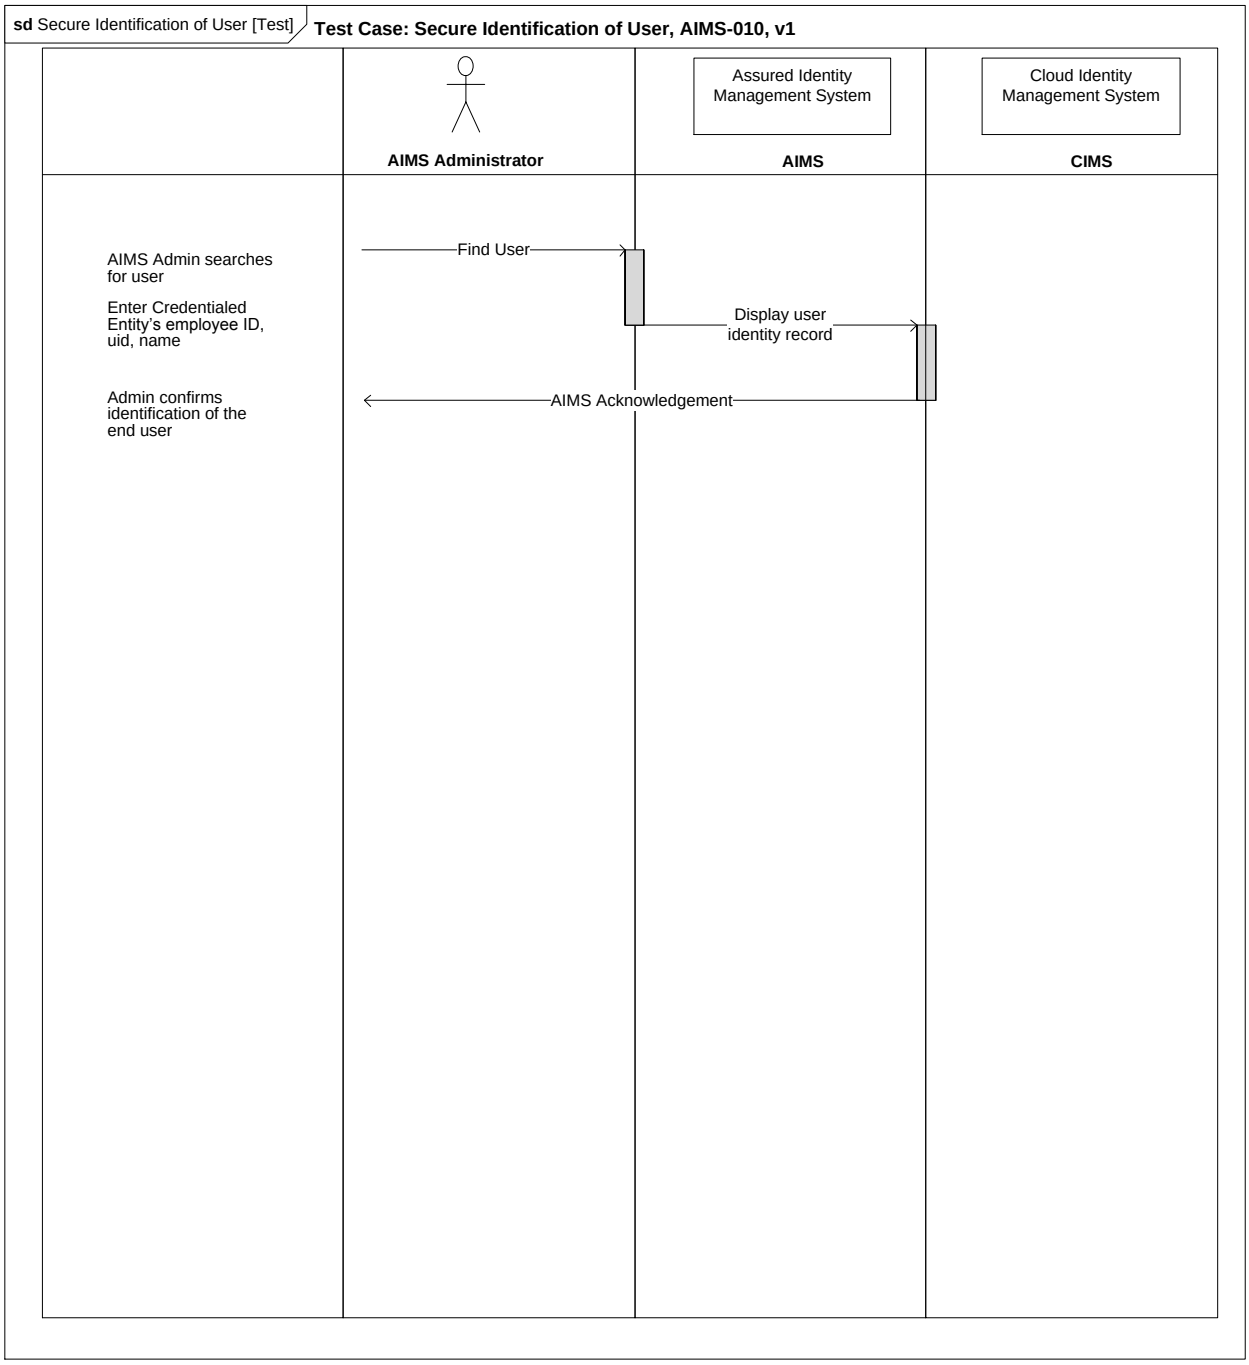

Figure 39.

Secure Identification of User Sequence Diagram TC-010, SysML

Constraints

A number of constraints were recognized within this investigation. SysML does not include timing or communications diagrams. Sequence diagrams provided a step-by step process the system actors followed for execution of the scenario. Swim lanes indicated the sequence of the actions, but could not account for the timing of each subsystems component. There is no timing diagram included in the SysML specification and traces to validate timing mechanism were not an option with the model. Timing diagrams were excluded from the specification due to "concerns about their maturity sustainability for systems engineering needs" (OMG, 2010).

The "on-premise" verification process was a constraint of this investigation. The model created for provisioning and identity assurance in the cloud requires "on-premise" verification for access to cloud environments. The scenario relied on the human resources function to verify and confirm identity as a condition of employment with the organization. Verification was done using the US Federal Government I-9 form which required a combination of identity credentials to be presented for confirmation of eligibility for employment. The I-9 form applied to U.S. citizens and non-citizens alike. In accordance with federal law, the employer must document the title, issuing authority, document number, expiration date, and the date of employment. I-9 records are kept by the employer and made available to the U.S. government for auditing purposes.

The investigation holds that the human resources organization is in the best position to confirm the identity of a subject as part of the new employment process (sometimes considered orientation). The I-9 form becomes part of an information package presented to the new employee. Other contents of the package might include job description, salary, benefits, and similar information. The I-9 form has three columns called lists A, B, and C (see Figure 40).

LISTS OF ACCEPTABLE DOCUMENTS

All documents must be unexpired

LIST A Documents that Establish Both Identity and Employment Authorization	OR	LIST B Documents that Establish Identity	AND	LIST C Documents that Establish Employment Authorization
1. U.S. Passport or U.S. Passport Card		1. Driver's license or ID card issued by a State or outlying possession of the United States provided it contains a photograph or information such as name, date of birth, gender, height, eye color, and address		1. Social Security Account Number card other than one that specifies on the face that the issuance of the card does not authorize employment in the United States
2. Permanent Resident Card or Alien Registration Receipt Card (Form I-551)				2. Certification of Birth Abroad issued by the Department of State (Form FS-545)
3. Foreign passport that contains a temporary I-551 stamp or temporary I-551 printed notation on a machine-readable immigrant visa		2. ID card issued by federal, state or local government agencies or entities, provided it contains a photograph or information such as name, date of birth, gender, height, eye color, and address		3. Certification of Report of Birth issued by the Department of State (Form DS-1350)
4. Employment Authorization Document that contains a photograph (Form I-766)		3. School ID card with a photograph		
		4. Voter's registration card		4. Original or certified copy of birth certificate issued by a State, county, municipal authority, or territory of the United States bearing an official seal
5. In the case of a nonimmigrant alien authorized to work for a specific employer incident to status, a foreign passport with Form I-94 or Form I-94A bearing the same name as the passport and containing an endorsement of the alien's nonimmigrant status, as long as the period of endorsement has not yet expired and the proposed employment is not in conflict with any restrictions or limitations identified on the form		5. U.S. Military card or draft record		
		6. Military dependent's ID card		
		7. U.S. Coast Guard Merchant Mariner Card		5. Native American tribal document
		8. Native American tribal document		6. U.S. Citizen ID Card (Form I-197)
		9. Driver's license issued by a Canadian government authority		
		For persons under age 18 who are unable to present a document listed above:		7. Identification Card for Use of Resident Citizen in the United States (Form I-179)
6. Passport from the Federated States of Micronesia (FSM) or the Republic of the Marshall Islands (RMI) with Form I-94 or Form I-94A indicating nonimmigrant admission under the Compact of Free Association Between the United States and the FSM or RMI		10. School record or report card		8. Employment authorization document issued by the Department of Homeland Security
		11. Clinic, doctor, or hospital record		
		12. Day-care or nursery school record		

Illustrations of many of these documents appear in Part 8 of the Handbook for Employers (M-274)

Figure 40.

I-9 List of Acceptable Documents, US Department of Homeland Security

List A has examples of documents that establish identity and employment authorization; List B includes documents that establish identity; List C documents that establish employment authorization. An employee may present a document from List A, or a combination of one from List B and on from List C. Regardless of the combination of credentials presented, the study maintains the need for a physical "on-premise" verification of identity as a method of assurance. From this procedure, the secure identity assurance flowed from human resources to AIMS and through CIMS.

The purpose of this investigation was to study identity assurance within cloud computing environments. The research project introduced fully dressed use cases based on the SysML modeling language. A prototype identity provisioning system called Assured Identity Management System (AIMS) was designed to evaluate the use cases. The investigation contributes to the body of knowledge by designing a new style of use cases focused on identity assurance in the cloud. There is now the ability for proposal, design, and research teams to take these use cases and import them into future implementations. Given that the models were created in SysML, the contents are extensible and can be modified or built upon to meet specific design requirements. The capability of SysML allowed the work packages to be shared among a variety of engineering tools and environments. The documented architecture framework, concept, requirements, and design artifacts allow the investigation to be implemented using these reusable components. In order to understand how identity assurance was implemented with a cloud-based system, ten use cases were developed to illustrate scenarios. The use cases were tested against the prototype environment using sequence diagrams.

Determining who is on the Cloud

To determine who is authorized on the cloud, the study focused on account provisioning methods between the cloud subscriber and the cloud provider. Identity assurance begins at the time of provisioning. Human resources personnel confirmed the identity of the subject using PIV guidance with the I-9 form and created an account in AIMS at the time of issuance. The account was provisioned in the cloud provider: AIMS (Assured Identity Management System) and then securely transmitted to the cloud subscriber system: CIMS (Cloud Identity Management System). In this manner the account controls were assured at issuance and not modified by the subscriber. The accounts in CIMS were created in AIMS and securely exchanged between the cloud systems. AIMS maintained the position as the record of authority for account and credentialing activities. The trust fabric was established at the initial identity credentialing phase with HR and carried through to CIMS. The AIMS framework solved the identity assurance problem by acting as a secure conduit from the organization through the cloud, while leveraging existing organizational structures (human resources), and federal policy and guidelines for identity verification.

Providing Assured Identity in the Cloud

The approach for this study combined processes, policy, virtual server components, and cloud computing systems to build a framework in which to manage identity records securely. Following the systems engineering process from concept to requirements, design, and implementation, the AIMS prototype system served as a foundation on which to provide identity assurance for an enterprise organization. The use case scenarios built in SysML are a key deliverable from the study. The analysis complete with actors and subsystem interactions, combined with the ability to visually depict identity lifecycle scenarios were critical to providing

identity assurance. The investigation identified a series of policies driven in part by the federal government that establish identity verification (PIV), identity processing (FIPS), and work authorization (I-9). These policies represented key items for identity assurance in the cloud from an auditing perspective. The technology deployed in this investigation provided the ability to manage identity data across the cloud through the use of standard protocols including SPML2, SOAP, XML, HTTPS, and TCP/IP.

Mechanisms for Identity Management and Access

The AIMS architecture reference model complete with artifacts and work packages offers a solution for providing assured identity for enterprise cloud computing environments. AIMS solved the problem by playing a pivotal role in providing a secure connection to exchange data between clouds throughout the identity lifecycle processes. From a design perspective the components were standards based and compatible with emerging cloud standards. From an operations and sustaining engineering perspective, virtualized systems deployed into the cloud were representative of common applications deployments in a platform-as-a-service (PAAS) model. The SysML deliverables allowed simulation in the creation of the use case. Through visual depiction, the data flows in terms of direction, frequency, and data type were analyzed for the communications path between actors and systems. The mechanisms used to build an assured identity management system included policy, process, architecture frameworks, requirements, and technical cloud components such as EC2, Ubuntu, SUSE, and OVF. These standards based technical mechanisms combined to create a secure solution for identity assurance.

Interoperability in the Global Enterprise

The systems in the investigation were chosen to be reflective of those found in global enterprise computing environments. Sun's Identity Manager platform was consistently in the Gartner leadership quadrant for identity and access management tools in recent years. Since the Sun-Oracle merger in 2010, the Identity Manager user community anticipates an upgrade path to the newly rebadged Oracle Waveset platform. Active Directory is a common component with the majority of the corporate directory services market share. The human resources component was simulated with a generic flat file feed, representative of any number of HR platforms including Oracle PeopleSoft, SAP Human Capital Management, or other brand. The Solaris operating system has been in enterprise data centers for many years. Windows 7 and the Ubuntu 8.1 operating systems are well known and used frequently by engineers, applications developers, and end-users. The virtualization and cloud tool set utilized Amazon EC2, a pioneer in cloud computing services. The OVF tool is an open source tool used to convert virtual disks to the OVF standard format. The identity and access mechanisms were intended to be reflective of enterprise computing environments with a mix of Linux, Solaris, Windows, Active Directory, and Sun Identity Manager (or similar SPML2-based identity management platform).

The interoperability and communications interfaces were designed to use standard TCP/IP connection methods to exchange data. Simple Object Access Protocol (SOAP) envelopes were used to manage the SPML2-based messages for provision, deprovision, and changes in identity record data. These standard technologies were compatible with a variety of compliant vendor applications. For example, Sun Identity Manager deployed on a cloud-based Ubuntu instance could be replaced with Novell Identity Manager deployed on a cloud-based SUSE instance. The two stacks have common interoperable components in terms of protocol,

communication methods, and compatibility. The system components and technologies deployed in this investigation were designed to be interoperable and portable with multiple industry standards and software configurations.

CHAPTER 5

DISCUSSION AND IMPLICATIONS

Recommendations for Practice

One of the contributions of this study is the development of fully dressed use cases. The computing and research community has recognized that cloud adoption needs a jump start. Created by Badger and Grance, the Standards Acceleration Jumpstarting Adoption of Cloud Computing (SAJACC) specifically identifies the use case method as a way to provide insight on how clouds can work (2010). Organizations under pressure to design and implement systems within tight development schedules and budgetary constraints may not have the resources to research how identities in the cloud interoperate, nor would they have time to explore the various security methods associate with identity management. One of the benefits of this study is the contribution is the end product as a set of portable, customizable design package of detailed use cases that may be used across business domain, industry, and project. Design teams can start with these use cases and modify to program specific requirements and security practices.

The series of use cases likely will not be a “one-for-one” fit with all organizations, but there are advantages to having a model to start from, even if the consumer applies use cases or sequence diagrams in piecemeal style. These artifacts would not be created from scratch or from blank templates, but from existing SysML artifacts provided in this study.

Choosing the right SysML design tools is a critical step early in the investigation. Practitioners and researchers will want to evaluate and perform due diligence for the design environments and tools available. There are many options, including the recommended applications available from the OMG SysML resource list. The investigator used three different tools to create the SysML models for this study. It would have been much easier and more efficient to choose a mature product and carry it through the life of the project. Future work in the tools area would include evaluating the capabilities and feature sets of various SysML design products. For example, this study used a mix of open source and COTS tools to generate the models. The tools varied in supportability and deployment model. Some were available though the Eclipse framework, others included the Eclipse framework bundled with the tool, and still others were standalone third party packages. Trade evaluations and comparisons on the availability and capability of SysML modeling tools in the enterprise is an area that could be studied.

Recommendations for Future Research

Building on the research conducted in this study, future researchers may investigate cloud adoption rates with and without the foundation of existing, customizable fully dressed use cases. Given the projections of cloud adopters over the next 2-3 years, and as identity assurance in the cloud matures, researchers should have a field with many candidate systems to evaluate. Hybrid cloud structures with an internal identity management component and an external cloud based component will become more common. Large organizations will carry forward legacy identity management and access systems and seek secure methods of assurance to build trusted relationships. An investigation that evaluates the benefits of starting with pre-existing extensible work packages against those starting from scratch would be helpful in promoting cloud adoption.

Design and implementation using a different applications suite is another direction for future studies. As mentioned in Chapter 4, Sun's Identity Manager is merging with the Oracle Waveset product line. The new release is due out in 2011; an evaluation of the capabilities in SPML2, SOAP, and web services adapter technology compared to the heritage Sun product provides an opportunity for further analysis. One major change is the work flow language XPRESS will be deprecated and the new standards-based language called Business Process Execution Language (BPEL) will replace it. Other platforms such as Novell's Identity Manager and IBM's Tivoli Federated Identity Manager are alternatives to the Sun-Oracle product lines. A follow-on study could take the SysML use case artifacts, requirements, and sequence diagrams and build a simulation of AIMS on these platforms. This would further demonstrate the transparency and portability of the prototype system and code base.

Conclusion

In summary, this research project designed and evaluated an approach to enabling assured identity in cloud computing environments. The AIMS use case framework provides a launch point for identity assurance in the cloud. The SysML artifacts including requirements, use cases, context diagrams, and sequence diagrams represent an *extensible* model intended to jumpstart systems design efforts. The framework serves the purpose of promoting *reusable* components to further the adoption of cloud computing as standards are created and implemented as indicated by NIST. The use case method provides insight into how clouds work and an identity assurance approach to managing account lifecycle processes in the cloud.

In addition to distributing identity attributes and meta-data, the AIMS framework provides a homogeneous security context for cloud systems to manage identities. Using SPML2 and the AIMS web service, a common trust fabric is shared across applications and services

deployed in the cloud. The policies governing AIMS present it as the record of authority for identity data.

The AIMS framework demonstrates qualities outlined by federal CIO Vivek Kundra in the Federal Cloud Computing Strategy: "platform strength resulting from greater uniformity and homogeneity, and resulting in *improved information assurance*, security response, system management, reliability, and maintainability" (2011). This investigation addresses one of the largest security concerns in cloud computing design and implementation: identity and access management. The research project contributes to the body of knowledge in systems management, security, cloud computing and virtualization.

REFERENCES

Ackoff, Russell L. (1962). *Scientific Method*. New York: John Wily & Sons. p6.

Advanced Information Assurance Handbook. (2004). Carnegie Mellon Software Engineering Institute. Pittsburgh, PA.

"After Bill." (28 June, 2008). *The Economist*, p. 77.

Aldrin, Buzz and Ken Abraham. (2009). *Magnificent Desolation: The Long Journey Home from the Moon*. New York: Random House, Inc.

Amazon, Inc. SEC Filing, Form 10-K. (28 January, 2010). Retrieved 10 September, 2010, from http://www.sec.gov/Archives/edgar/data/1018724/000119312510016098/0001193125-10-016098-index.htm

Armbrust, M. et al. (April, 2010). A view of cloud computing. *Communications of the ACM*, ACM 53, 4, 50–58.

Badger, Lee and Tim Grance. (20 May, 2010). Standards Acceleration to Jumpstart Adoption of Cloud Computing (SAJACC). Retrieved 16 September, 2010, from: http://csrc.nist.gov/groups/SNS/cloud-computing/documents/forumworkshop-may2010/nist_cloud_computing_forum-badger_grance.pdf

Barham, Paul, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew Warfield. (October 2003). Xen and the art of virtualization. *In Proceedings of 19th Symposium on Operating Systems Principles*. Bolton Landing, NY, USA. 164–177.

Blakely, Bob and Drue Reeves. (25 March, 2010). Defining Cloud Computing. Burton Group Management Briefing. Retrieved 25 September, 2010 from http://www.burtongroup.com/Client/Research/Document.aspx?cid=1951

Brynjolfsson, Erik, Paul Hofmann, and John Jordan. (May, 2010). "Cloud computing and electricity: beyond the utility model." Communications of the ACM Volume 53 Issue 5.

Brynjolfsson, E. and Saunders, A. (2010). *Wired for Innovation: How IT is Reshaping the Economy*. MIT Press, Cambridge, MA.

Christodorescu, M., Sailer, R., Schales, D. L., Sgandurra, D., and Zamboni, D. (2009). Cloud security is not (just) virtualization security: a short paper. In *Proceedings of the 2009 ACM workshop on Cloud computing security* (pp. 97-102). Chicago, Illinois, USA.

Cloud Computing Information Assurance Framework. (2009). European Network and Information Security Agency (ENISA).

Cloud Computing Use Cases Whitepaper. (02 July, 2010). Cloud Computing Use Case Discussion Group. Retrieved 12 September, 2010 from http://opencloudmanifesto.org/Cloud_Computing_Use_Cases_Whitepaper-4_0.pdf

Cloud Computing Manifesto. (2009). Retrieved 19 November, 2009, from http://www.opencloudmanifesto.org/

Cockburn, Alistair. Use cases, ten years later. Originally published in *STQE magazine*, Mar/Apr 2002. Retrieved 16 September, 2010, from http://alistair.cockburn.us/Use+cases%2c+ten+years+later

de Lemos, Rogério. (September 2008). *Architecting Dependable Systems V*, Volume 5. Springer-Verlag New York, LLC. ISBN 354085570X. 231.

Dashofy, Eric Matthew. (2007). "Supporting Stakeholder-driven, Multi-view Software Architecture Modeling." University of California, Irvine.

Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002. Official Journal L 201, 31/07/2002 P. 0037 - 0047. Retrieved 03 March, 2010, from http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32002L0058:EN:HTML

Directive 8500.1 Information Assurance. (24 October, 2002). Department of Defense. 20.

Durkee, Dave. 2010. "Why the Cloud Will Never Be Free." *ACM Queue Distributed Computing*. Communications of the ACM Vol. 53 No. 5. 62-69.

E-Authentication Guidance for Federal Agencies. (16 December, 2003). Memorandum to the heads of all departments and agencies. M-04-04. Office of Management and Budget (OMB).

Electronic Authentication Guideline. (08 December, 2008). National Institute of Standards and Technology (NIST) 153 Special Publication 800-63 version 1.0.2 [NIST800-63-1]. U.S. Department of Commerce.

Farber, Dan. Amazon's Blueprint for Cloud Computing. (25 June, 2008). Retrieved 12 October, 2010, from http://news.cnet.com/8301-13953_3-9977100-80.html?tag=mncol

Farber, Dan. Defining Cloud Computing. (07 May, 2008). Interview with Kevin Marks of Google. Retrieved 09 June, 2010, from http://news.cnet.com/8301-13953_3-9938949-80.html?tag=mncol%20video%20interview

Federal Chief Information Officer (CIO) Council. (February 2001). Practical Guide to Enterprise Architecture. Version 1.0,

Federal Information Security Management Act of 2002, E-Government Act of 2002. (2002). Retrieved 04 November, 2010, from http://www.gpo.gov/fdsys/pkg/PLAW-107publ347/content-detail.html

Friedenthal, Sanford, Alan Moore, and Rick Steiner. (2008). A Practical Guide to SysML: Systems Model Language. Morgan Kaufmann Publishers. Burlington, MA 01803.

Gardner, W. David. (April 15, 2010). "Enterprises See Risks In Cloud." *InformationWeek*. Retrieved 04 September, 2010, from http://www.informationweek.com/news/security/app-security/224400386

Gemmill, Jill. (2006). A trust-relationship management framework for federated virtual organizations. The University of Alabama at Birmingham, 141 pages. AAT 3226740.

Goldberg, R. P. (1973). Architecture of Virtual Machines. In Proceedings of AFIPS National Computer Conference. New York, New York

Google, Inc. SEC Filing, Form 10-Q. (30 June, 2009). Retrieved 10 September, 2010, from http://sec.gov/Archives/edgar/data/1288776/000119312509163845/d10q.htm#tx66132_12

Greene, Tim. (23 April, 2009). Cloud security stokes concerns at RSA. *Network World*. Retrieved 05 April, 2010, from: http://www.networkworld.com/news/2009/042309-rsa-cloud-security.html

Geurs, Karl (Director, Producer). (5 August, 1997). *Pooh's Grand Adventure: The Search for Christopher Robin* [Motion Picture]. United States of America: Walt Disney Television Animation

Guide to Protecting the Confidentiality of Personally Identifiable Information (PII). (April, 2010). National Institute of Standards and Technology (NIST) Special Publication 800-122 [NIST800-122]. U.S. Department of Commerce.

H.R. 2647: National Defense Authorization Act for Fiscal Year 2010. (28 October, 2009). Public Law. National Defense Authorization Act for Fiscal Year 2010. Retrieved 13 June, 2010, from http://www.govtrack.us/congress/bill.xpd?bill=h111-2647&tab=summary

Hamilton, J. (September 2008). "Internet-scale service efficiency." In Proceedings of the Large-Scale Distributed Systems and Middleware (LADIS) Workshop.

Hayes, James. (2009). Clout of the Cloud. *Engineering & Technology*. Vol. 4. 60-62.

Hunter, Philip. (09 October, 2009). Cloud Aloud. *Engineering and Technology*. Volume 4, Issue 16. ISSN:1750-9637. 54-56.

Identity Assurance Framework: Assurance Levels. (24 April, 2010). Version 2.0. Kantara Initiative.

INCOSE Systems Engineering Handbook. (June, 2004). Version 2a.

Kanaskar, N., Topaloglu, U., and Bayrak, C. (2005). Globus security model for grid environment. In Proceedings of ACM SIGSOFT Software Engineering Notes. 1-9.

Keller, John. (16 July, 2010). "Air Force to establish research center of excellence in assured cloud computing." *Military & Aerospace Electronics*.

Kolowich, Steve. (October 6, 2010). Mixing Work and Play on Facebook. *Inside Higher Ed*. Retrieved 07 October, 2010, from http://www.insidehighered.com/news/2010/10/06/facebook

Kundra, Vivek. (08 February, 2011). Federal Cloud Computing Strategy. Retrieved 15 February, 2011 from http://www.cio.gov/documents/Federal-Cloud-Computing-Strategy.pdf

Landau, Susan. (19 June, 2009). Privacy, Trust, and Security in Technology. Discussion with Sun Distinguished Engineer Susan Landau. System News. #21871. Retrieved 29 September, 2010 from https://sun.systemnews.com/articles/136/3/sec

Latham, Donald C. (2004). Vulnerabilities to Electromagnetic Attack of the Civil Infrastructure. In Jacques S. Gansler and Hans Binnendijk, Information assurance: trends in vulnerabilities, threats, and technologies. Washington DC: National Defense University Center for Technology and National Security Policy. 10

Lombardi, Flavio and Roberto Di Pietro. (07 June, 2010). Secure Virtualization for Cloud Computing. Journal of Network and Computer Applications. Volume 34, Issue 4, July 2011, Pages 1113-1122. doi:10.1016/j.jnca.2010.06.008

Maier, Mark. (1998-2007). The Art and Science of Systems Architecting. The Aerospace Corporation.

McLaughlin, Laurianne. (21 October, 2008). Cloud Computing Survey: IT Leaders See Big Promise, Have Big Security Questions. CIO Magazine. Retrieved 25 August, 2010, from http://www.cio.com/article/455832/Cloud_Computing_Survey_IT_Leaders_See_Big_Promise_Have_Big_Security_Questions?page=1&taxonomyId=3112

Mell, P. and Grance, T. (2009). The NIST Definition of Cloud Computing. U.S. Department of Commerce, National Institute of Standards and Technology, Information Technology Laboratory. Version 15, 10-7-09. Retrieved 08 July, 2010, from http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc

Mell, P. and Grance, T. (2009). Effectively and Securely Using the Cloud Computing Paradigm. U.S. Department of Commerce, National Institute of Standards and Technology, Information Technology Laboratory. 10-7-09. Retrieved 05 May, 2010 from http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-computing-v26.ppt

Messmer, Ellen. (10 June, 2010). "Identity Management Top Security Priority In Gartner Survey." Retrieved 17 August, 2010, from http://www.thestreet.com/story/10780483/identity-management-top-security-priority-in-gartner-survey.htm

Nagesh, Gautham. (26, November 2008). Local technology czar could be headed to Obama administration. Nextgov. Retrieved 12 October, 2010, from http://www.nextgov.com/nextgov/ng_20081126_1117.php .

O'Dell, Jolie. (11 June, 2010). Experts say we'll be working in the 'cloud' by 2020. CNN. Retrieved 18 September, 2010 from http://www.cnn.com/2010/TECH/web/06/11/cloud.mashable/index.html

OMG Systems Modeling Language (OMG SysML). (02 November, 2008). Retrieved 12 October, 2009, from http://www.omg.org/spec/SysML/1.1

Oo, May Phyo and Thinn Thu Naing. (2007). Access Control System for Grid Security Infrastructure. In Proceedings of the 2007 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology Workshops.

Oracle Technology Network Developer License Terms. (2009). Retrieved 30 November, 2010, from http://www.oracle.com/technetwork/testcontent/standard-license-088383.html

Pachner, Joanna. (01 March, 2010). "Cloud Computing Heavy with Hype." Canadian Business. Vol. 83 Issue 2, p13-14.

Park, Alvin R, and Brian Gammage. (13 October, 2005). "Microsoft Updates Server Licensing to Enable Virtualization." ID Number G00132810. Gartner Group. Stamford, CT.

Perry, Geva. (28 February, 2008). How Cloud Computing and Utility Computing are Different. Gigaom. Retrieved 26 September, 2010 from http://gigaom.com/2008/02/28/how-cloud-utility-computing-are-different/

Pew Internet & American Life Project. (September, 2008). Use of Cloud Computing Applications and Services. Retrieved 24 September, 2010 from http://www.pewinternet.org/~/media/Files/Reports/2008/PIP_Cloud.Memo.pdf.pdf

"Pre-Milestone A and Early-Phase Systems Engineering: A Retrospective Review and Benefits for Future Air Force Systems Acquisition." (2008). Air Force Studies Board, National Research Council of the National Academies, Washington DC, 2008. Original Publication Andres, Richard. "An Overview of the Acquisition Logistics," Defense Acquisition University, Fort Belvoir, VA, 2003.

"Private Clouds Showing up on IT's Agenda." (15 December, 2008). *CIO Magazine*. Retrieved 12 October, 2010, from http://www.cio.de/news/cio_worldnews/867008/index5.html

Ramo, Simon and Robin K. St.Clair. (1998). The Systems Approach: Fresh Solutions to Complex Problems through Combining Science and Practical Common Sense. Anaheim, CA: KNI, Inc. Retrieved 27 September, 2010, from http://www.incose.org/ProductsPubs/DOC/SystemsApproach.pdf

Ross, Jeanne W. (08 January 2007). "Enterprise Architecture as a Strategy." Center for Information Systems Research. MIT Sloan-CISR.

Salmon, John. (24 September, 2008). "Clouded in uncertainty – the legal pitfalls of cloud computing." *Computing Magazine*. Incisive Media. London. http://www.computing.co.uk/computing/features/2226701/clouded-uncertainty-4229153

Segaller, S. (1998). *Nerds: A Brief History of the Internet*. New York: TV Books.

Siegel, Del. (2010). Cloud Computing: A Free Technology Option to Promote Collaborative Learning. *Gifted Child Today*. Volume 33 n4. 41-45.

Slack, S. E. (31 March, 2009). Is There Value in Cloud Computing? Retrieved 12 October, 2010, from http://www.ibm.com/developerworks/architecture/library/ar-valuecloudcomputing//?S_TACT=105AGX01&S_CMP=HP

Srodawa, Ronald J and Lee A. Bates. (1973). An Efficient Virtual Machine Implementation. In Proceedings of AFIPS National Computer Conference. New York, New York.

Sun Microsystems, Inc. (2009). Introduction to Cloud Computing Architecture.

Talbot, David. (January 2009). Security in the Ether. *Technology Review*. Retrieved 20 September, 2010, from http://www.technologyreview.in/computing/24284/

Tufte, Edward. (January 2001). *The Visual Display of Quantitative Information*. Graphics Press. ISBN 0961392142.

Tucker, A. and Comay, D. (2004). Solaris Zones: Operating System Support for Server Consolidation. In Proceedings of Virtual Machine Research and Technology Symposium.

U.S. Department of Health and Human Services Homeland Security Presidential Directive-12 (HSPD-12). (30 January, 2007). HSPD-12 Standard Operating Procedures and Use Cases Program Office. Washington, D.C.

Varian, Melinda. (1997). "VM and the VM Community: Past, Present, and Future." 1997. Office of Computing and Information Technology. Princeton University. Princeton, NJ 08544 USA.

Valez, Jennyfer. (2009). *Cloud Computing: Not So Cloudy Anymore*. Frost and Sullivan.

Vouk, M. (2008). Cloud Computing: Issues, Research, and Implementations. *Journal of Computing and Information Technology – CIT*. 4, 235–246. doi:10.2498/cit.1001391

Waldrop, M. (Jan/Feb 2000). "Computing's Johnny Appleseed." *Technology Review.* Retrieved 13 October, 2010, from http://www.techreview.com/articles/jan00/waldrop.htm

Weapon Systems Acquisition Reform Act of 2009 (DTM 09-027) (WSARA). (22 May, 2009). Public Law 111-23, "Weapon Systems Acquisition Reform Act of 2009," Retrieved 10, June, 2010 from https://acc.dau.mil/wsara

Whitaker, A, M. Shaw, and S. Gribble. (September, 2002). Denali: A Scalable Isolation Kernel. In Proceedings of ACM SIGOPS European Workshop.

Wittmann, Art. (15 May, 2010). Our Maturing View of Cloud Computing. *Information Week Analytics*. Retrieved 02 September, 2010, from http://www.informationweek.com/news/cloud-computing/software/224701815