Draft Special Publication 800-144

National Institute of
Standards and Technology
U.S. Department of Commerce

Guidelines on Security and Privacy in Public Cloud Computing

Wayne Jansen
Timothy Grance

Draft NIST Special Publication

Guidelines on Security and Privacy in Public Cloud Computing

Wayne Jansen

Timothy Grance

C O M P U T E R S E C U R I T Y

Computer Security Division
Information Technology Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899-8930

January 2011



U.S. Department of Commerce
Gary Locke, Secretary

National Institute of Standards and Technology
Patrick D. Gallagher, Director

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analysis to advance the development and productive use of information technology. ITL's responsibilities include the development of technical, physical, administrative, and management standards and guidelines for the cost-effective security and privacy of sensitive unclassified information in Federal computer systems. This Special Publication discusses ITL's research, guidance, and outreach efforts in computer security, and its collaborative activities with industry, government, and academic organizations.

National Institute of Standards and Technology Draft Special Publication 800-144
60 pages (Jan. 2011)

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

Abstract

Cloud computing can and does mean different things to different people. The common characteristics most share are on-demand scalability of highly available and reliable pooled computing resources, secure access to metered services from nearly anywhere, and dislocation of data from inside to outside the organization. While aspects of these characteristics have been realized to a certain extent, cloud computing remains a work in progress. This publication provides an overview of the security and privacy challenges pertinent to public cloud computing and points out considerations organizations should take when outsourcing data, applications, and infrastructure to a public cloud environment.

Keywords: Cloud Computing; Computer Security and Privacy; Information Technology Outsourcing

Table of Contents

Executive Summary vi

1. Introduction 1

1.1 Authority 1
1.2 Purpose and Scope 1
1.3 Audience 1
1.4 Document Structure 2

2. Background 3

3. Public Cloud Services 7

3.1 Service Agreements 7
3.2 The Security Upside 8
3.3 The Security Downside 10

4. Key Security and Privacy Issues 13

4.1 Governance 13
4.2 Compliance 14
4.3 Trust 16
4.4 Architecture 19
4.5 Identity and Access Management 21
4.6 Software Isolation 22
4.7 Data Protection 24
4.8 Availability 25
4.9 Incident Response 28
4.10 Summary of Recommendations 28

5. Public Cloud Outsourcing 30

5.1 General Concerns 31
5.2 Preliminary Activities 33
5.3 Initiating and Coincident Activities 35
5.4 Concluding Activities 35
5.5 Summary of Recommendations 36

6. Conclusion 38

7. References 39

Appendix A—Acronyms 51

Appendix B—Online Resources 52

Executive Summary

Cloud computing has been defined by NIST as 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 cloud provider interaction [Mel09]. Cloud computing technologies can be implemented in a wide variety of architectures, under different service and deployment models, and can coexist with other technologies and software design approaches. The security challenges cloud computing presents, however, are formidable, especially for public clouds whose infrastructure and computational resources are owned by an outside party that sells those services to the general public.

The emergence of cloud computing promises to have far-reaching effects on the systems and networks of federal agencies and other organizations. Many of the features that make cloud computing attractive, however, can also be at odds with traditional security models and controls. The primary purpose of this report is to provide an overview of public cloud computing and the security and privacy considerations involved. More specifically, this document describes the threats, technology risks, and safeguards surrounding public cloud environments, and their treatment.

The key guidelines from the report are summarized and listed below and are recommended to federal departments and agencies.

Carefully plan the security and privacy aspects of cloud computing solutions before engaging them.

As with any emerging information technology area, cloud computing should be approached carefully with due consideration to the sensitivity of data. Planning helps to ensure that the computing environment is as secure as possible and is in compliance with all relevant organizational policies and that data privacy is maintained. It also helps to ensure that the agency derives full benefit from information technology spending.

The security objectives of an organization are a key factor for decisions about outsourcing information technology services and, in particular, for decisions about transitioning organizational data, applications, and other resources to a public cloud computing environment. The information technology governance practices of the organizations that pertain to the policies, procedures, and standards used for application development and service provisioning, as well as the design, implementation, testing, and monitoring of deployed or engaged services, should be extended to cloud computing environments.

To maximize effectiveness and minimize costs, security and privacy must be considered from the initial planning stage at the start of the systems development life cycle. Attempting to address security after implementation and deployment is not only much more difficult and expensive, but also more risky.

Understand the public cloud computing environment offered by the cloud provider and ensure that a cloud computing solution satisfies organizational security and privacy requirements.

Cloud providers are generally not aware of a specific organization's security and privacy needs. Adjustments to the cloud computing environment may be warranted to meet an organization's requirements. Organizations should require that any selected public cloud computing solution is configured, deployed, and managed to meet their security, privacy, and other requirements.

Non-negotiable service agreements in which the terms of service are prescribed completely by the cloud provider are generally the norm in public cloud computing. Negotiated service agreements are also possible. Similar to traditional information technology outsourcing contracts used by agencies, negotiated agreements can address an organization's concerns about security and privacy details, such as the vetting of employees, data ownership and exit rights, isolation of tenant applications, data encryption and segregation, tracking and reporting service effectiveness, compliance with laws and regulations, and the use of validated products meeting federal or national standards (e.g., Federal Information Processing Standard 140).

Critical data and applications may require an agency to undertake a negotiated service agreement in order to use a public cloud. Points of negotiation can negatively affect the economies of scale that a non-negotiable service agreement brings to public cloud computing, however, making a negotiated alternative less cost effective. As an alternative, the organization may be able to employ compensating controls to work around identified shortcomings in the public cloud service. Other alternatives include cloud computing environments with a more suitable deployment model, such as a private cloud, which offers an organization greater oversight and control over security and privacy.

Ensure that the client-side computing environment meets organizational security and privacy requirements for cloud computing.

Cloud computing encompasses both a server and a client side. With emphasis typically placed on the former, the latter can be easily overlooked. Maintaining physical and logical security over clients can be troublesome, especially with embedded mobile devices such as smart phones. Their size and portability can result in the loss of physical control. Built-in security mechanisms often go unused or can be overcome or circumvented without difficulty by a knowledgeable party to gain control over the device.

Because of their ubiquity, Web browsers are a key element for client-side access to cloud computing services. Clients may also entail small lightweight applications that run on desktop and mobile devices to access services. The various available plug-ins and extensions for Web browsers are notorious for their security problems. Many browser add-ons also do not provide automatic updates, increasing the persistence of any existing vulnerabilities. Similar problems exist for other types of clients.

The increased availability and use of social media, personal Webmail, and other publicly available sites are a concern, since they can negatively impact the security of the client, its underlying platform, and cloud services accessed, through social engineering attacks. Having a backdoor Trojan, keystroke logger, or other type of malware running on a client device

undermines the security of cloud or other Web-based services. As part of the overall cloud computing security architecture, organizations should review existing measures and employ additional ones, if necessary, to secure the client side.

Maintain accountability over the privacy and security of data and applications implemented and deployed in public cloud computing environments.

Organizations should employ appropriate security management practices and controls over cloud computing. Strong management practices are essential for operating and maintaining a secure cloud computing solution. Security and privacy practices entail monitoring the organization's information system assets and assessing the implementation of policies, standards, procedures, and guidelines that are used to establish and preserve the confidentiality, integrity, and availability of information system resources.

Assessing and managing risk in cloud computing systems can be a challenge. Both qualitative and quantitative factors apply in a risk analysis. Risks must be carefully weighed against the available technical, management, and operational safeguards and the necessary steps must be taken to reduce risk to an acceptable level. The organization must also ensure that security and privacy controls are implemented correctly, operate as intended, and meet its requirements.

Establishing a level of confidence about a cloud service environment depends on the ability of the cloud provider to provision the security controls necessary to protect the organization's data and applications, and also the evidence provided about the effectiveness of those controls [Jtf10]. Verifying the correct functioning of a subsystem and the effectiveness of security controls as extensively as with an organizational system may not be feasible in some cases, however, and other factors such as third-party audits may be used to establish a level of trust. Ultimately, if the level of confidence in the service falls below expectations and the organization is unable to employ compensating controls, it must either reject the service or accept a greater degree of risk.

In general, organizations should have security controls in place for cloud-based applications that are commensurate with or surpass those used if the applications were deployed in-house. Cloud computing is heavily dependent on the individual security of each of its many components, including those for self-service, quota management, and resource metering, plus the hypervisor, guest virtual machines, supporting middleware, deployed applications, and data storage. Many of the simplified interfaces and service abstractions belie the inherent complexity that affects security. Organizations should ensure to the extent practical that all of these elements are secure and that security is maintained based on sound security practices.

1. Introduction

Interest in cloud computing has rapidly grown in recent years due to the advantages of greater flexibility and availability of computing resources at lower cost. Security and privacy, however, are a concern for agencies and organizations considering migrating applications to public cloud computing environments, and form the impetus behind this document.

1.1 Authority

The National Institute of Standards and Technology (NIST) developed this document in furtherance of its statutory responsibilities under the Federal Information Security Management Act (FISMA) of 2002, Public Law 107-347.

NIST is responsible for developing standards and guidelines, including minimum requirements, for providing adequate information security for all agency operations and assets; but such standards and guidelines shall not apply to national security systems. This guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130, Section 8b(3), "Securing Agency Information Systems," as analyzed in A-130, Appendix IV: Analysis of Key Sections. Supplemental information is provided in A-130, Appendix III.

This guideline has been prepared for use by Federal agencies. It may be used by nongovernmental organizations on a voluntary basis and is not subject to copyright, though attribution is desired.

Nothing in this document should be taken to contradict standards and guidelines made mandatory and binding on Federal agencies by the Secretary of Commerce under statutory authority, nor should these guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or any other Federal official.

1.2 Purpose and Scope

The purpose of this document is to provide an overview of public cloud computing and the security and privacy challenges involved. The document discusses the threats, technology risks, and safeguards for public cloud environments, and provides the insight needed to make informed information technology decisions on their treatment.

1.3 Audience

The intended audience for this document includes the following categories of individuals:

- System managers, executives, and information officers making decisions about cloud computing initiatives
- Security professionals, including security officers, security administrators, auditors, and others with responsibility for information technology security
- Information technology program managers concerned with security and privacy measures for cloud computing

- System and network administrators
- Users of public cloud computing services.

This document, while technical in nature, provides background information to help readers understand the topics that are covered. The material presumes that readers have some minimal operating system and networking expertise and a basic understanding of cloud computing. Because of the evolving nature of security and privacy considerations in cloud computing, readers are expected to take advantage of other resources for more detailed and current information. These resources include the various publications listed or referenced in this document, the majority of which are available on-line.

1.4 Document Structure

The remainder of this document is organized into the following chapters:

- Chapter 2 presents an overview of public cloud computing.
- Chapter 3 discusses the benefits and drawbacks of public cloud services from a security and privacy perspective.
- Chapter 4 discusses key security and privacy issues in public cloud computing and precautions that can be taken to mitigate them.
- Chapter 5 provides guidance on addressing security and privacy issues when outsourcing support for data and applications to a cloud provider.
- Chapter 6 presents a short conclusion.
- Chapter 7 contains a list of references.

The document also has appendices that contain supporting material: A list of acronyms is given in Appendix A and a list of other resources can be found in Appendix B.

2. Background

Cloud computing has been defined by NIST as 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 cloud provider interaction [Mel09]. Cloud computing can be considered a new computing paradigm insofar as it allows the utilization of a computing infrastructure at one or more levels of abstraction, as an on-demand service made available over the Internet or other computer network. Because of the implications for greater flexibility and availability at lower cost, cloud computing is a subject that has been receiving a good deal of attention lately.

Cloud computing services benefit from economies of scale achieved through versatile use of resources, specialization, and other practicable efficiencies. However, cloud computing is an emerging form of distributed computing that is still in its infancy. The term itself is often used today with a range of meanings and interpretations [Fow09]. Much of what has been written about cloud computing is definitional, aimed at identifying important paradigms of use and providing a general taxonomy for conceptualizing important facets of service.

Public cloud computing is one of several deployment models that have been defined. A public cloud is one in which the infrastructure and other computational resources that it comprises are made available to the general public over the Internet. It is owned by a cloud provider selling cloud services and, by definition, is external to an organization. At the other end of the spectrum are private clouds. A private cloud is one in which the computing environment is operated exclusively for an organization. It may be managed either by the organization or a third party, and may be hosted within the organization's data center or outside of it. A private cloud gives the organization greater control over the infrastructure and computational resources than does a public cloud.

Two other deployment models that fall between public and private clouds are community clouds and hybrid clouds. A community cloud is somewhat similar to a private cloud, but the infrastructure and computational resources are shared by several organizations that have common privacy, security, and regulatory considerations, rather than for the exclusive use of a single organization. A hybrid cloud is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables interoperability.

Just as the different deployment models affect an organization's scope and control over the computational environment of a cloud, so too does the service model supported by the cloud affect them. Three well-known and frequently-used service models are the following [Lea09, Vaq09, You08]:

- **Software-as-a-Service.** Software-as-a-Service (SaaS) is a model of software deployment whereby one or more applications and the computational resources to run them are provided for use on demand as a turnkey service. Its main purpose is to reduce the total cost of hardware and software development, maintenance, and operations. Security

provisions are carried out mainly by the cloud provider. The cloud subscriber does not manage or control the underlying cloud infrastructure or individual applications, except for preference selections and limited administrative application settings.

- **Platform-as-a-Service.** Platform-as-a-Service (PaaS) is a model of software deployment whereby the computing platform is provided as an on-demand service upon which applications can be developed and deployed. Its main purpose is to reduce the cost and complexity of buying, housing, and managing the underlying hardware and software components of the platform, including any needed program and database development tools. The development environment is typically special purpose, determined by the cloud provider and tailored to the design and architecture of its platform. The cloud subscriber has control over applications and application environment settings of the platform. Security provisions are split between the cloud provider and the cloud subscriber.

- **Infrastructure-as-a-Service.** Infrastructure-as-a-Service (IaaS) is a model of software deployment whereby the basic computing infrastructure of servers, software, and network equipment is provided as an on-demand service upon which a platform to develop and execute applications can be established. Its main purpose is to avoid purchasing, housing, and managing the basic hardware and software infrastructure components, and instead obtain those resources as virtualized objects controllable via a service interface. The cloud subscriber generally has broad freedom to choose the operating system and development environment to be hosted. Security provisions beyond the basic infrastructure are carried out mainly by the cloud subscriber.

Figure 1 illustrates the differences in scope and control between the cloud subscriber and cloud provider, for each of the service models discussed above. Five conceptual layers of a generalized cloud environment are identified in the center diagram and apply to public clouds, as well as each of the other deployment models. The arrows at the left and right of the diagram denote the approximate range of the cloud provider's and user's scope and control over the cloud environment for each service model. In general, the higher the level of support available from a cloud provider, the more narrow the scope and control the cloud subscriber has over the system.

The two lowest layers shown denote the physical elements of a cloud environment, which are under the full control of the cloud provider, regardless of the service model. Heating, ventilation, air conditioning (HVAC), power, communications, and other aspects of the physical plant comprise the lowest layer, the facility layer, while computers, network and storage components, and other physical computing infrastructure elements comprise the hardware layer.

The remaining layers denote the logical elements of a cloud environment. The virtualized infrastructure layer entails software elements, such as hypervisors, virtual machines, virtual data storage, and supporting middleware components used to realize the infrastructure upon which a computing platform can be established. While virtual machine technology is commonly used at this layer, other means of providing the necessary software abstractions are not precluded. Similarly, the platform architecture layer entails compilers, libraries, utilities, and other software tools and development environments needed to implement applications. The application layer

represents deployed software applications targeted towards end-user software clients or other programs, and made available via the cloud.

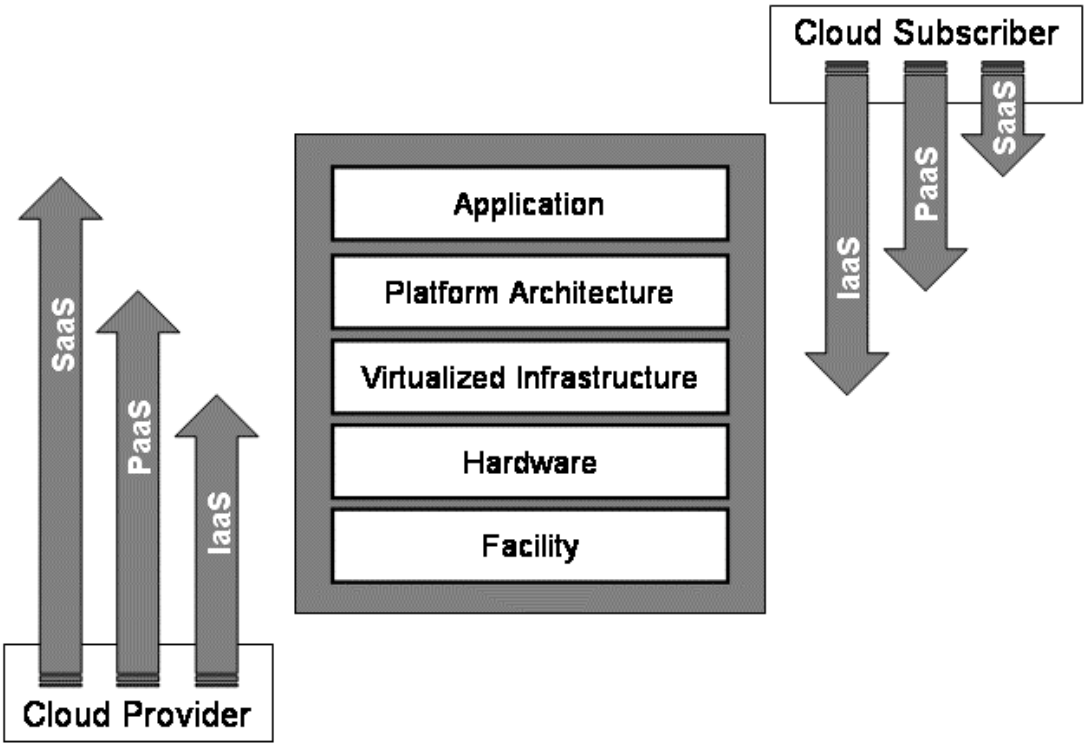


Figure 1: Differences in Scope and Control among Cloud Service Models

Some have argued that the distinction between IaaS and PaaS is fuzzy, and in many commercial offerings, the two are more alike than different [Arm10]. Nevertheless, these terms do serve a purpose, distinguishing between very basic support environments and environments having greater levels of support, and accordingly different allocations of control and responsibility between the cloud subscriber and the cloud provider.

While cloud computing can be implemented exclusively for an organization as a private internal cloud, its main thrust has been to provide a vehicle for outsourcing parts of that environment to an outside party as a public cloud. As with any outsourcing of information technology services, concerns exist about the implications for computer security and privacy. The main issue centers on the risks associated with moving important applications or data from within the confines of the organization's computing center to that of another organization (i.e., a public cloud), which is readily accessible by the general public.

Reducing cost and increasing efficiency are primary motivations for moving towards a public cloud, but reducing responsibility for security should not be. Ultimately, the organization is accountable for the overall security of the outsourced service. Monitoring and addressing security issues that arise remain in the purview of the organization, as does oversight over other important issues such as performance and availability. Because cloud computing brings with it

new security challenges, it is essential for an organization to oversee and manage how the cloud provider secures and maintains the computing environment and ensures data is kept secure.

3. Public Cloud Services

The outlook on cloud computing services can vary significantly among organizations, because of inherent differences in such things as their purpose, assets held, exposure to the public, threats faced, and tolerance to risk. For example, a government organization that mainly handles data about individual citizens of the country has different security objectives than a government organization that does not. Similarly, the security objectives of a government organization that prepares and disseminates information for public consumption are different from one that deals mainly with classified information for its own internal use. From a risk perspective, determining the suitability of cloud services for an organization is not possible without understanding the context in which the organization operates and the consequences from the plausible threats it faces.

The set of security objectives of an organization, therefore, is a key factor for decisions about outsourcing information technology services and, in particular, for decisions about transitioning organizational resources to public cloud computing and a specific provider's services and service arrangements. What works for one organization may not necessarily work for another. In addition, practical considerations apply—most organizations cannot afford financially to protect all computational resources and assets at the highest degree possible and must prioritize available options based on cost as well as criticality and sensitivity. When considering the potential benefits of public cloud computing, it is important to keep the organizational security objectives in mind and to act accordingly. Ultimately, a decision on cloud computing rests on a risk analysis of the tradeoffs involved.

3.1 Service Agreements

Specifications for public cloud services and service arrangements are generally called Service Level Agreements (SLAs). An SLA represents the understanding between the cloud subscriber and cloud provider about the expected level of service to be delivered and, in the event that the provider fails to deliver the service at the level specified, the compensation available to the cloud subscriber. An SLA, however, typically forms only a part of the terms of service stipulated in the overall service contract or service agreement. The terms of service cover other important details such as licensing of services, criteria for acceptable use, service suspension and termination, limitations on liability, privacy policy, and modifications to the terms of service. For the purpose of this report, the term SLA is used to refer to the service contract in its entirety.[1]

Two types of SLAs exist: predefined non-negotiable agreements and negotiated agreements [Bra10, UCG10]. Non-negotiable agreements are in many ways the basis for the economies of scale enjoyed by public cloud computing. Not only are the terms of service prescribed completely by the cloud provider, but with some offerings, the provider can also make modifications to the terms unilaterally without giving any direct notification to the cloud subscriber (e.g., by posting an updated version online) [Bra10]. Negotiated SLAs are more like

[1] In some contexts, an SLA is used to mean the overall service agreement or contract and not merely a subset (e.g., [Kan09]).

traditional information technology outsourcing contracts. They can be used to address an organization's concerns about security and privacy policy, procedures, and technical controls, such as the vetting of employees, data ownership and exit rights, isolation of tenant applications, data encryption and segregation, tracking and reporting service effectiveness, compliance with laws and regulations (e.g., Federal Information Security Management Act), and the use of validated products meeting national or international standards (e.g., Federal Information Processing Standard 140-2 for cryptographic modules).

Critical data and applications may require an agency to undertake a negotiated SLA [Wall0]. Since points of negotiation can significantly perturb and negatively affect the economies of scale that a non-negotiable SLA brings to public cloud computing, a negotiated SLA is normally less cost effective. The outcome of a negotiation is also dependent on the size of the organization and the influence it can exert. Regardless of the type of SLA, obtaining adequate legal and technical advice is recommended to ensure that the terms of service adequately meet the needs of the organization.

3.2 The Security Upside

While the biggest obstacle facing public cloud computing is security, the cloud computing paradigm provides opportunities for innovation in provisioning security services that hold the prospect of improving the overall security of some organizations. The biggest beneficiaries are likely to be smaller organizations that have limited numbers of information technology administrators and security personnel, and lack the economies of scale available to larger organizations with sizeable data centers.

Potential areas of improvement where organizations may derive security benefits from transitioning to a public cloud computing environment include the following:

- **Staff Specialization.** Cloud providers, just as organizations with large-scale computing facilities, have an opportunity for staff to specialize in security, privacy, and other areas of high interest and concern to the organization. Increases in the scale of computing induce specialization, which in turn allows security staff to shed other duties and concentrate exclusively on security issues. Through increased specialization, there is an opportunity for staff members gain in-depth experience, take remedial actions, and make security improvements more readily than otherwise would be possible with a diverse set of duties.

- **Platform Strength.** The structure of cloud computing platforms is typically more uniform than that of most traditional computing centers. Greater uniformity and homogeneity facilitate platform hardening and enable better automation of security management activities like configuration control, vulnerability testing, security audits, and security patching of platform components. Information assurance and security response activities also profit from a uniform, homogeneous cloud infrastructure, as do system management activities, such as fault management, load balancing, and system maintenance. Many cloud providers meet standards for operational compliance and certification in areas like healthcare (e.g., Health Insurance Portability and Accountability

Act (HIPAA)), finance (e.g., Payment Card Industry Data Security Standard (PCI DSS)) and audit (e.g., Statement on Auditing Standards No. 70 (SAS 70)).

- **Resource Availability.** The scalability of cloud computing facilities allows for greater availability. Redundancy and disaster recovery capabilities are built into cloud computing environments and on-demand resource capacity can be used for better resilience when facing increased service demands or distributed denial of service attacks, and for quicker recovery from serious incidents. When an incident occurs, an opportunity also exists to capture information more readily, with greater detail and less impact on production. In some cases, however, such resiliency can have a downside. For example, an unsuccessful distributed denial of service attack can quickly consume large amounts of resources to defend against and cause charges to soar, inflicting serious financial damage to an organization.

- **Backup and Recovery.** The backup and recovery policies and procedures of a cloud service may be superior to those of the organization and, if copies are maintained in diverse geographic locations, may be more robust. Data maintained within a cloud can be more available, faster to restore, and more reliable in many circumstances than that maintained in a traditional data center. Under such conditions, cloud services could also serve as a means for offsite backup storage for an organization's data center, in lieu of more traditional tape-based offsite storage [Kum08]. However, network performance over the Internet and the amount of data involved are limiting factors that can affect restoration.

- **Mobile Endpoints.** The architecture of a cloud solution extends to the client at the service endpoint, used to access hosted applications. Cloud clients can be browser-based or applications-based. Since the main computational resources needed are held by the cloud provider, clients are generally lightweight computationally and easily supported on laptops, notebooks, and netbooks, as well as embedded devices such as smart phones, tablets, and personal digital assistants.[2]

- **Data Concentration.** Data maintained and processed in the cloud can present less of a risk to an organization with a mobile workforce than having that data dispersed on portable computers or removable media out in the field, where theft and loss of devices routinely occur. Many organizations have already made the transition to support access to organizational data from mobile devices to improve workflow management and gain other operational efficiencies.

Besides providing a computing platform or substitute for in-house applications, public cloud services, such as the following, can also be focused on provisioning security to other computing environments:

[2] While not a security benefit per se, this relates to the next bulleted item.

- **Data Center Oriented.** Cloud services can be used to improve the security of data centers. For example, electronic mail can be redirected to a cloud provider via mail exchange (MX) records, examined and analyzed collectively with similar transactions from other data centers to discover widespread spam, phishing, and malware campaigns, and to carry out remedial action (e.g., quarantining suspect messages and content) more comprehensively than a single organization would be able to do. Researchers have also successfully demonstrated a system architecture for provisioning cloud-based antivirus services, as an alternative to host-based antivirus solutions [Obe08b].

- **Cloud Oriented.** Cloud services are available to improve the security of other cloud environments. For example, reverse proxy products are available that enable unfettered access to a SaaS environment, yet maintain the data stored in that environment in encrypted form [Nav10]. Cloud-based identity management services also exist, which can be used to augment or replace an organization's directory service for identification and authentication of users to a cloud.

3.3 The Security Downside

Besides its many potential benefits for security and privacy, public cloud computing also brings with it potential areas of concern, when compared with computing environments found in traditional data centers. Some of the more fundamental concerns include the following:

- **System Complexity.** A public cloud computing environment is extremely complex compared with that of a traditional data center. Many components comprise a public cloud, resulting in a large attack surface. Besides components for general computing, such as deployed applications, virtual machine monitors, guest virtual machines, data storage, and supporting middleware, there are also components that comprise the management backplane, such as those for self-service, resource metering, quota management, data replication and recovery, workload management, and cloud bursting.[3] Cloud services themselves may also be realized through nesting and layering with services from other cloud providers. Components change over time as upgrades and feature improvements occur, confounding matters further.

  Security depends not only on the correctness and effectiveness of many components, but also on the interactions among them. The number of possible interactions between components increases as the square of the number of components, which pushes the level of complexity upward. Complexity typically relates inversely to security, with greater complexity giving rise to vulnerabilities [Avo00, Gee08, Sch00].

- **Shared Multi-tenant Environment.** Public cloud services offered by providers have a serious underlying complication—subscribing organizations typically share components and resources with other subscribers that are unknown to them. Threats to network and

[3] Cloud bursting involves the deployment and launching of an application at a cloud and the redirection of requests to it, in the event that computing resources at organization's data center become saturated.

computing infrastructures continue to increase each year and have become more sophisticated. Having to share an infrastructure with unknown outside parties can be a major drawback for some applications and requires a high level of assurance for the strength of the security mechanisms used for logical separation. While not unique to cloud computing, logical separation is a non-trivial problem that is exacerbated by the scale of cloud computing. Access to organizational data and resources could inadvertently be exposed to other subscribers through a configuration or software error. An attacker could also pose as a subscriber to exploit vulnerabilities from within the cloud environment to gain unauthorized access.

- **Internet-facing Services.** Public cloud services are delivered over the Internet, exposing both the administrative interfaces used to self-service an account and the interfaces for users and applications to access other available services. Applications and data that were previously accessed from the confines an organization's intranet, but moved to the cloud, must now face increased risk from network threats that were previously defended against at the perimeter of the organization's intranet and from new threats that target the exposed interfaces. The effect is somewhat analogous to the inclusion of wireless access points into an organization's intranet at the onset of that technology. Requiring remote administrative access as the sole means to manage the assets of the organization held by the cloud provider also increases risk, compared with a traditional data center, where administrative access to platforms can be restricted to direct or internal connections.

- **Loss of Control.** While security and privacy concerns in cloud computing services are similar to those of traditional non-cloud services, they are amplified by external control over organizational assets and the potential for mismanagement of those assets. Migrating to a public cloud requires a transfer of control to the cloud provider over information as well as system components that were previously under the organization's direct control. Loss of control over both the physical and logical aspects of the system and data diminishes the organization's ability to maintain situational awareness, weigh alternatives, set priorities, and effect changes in security and privacy that are in the best interest of the organization.

A more detailed discussion of the security and privacy issues that stem from these fundamental concerns is given in the next chapter.

As with any technology, cloud computing services can be turned towards improper or illicit activities. A couple of noteworthy instances have already occurred that give a sense of what might be expected in the future:

- **Botnets.** In many ways, botnets assembled and controlled by hackers are an early form of cloud computing. Cost reduction, dynamic provisioning, redundancy, security, and many other characteristics of cloud computing apply. Botnets have been used for sending spam, harvesting login credentials, and launching injection attacks against Websites [Pro09]. Botnets could be used to launch a denial of service attack against the infrastructure of a cloud provider. The possibility that a cloud service could become infiltrated by a botnet has already occurred; in 2009, a command-and-control node was

discovered operating from within an IaaS cloud [Mcm09a, Whi09]. Spammers have also purchased cloud services directly and launched phishing campaigns, ensnaring recipients with malware via social engineering techniques [Kre08].

- **Mechanism Cracking.** WiFi Protected Access (WPA) Cracker, a cloud service ostensibly for penetration testers, is an example of harnessing cloud resources on demand to determine the encrypted password used to protect a wireless network. With cloud computing, a task that would take five days to run on a single computer takes only 20 minutes to accomplish on a cluster of 400 virtual machines [Rag09]. Because cryptography is used widely in authentication, data confidentiality and integrity, and other security mechanisms, these mechanisms become, in effect, less effective with the availability of cryptographic key cracking cloud services. Both cloud-based and traditional types of systems are possible targets. CAPTCHA cracking is another area where cloud services could be applied to bypass verification meant to thwart abusive use of Internet services by automated software.[4]

[4] CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) involves the solution of a simple test by a user before gaining service, as a means of thwarting unwanted automated access.

4. Key Security and Privacy Issues

Although the emergence of cloud computing is a recent development, insights into critical aspects of security can be gleaned from reported experiences of early adopters and also from researchers analyzing and experimenting with available cloud provider platforms and associated technologies. The sections below highlight privacy and security-related issues that are believed to have long-term significance for cloud computing. Where possible, to illustrate an issue, examples are given of problems previously exhibited or demonstrated. Note that security and privacy considerations that stem from information technology outsourcing are covered in the next chapter and complement the material below.

Because cloud computing has grown out of an amalgamation of technologies, including service oriented architecture, virtualization, Web 2.0, and utility computing, many of the privacy and security issues involved can be viewed as known problems cast in a new setting. The importance of their combined effect, however, should not be discounted. Cloud computing does represent a thought-provoking paradigm shift that goes beyond conventional norms to de-perimeterize the organizational infrastructure—*at the extreme, displacing applications from one organization's infrastructure to the infrastructure of another organization, where the applications of potential adversaries may also operate.*

4.1 Governance

Governance implies control and oversight over policies, procedures, and standards for application development, as well as the design, implementation, testing, and monitoring of deployed services. With the wide availability of cloud computing services, lack of organizational controls over employees engaging such services arbitrarily can be a source of problems. While cloud computing simplifies platform acquisition, it doesn't alleviate the need for governance; instead, it has the opposite effect, amplifying that need.

The ability to reduce capital investment and transform it into operational expenses is an advantage of cloud computing. Cloud computing can lower the initial cost of deploying new services and thus align expense with actual use.[5] However, the normal processes and procedures set in place by an organization for acquiring computational resources as capital expenditures may be easily bypassed by a department or an individual and the action obscured as operational expenses. If such actions are not governed by an organization, its policies and procedures for privacy, security, and oversight could be overlooked and the organization put at risk. For example, vulnerable systems could be deployed, legal regulations could be ignored, charges could amass quickly to unacceptable levels, resources could be used for unsanctioned purposes, or other untoward effects could occur.

[5] Many businesses also prefer operational expenses over capital expenditures, because of tax considerations (e.g., the ability to manage the cost of capital better and deduct operational expenses in the accounting period in which they are incurred versus depreciating the capital expenditure over time).

A study involving more than nine hundred information technology professionals in Europe and the United States indicates a strong concern by participants that cloud computing services may have been deployed without their knowledge in parts of their respective organization [Pon10]. The issue is somewhat akin to the problem with individuals setting up rogue wireless access points tied into the organizational infrastructure—without proper governance, the organizational computing infrastructure could be transformed into a sprawling, unmanageable mix of insecure services. Organizational practices pertaining to the policies, procedures, and standards used for application development and service provisioning, as well as the design, implementation, testing, and monitoring of deployed or engaged services, should be extended to cover cloud computing environments.

Dealing with cloud services requires attention to the roles and responsibilities involved, particularly with respect to managing risks. Ensuring systems are secure and risk is managed is challenging in any environment and even more daunting with cloud computing. Audit mechanisms and tools should be in place to determine how data is stored, protected, and used; to validate services; and to verify policy enforcement. A risk management program should also be in place that is flexible enough to deal with the continuously evolving and shifting risk landscape.

4.2 Compliance

Compliance involves conformance with an established specification, standard, regulation, or law. Various types of security and privacy laws and regulations exist within different countries at the national, state, and local levels, making compliance a potentially complicated issue for cloud computing.

- **Data Location.** One of the most common compliance issues facing an organization is data location [Bin09, Kan09, Ove10]. Use of an in-house computing center allows an organization to structure its computing environment and to know in detail where data is stored and what safeguards are used to protect the data. In contrast, a characteristic of many cloud computing services is that detailed information about the location of an organization's data is unavailable or not disclosed to the service subscriber. This situation makes it difficult to ascertain whether sufficient safeguards are in place and whether legal and regulatory compliance requirements are being met. External audits and security certifications can to some extent alleviate this issue, but they are not a panacea [Mag10].

  When information crosses borders, the governing legal, privacy, and regulatory regimes can be ambiguous and raise a variety of concerns (e.g., [CBC04]). Consequently, constraints on the trans-border flow of sensitive data, as well as the requirements on the protection afforded the data, have become the subject of national and regional privacy and security laws and regulations [Eis05]. Among the concerns to be addressed are whether the laws in the jurisdiction where the data was collected permit the flow, whether those laws continue to apply to the data post transfer, and whether the laws at the destination present additional risks or benefits [Eis05]. Technical, physical and administrative safeguards, such as access controls, often apply. For example, European

data protection laws may impose additional obligations on the handling and processing of data transferred to the U.S. [DoC00].

The main compliance concerns with trans-border data flows include whether the laws in the jurisdiction where the data was collected permit the flow, whether those laws continue to apply to the data post transfer, and whether the laws at the destination present additional risks or benefits [Eis05]. Technical, physical and administrative safeguards, such as access controls, often apply. For example, European data protection laws may impose additional obligations on the handling and processing of data transferred to the U.S. [DoC00].

- **Law and Regulations.** For U.S. Federal agencies, the major security and privacy compliance concerns include the Clinger-Cohen Act of 1996, the Office of Management and Budget (OMB) Circular No. A-130, particularly Appendix III, the Privacy Act of 1974, and the Federal Information Security Management Act (FISMA) of 2002. Also of importance are National Archives and Records Administration (NARA) statues, including the Federal Records Act (44 U.S.C. Chapters 21, 29, 31, 33) and NARA regulations (Title 36 of the Code of Federal Regulations, Chapter XII, Subchapter B).

  The Clinger-Cohen Act assigns responsibilities for the efficiency, security, and privacy of computer systems within the federal government and establishes a comprehensive approach for executive agencies to improve the acquisition and management of their information resources. As part of OMB's responsibilities under the Clinger-Cohen act, various circulars have been issued. Circular A-130 establishes policy for the management of Federal information resources, including procedural and analytic guidelines for implementing specific aspects of these policies. Appendix III requires that adequate security is provided for all agency information that is collected, processed, transmitted, stored, or disseminated in general support systems and major applications. The Privacy Act likewise governs the collection, maintenance, use, and dissemination of personally identifiable information about individuals that is maintained in systems of records by federal agencies.

  FISMA requires federal agencies to adequately protect their information and information systems against unauthorized access, use, disclosure, disruption, modification, or destruction [HR2458]. That mandate includes protecting information systems used or operated by an agency or by a contractor of an agency or other organization on behalf of an agency. That is, any external provider handling federal information or operating information systems on behalf of the federal government must meet the same security requirements as the source federal agency. The security requirements also apply to external subsystems storing, processing, or transmitting federal information and any services provided by or associated with the subsystem.

  Under the Federal Records Act and NARA regulations, agencies are responsible for managing federal records effectively throughout their lifecycle, including records in electronic information systems and in contracted environments. If a contractor holds federal records, the contractor must manage them in accordance with all applicable

records management laws and regulations. Managing the records includes secure storage, retrievability, and proper disposition, including transfer of permanently valuable records to NARA in an acceptable format [Fer10].

Other government and industry-association requirements, such as the Health Insurance Portability and Accountability Act (HIPAA) and the Payment Card Industry Data Security Standard (PCI DSS), may apply to a particular organization. For example, the Veterans Health Administration falls under HIPAA standards for private and public health care facilities and applies to both employees and contractors [DVA]. HIPAA requires both technical and physical safeguards for controlling access to data, which may create compliance issues for some cloud providers.

Cloud providers are becoming more sensitive to legal and regulatory concerns, and may be willing to commit to store and process data in specific jurisdictions and apply required safeguards for security and privacy. However, the degree to which they will accept liability for exposure of content under their control remains to be seen. Even so, organizations are ultimately accountable for the security and privacy of data held by a cloud provider on their behalf.

- **Electronic Discovery.** Electronic discovery involves the identification, collection, processing, analysis, and production of electronic documents in the discovery phase of litigation [Daw05]. Organizations also have other incentives and obligations to preserve and produce electronic documents, such as complying with audit and regulatory information requests, and for government organizations, with Freedom of Information Act (FOIA) requests. Documents not only include electronic mail, attachments, and other data objects stored on a computer system or storage media, but also any associated metadata, such as dates of object creation or modification, and non-rendered file content (i.e., data that is not explicitly displayed for users).

  The capabilities and process of a cloud provider, such as the form in which data is maintained and the electronic discovery-related tools available, affect the ability of the organization to meet its obligations in a cost effective, timely, and compliant manner [Mcd10]. For example, a cloud provider's archival capabilities may not preserve the original metadata as expected, causing spoliation (i.e., the intentional, reckless, or negligent destruction, loss, material alteration, or obstruction of evidence that is relevant to litigation), which could negatively impact litigation.

4.3 Trust

Under the cloud computing paradigm, an organization relinquishes direct control over many aspects of security and, in doing so, confers an unprecedented level of trust onto the cloud provider.

- **Insider Access.** Data processed or stored outside the confines of an organization, its firewall, and other security controls bring with it an inherent level of risk. The insider security threat is a well-known issue for most organizations and, despite the name, applies as well to outsourced cloud services [Ash10, Cap09, Kow08]. Insider threats go

beyond those posed by current or former employees to include contractors, organizational affiliates, and other parties that have received access to an organization's networks, systems, and data to carry out or facilitate operations. Incidents may involve various types of fraud, sabotage of information resources, and theft of confidential information. Incidents may also be caused unintentionally—for instance, a bank employee sending out sensitive customer information to the wrong Google mail account [Zet09b].

Moving data and applications to a cloud computing environment operated by a cloud provider expands the insider security risk not only to the cloud provider's staff, but also potentially among other customers using the service. For example, a denial of service attack launched by a malicious insider was demonstrated against a well-known IaaS cloud [Sla09]. The attack involved a cloud subscriber creating an initial 20 accounts and launching virtual machine instances for each, then using those accounts to create an additional 20 accounts and machine instances in an iterative fashion, exponentially growing and consuming resources beyond set limits.

- **Data Ownership.** The organization's ownership rights over the data must be firmly established in the service contract to enable a basis for trust. The continuing controversy over privacy and data ownership rights for social networking users illustrates the impact that ambiguous terms can have on the parties involved (e.g., [Goo10, Rap09]). Ideally, the contract should state clearly that the organization retains ownership over all its data; that the cloud provider acquires no rights or licenses through the agreement to use the data for its own purposes, including intellectual property rights or licenses; and that the cloud provider does not acquire and may not claim any security interest in the data [Mcd10]. For these provisions to work as intended, the terms of data ownership must not be subject to unilateral amendment by the cloud provider.

- **Composite Services.** Cloud services themselves can be composed through nesting and layering with other cloud services. For example, a SaaS provider could build its services upon the services of a PaaS or IaaS cloud. The level of availability of the SaaS cloud would then depend on the availability of those services. Cloud services that use third-party cloud providers to outsource or subcontract some of their services should raise concerns, including the scope of control over the third party, the responsibilities involved, and the remedies and recourse available should problems occur. Trust is often not transitive, requiring that third-party arrangements be disclosed in advance of reaching an agreement with the cloud provider, and that the terms of these arrangements are maintained throughout the agreement or until sufficient notification can be given of any anticipated changes.

  Liability and performance guarantees can become a serious issue with composite cloud services. For example, a consumer storage-based social networking service closed down after losing access to a significant amount of data from 20,000 of its subscribers. Because it relied on another cloud provider to host historical data, and on yet another cloud provider to host its newly launched application and database, direct responsibility for the cause of the failure was unclear and never resolved [Bro08].

- **Visibility.** Migration to public cloud services relinquishes control to the cloud provider for securing the systems on which the organization's data and applications operate. Management, procedural, and technical controls used in the cloud must be commensurate with those used for internal organizational systems or surpass them, to avoid creating gaps in security. Since metrics for comparing two computer systems are an ongoing area of research, making such comparisons can be a formidable task [Jan09]. Cloud providers are typically reluctant to provide details of their security and privacy, since such information might be used to devise an avenue of attack. Moreover, detailed network and system level monitoring by a cloud subscriber is generally not part of most service arrangements, limiting visibility and the means to audit operations directly (e.g., [Bro09, Dig08, Met09]).

  Transparency in the way the cloud provider operates is a vital ingredient for effective oversight over system security and privacy by an organization. To ensure that policy and procedures are being enforced throughout the system lifecycle, service arrangements should include some means for gaining visibility into the security controls and processes employed by the cloud provider and their performance over time. Ideally, the organization would have control over aspects of the means of visibility, such as the threshold for alerts and notifications or the level of detail and schedule for reports, to accommodate its needs.

- **Risk Management.** With cloud-based services, some subsystems or subsystem components are outside of the direct control of a subscribing organization. Many people feel more comfortable with risk when they have more control over the processes and equipment involved. At a minimum, a high degree of control provides the option to weigh alternatives, set priorities, and act decisively in the best interest of the organization when faced with an incident. Risk management is the process of identifying and assessing risk, and taking the necessary steps to reduce it to an acceptable level [Sto02]. Public cloud-based systems, as with traditional information systems, require that risks are managed throughout the system lifecycle.

  Assessing and managing risk in systems that use cloud services can be a challenge. To the extent practical, the organization should ensure that security controls are implemented correctly, operate as intended, and meet its security requirements. Establishing a level of trust about a cloud service is dependent on the degree of control an organization is able to exert on the provider to provision the security controls necessary to protect the organization's data and applications, and also the evidence provided about the effectiveness of those controls [Jtf10]. However, verifying the correct functioning of a subsystem and the effectiveness of security controls as extensively as with an organizational system may not be feasible in some cases, and other means (e.g., third-party audits) may be used to establish a level of trust. Ultimately, if the level of trust in the service falls below expectations and the organization is unable to employ compensating controls, it must either reject the service or accept a greater degree of risk.

4.4 Architecture

The architecture of the software systems used to deliver cloud services comprises hardware and software residing in the cloud. The physical location of the infrastructure is determined by the cloud provider as is the implementation of the reliability and scalability logic of the underlying support framework. Virtual machines often serve as the abstract unit of deployment and are loosely coupled with the cloud storage architecture. Applications are built on the programming interfaces of Internet-accessible services, which typically involve multiple cloud components communicating with each other over application programming interfaces. Many of the simplified interfaces and service abstractions belie the inherent complexity that affects security.

- **Attack Surface.** The hypervisor or virtual machine monitor is an additional layer of software between an operating system and hardware platform that is used to operate multi-tenant virtual machines. Besides virtualized resources, the hypervisor normally supports other application programming interfaces to conduct administrative operations, such as launching, migrating, and terminating virtual machine instances. Compared with a traditional non-virtualized implementation, the addition of a hypervisor causes an increase in the attack surface.

  The complexity in virtual machine environments can also be more challenging than their traditional counterparts, giving rise to conditions that undermine security [Gar05]. For example, paging, checkpointing, and migration of virtual machines can leak sensitive data to persistent storage, subverting protection mechanisms in the hosted operating system intended to prevent such occurrences. Moreover, the hypervisor itself can potentially be compromised. For instance, a vulnerability that allowed specially crafted File Transfer Protocol (FTP) requests to corrupt a heap buffer in the hypervisor, which could allow the execution of arbitrary code at the host, was discovered in a widely used virtualization software product, in a routine for Network Address Translation (NAT) [Sec05, She05].

- **Virtual Network Protection.** Most virtualization platforms have the ability to create software-based switches and network configurations as part of the virtual environment to allow virtual machines on the same host to communicate more directly and efficiently. For example, for virtual machines requiring no external network access, the virtual networking architectures of most virtualization software products support same-host networking, in which a private subnet is created for intra-host communications. Traffic over virtual networks may not be visible to security protection devices on the physical network, such as network-based intrusion detection and prevention systems [Vie09]. To avoid a loss of visibility and protection against intra-host attacks, duplication of the physical network protection capabilities may be required on the virtual network [Ref10, Vmw10].

- **Ancillary Data.** While the focus of protection is placed mainly on the application data, as guardians of the realm, cloud providers hold significant details about the service users' accounts that could be compromised and used in subsequent attacks. Payment information is one example; other, more subtle types of information, can also be involved. For example, a database of contact information stolen from a SaaS cloud

provider, via a targeted phishing attack against one of its employees, was used in turn to launch successful targeted electronic mail attacks against subscribers of the cloud service [Kre07, Mcm07]. The incident illustrates the need for cloud providers to promptly report security breaches occurring not only in the data the cloud provider holds for its subscribers, but also the data it holds *about* its subscribers.

Another type of ancillary data held by IaaS cloud providers is virtual machine images. A virtual machine image entails the software stack, including installed and configured applications, used to boot the virtual machine into an initial state or the state of some previous checkpoint. Sharing virtual machine images is a common practice in some cloud computing environments. Image repositories must be carefully managed and controlled to avoid problems.

The provider of an image faces risks, since an image can contain proprietary code and data and embody vulnerabilities. An attacker may attempt to examine images to determine whether they leak information or provide an avenue for attack [Wei09]. This is especially true of development images that are accidentally released. The reverse may also occur—an attacker may attempt to supply a virtual machine image containing malware to users of a cloud computing system [Jen09, Wei09].[6] For example, researchers demonstrated that by manipulating the registration process to gain a first-page listing, they could readily entice cloud users to run virtual machine images they contributed to the image repository of a popular cloud provider [Mee09]. The risks for users running tainted images include theft and corruption of data.

- **Client-Side Protection.** A successful defense against attacks requires securing both the client and server side of cloud computing. With emphasis typically placed on the latter, the former can be easily overlooked. Web browsers, a key element for many cloud computing services, and the various available plug-ins and extensions for them are notorious for their security problems [Jen09, Ker10, Pro07, Pro09]. Moreover, many browser add-ons do not provide automatic updates, increasing the persistence of any existing vulnerabilities.

  Maintaining physical and logical security over clients can be troublesome, especially with embedded mobile devices such as smart phones. Their size and portability can result in the loss of physical control. Built-in security mechanisms often go unused or can be overcome or circumvented without difficulty by a knowledgeable party to gain control over the device [Jan08]. Smart phones are also treated more as fixed appliances with a limited set of functions, than as general-purpose systems. No single operating system dominates and security patches and updates for system components and add-ons are not as frequent as for desktop clients, making vulnerabilities more persistent with a larger window of opportunity for exploitation.

[6] For PaaS and SaaS environments, a malicious implementation module is supplied.

The increased availability and use of social media, personal Webmail, and other publicly available sites also have associated risks that are a concern, since they can negatively impact the security of the browser, its underlying platform, and cloud services accessed, through social engineering attacks. For example, spyware was reportedly installed in a hospital system via an employee's personal Webmail account and sent the attacker more than 1,000 screen captures, containing financial and other confidential information, before being discovered [Mcm09b]. Having a backdoor Trojan, keystroke logger, or other type of malware running on a client does not bode well for the security of cloud or other Web-based services it accesses [Fre08, MRG10]. As part of the overall security architecture for cloud computing, organizations need to review existing measures and employ additional ones, if necessary, to secure the client side. Banks are beginning to take the lead in deploying hardened browser environments that encrypt network exchanges and protect against keystroke logging [Dun10a, Dun10b].

- **Server-Side Protection.** Virtual servers and applications, much like their non-virtual counterparts, need to be secured in IaaS clouds, both physically and logically. Following organizational policies and procedures, hardening of the operating system and applications should occur to produce virtual machine images for deployment. Care must also be taken to provision security for the virtualized environments in which the images run [You07]. For example, virtual firewalls can be used to isolate groups of virtual machines from other hosted groups, such as production systems from development systems or development systems from other cloud-resident systems. Carefully managing virtual machine images is also important to avoid accidentally deploying images under development or containing vulnerabilities.

  Hybrid clouds are a type of composite cloud with similar protection issues. In a hybrid cloud the infrastructure consists of a private cloud composed with either a public cloud or another organization's private cloud. The clouds themselves remain unique entities, bound together by standardized or proprietary technology that enables unified service delivery, but also creates interdependency. For example, identification and authentication might be performed through an organization's private cloud infrastructure, as a means for its users to gain access to services provisioned in a public cloud. Preventing holes or leaks between the composed infrastructures is a major concern with hybrid clouds, because of increases in complexity and diffusion of responsibilities. The availability of the hybrid cloud, computed as the product of the availability levels for the component clouds, can also be a concern; if the percent availability of any one component drops, the overall availability suffers proportionately.

4.5 Identity and Access Management

Data sensitivity and privacy of information have become increasingly an area of concern for organizations and unauthorized access to information resources in the cloud is a major concern. One recurring issue is that the organizational identification and authentication framework may not naturally extend into the cloud and extending or changing the existing framework to support cloud services may be difficult [Cho09]. The alternative of employing two different authentication systems, one for the internal organizational systems and another for external cloud-based systems, is a complication that can become unworkable over time. Identity

federation, popularized with the introduction of service oriented architectures, is one solution that can be accomplished in a number of ways, such as with the Security Assertion Markup Language (SAML) standard or the OpenID standard.

- **Authentication.** A growing number of cloud providers support the SAML standard and use it to administer users and authenticate them before providing access to applications and data. SAML provides a means to exchange information, such as assertions related to a subject or authentication information, between cooperating domains. SAML request and response messages are typically mapped over the Simple Object Access Protocol (SOAP), which relies on the eXtensible Markup Language (XML) for its format. SOAP messages are digitally signed. For example, once a user has established a public key certificate for a public cloud, the private key can be used to sign SOAP requests.

  SOAP message security validation is complicated and must be carried out carefully to prevent attacks. For example, XML wrapping attacks have been successfully demonstrated against a public IaaS cloud [Gaj09, Gru09]. XML wrapping involves manipulation of SOAP messages. A new element (i.e., the wrapper) is introduced into the SOAP Security header; the original message body is then moved under the wrapper and replaced by a bogus body containing an operation defined by the attacker [Gaj09, Gru09]. The original body can still be referenced and its signature verified, but the operation in the replacement body is executed instead.

- **Access Control.** SAML alone is not sufficient to provide cloud-based identity and access management services. The capability to adapt cloud subscriber privileges and maintain control over access to resources is also needed. As part of identity management, standards like the eXtensible Access Control Markup Language (XACML) can be used by a cloud provider to control access to cloud resources, instead of using a proprietary interface. XACML focuses on the mechanism for arriving at authorization decisions, which complements SAML's focus on the means for transferring authentication and authorization decisions between cooperating entities. XACML is capable of controlling the proprietary service interfaces of most providers, and some cloud providers already have it in place. Messages transmitted between XACML entities are susceptible to attack by malicious third parties, making it important to have safeguards in place to protect decision requests and authorization decisions from possible attacks, including unauthorized disclosure, replay, deletion and modification [Kel05].

4.6 Software Isolation

High degrees of multi-tenancy over large numbers of platforms are needed for cloud computing to achieve the envisioned flexibility of on-demand provisioning of reliable services and the cost benefits and efficiencies due to economies of scale. To reach the high scales of consumption desired, cloud providers have to ensure dynamic flexible delivery of service and isolation of subscriber resources. Multi-tenancy in cloud computing is typically done by multiplexing the execution of virtual machines from potentially different users on the same physical server [Ris09]. It is important to note that applications deployed on guest virtual machines remain susceptible to attack and compromise, much the same as their non-virtualized counterparts. This

was dramatically exemplified by a botnet found operating out of an IaaS cloud computing environment [Mcm09a, Whi09].

- **Hypervisor Complexity.** The security of a computer system depends on the quality of the underlying software kernel that controls the confinement and execution of processes. A virtual machine monitor or hypervisor is designed to run multiple virtual machines, each hosting an operating system and applications, concurrently on a single host computer, and to provide isolation between the different guest virtual machines.

  A virtual machine monitor can, in theory, be smaller and less complex than an operating system. These characteristics generally make it easier to analyze and improve the quality of security, giving a virtual machine monitor the potential to be better suited for maintaining strong isolation between guest virtual machines than an operating system is for isolating processes [Kar08]. In practice, however, modern hypervisors can be large and complex, comparable to an operating system, which negates this advantage. For example, Xen, an open source x86 virtual machine monitor, incorporates a modified Linux kernel to implement a privileged partition for input/output operations, and KVM, another open source effort, transforms a Linux kernel into a virtual machine monitor [Kar08, Sha08, Xen08]. Understanding the use of virtualization by a cloud provider is a prerequisite to understanding the security risk involved.

- **Attack Vectors.** Multi-tenancy in virtual machine-based cloud infrastructures, together with the subtleties in the way physical resources are shared between guest virtual machines, can give rise to new sources of threat. The most serious threat is that malicious code can escape the confines of its virtual machine and interfere with the hypervisor or other guest virtual machines. Live migration, the ability to transition a virtual machine between hypervisors on different host computers without halting the guest operating system, and other features provided by virtual machine monitor environments to facilitate systems management, also increase software size and complexity and potentially add other areas to target in an attack.

  Several examples illustrate the types of attack vectors possible. The first is mapping the cloud infrastructure. While seemingly a daunting task to perform, researchers have demonstrated an approach with a popular IaaS cloud [Ris09]. By launching multiple virtual machine instances from multiple cloud subscriber accounts and using network probes, assigned IP addresses and domain names were analyzed to identify service location patterns. Building on that information and general technique, the plausible location of a specific target virtual machine could be identified and new virtual machines instantiated to be eventually co-resident with the target.

  Once a suitable target location is found, the next step for the guest virtual machine is to bypass or overcome containment by the hypervisor or to takedown the hypervisor and system entirely. Weaknesses in the provided programming interfaces and the processing of instructions are common targets for uncovering vulnerabilities to exploit [Fer07]. For example, a serious flaw that allowed an attacker to write to an arbitrary out-of-bounds memory location was discovered in the power management code of a hypervisor by

fuzzing emulated I/O ports [Orm07].[7] A denial of service vulnerability, which could allow a guest virtual machine to crash the host computer along with the other virtual machines being hosted, was also uncovered in a virtual device driver of a popular virtualization software product [Vmw09].

More indirect attack avenues may also be possible. For example, researchers developed a way for an attacker to gain administrative control of guest virtual machines during a live migration, employing a man-in-the-middle attack to modify the code used for authentication [Obe08a]. Memory modification during migration presents other possibilities, such as the potential to insert a virtual machine-based rootkit layer below the operating system [Kin06]. A zero-day exploit in HyperVM, an open source application for managing virtual private servers, purportedly led to the destruction of approximately 100,000 virtual server-based Websites hosted by a service provider [Goo09b]. Another example of an indirect attack involves monitoring resource utilization on a shared server to gain information and perhaps perform a side-channel attack, similar to attacks used in other computing environments [Ris09]. For example, an attacker could determine periods of high activity, estimate high-traffic rates, and possibly launch keystroke timing attacks to gather passwords and other data from a target server.

4.7 Data Protection

Data stored in the cloud typically resides in a shared environment collocated with data from other customers. Organizations moving sensitive and regulated data into the cloud, therefore, must account for the means by which access to the data is controlled and the data is kept secure.

- **Data Isolation.** Data can take many forms. For example, for cloud-based application development, it includes the application programs, scripts, and configuration settings, along with the development tools. For deployed applications, it includes records and other content created or used by the applications, as well as account information about the users of the applications. Access controls are one means to keep data away from unauthorized users; encryption is another. Access controls are typically identity-based, which makes authentication of the user's identity an important issue in cloud computing.

  Database environments used in cloud computing can vary significantly. For example, some environments support a multi-instance model, while others support a multi-tenant model. The former provide a unique database management system running on a virtual machine instance for each cloud subscriber, giving the subscriber complete control over role definition, user authorization, and other administrative tasks related to security. The latter provide a predefined environment for the cloud subscriber that is shared with other tenants, typically through tagging data with a subscriber identifier. Tagging gives the appearance of exclusive use of the instance, but relies on the cloud provider to establish and maintain a sound secure database environment.

[7] Fuzzing is a type of fault injection technique that involves sending pseudorandom data to an interface to discover flaws.

Various types of multi-tenant arrangements exist for databases. Each arrangement pools resources differently, offering different degrees of isolation and resource efficiency [Jac07, Wai08]. Other considerations also apply. For example, certain features like data encryption are only viable with arrangements that use separate rather than shared databases. These sorts of tradeoffs require careful evaluation of the suitability of the data management solution for the data involved. Requirements in certain fields, such as healthcare, would likely influence the choice of database and data organization used in an application. Privacy sensitive information, in general, is a serious concern [Pea09].

Data must be secured while at rest, in transit, and in use, and access to the data must be controlled. Standards for communications protocols and public key certificates allow data transfers to be protected using cryptography. Procedures for protecting data at rest are not as well standardized, however, making interoperability an issue due to the predominance of proprietary systems. The lack of interoperability affects the availability of data and complicates the portability of applications and data between cloud providers.

Currently, the responsibility for cryptographic key management falls mainly on the cloud service subscriber. Key generation and storage is usually performed outside the cloud using hardware security modules, which do not scale well to the cloud paradigm. NIST's Cryptographic Key Management Project is identifying scalable and usable cryptographic key management and exchange strategies for use by government, which could help to alleviate the problem eventually.[8] Protecting data in use is an emerging area of cryptography with little practical results to offer, leaving trust mechanisms as the main safeguard [Gre09].

- **Data Sanitization.** The data sanitization practices that a cloud provider implements have obvious implications for security. Sanitization is the removal of sensitive data from a storage device in various situations, such as when a storage device is removed from service or moved elsewhere to be stored. Data sanitization also applies to backup copies made for recovery and restoration of service, and also residual data remaining upon termination of service. In a cloud computing environment, data from one subscriber is physically commingled with the data of other subscribers, which can complicate matters. For instance, many examples exist of researchers obtaining used drives from online auctions and other sources and recovering large amounts of sensitive information from them (e.g., [Val08]). With the proper skills and equipment, it is also possible to recover data from failed drives that are not disposed of properly by cloud providers [Sob06].

4.8 Availability

In simple terms, availability is the extent to which an organization's full set of computational resources is accessible and usable. Availability can be affected temporarily or permanently, and a loss can be partial or complete. Denial of service attacks, equipment outages, and natural

[8] Cryptographic Key Management Project Website - http://csrc.nist.gov/groups/ST/key_mgmt/

disasters are all threats to availability. The concern is that most downtime is unplanned and can impact the mission of the organization.

- **Temporary Outages.** Despite employing architectures designed for high service reliability and availability, cloud computing services can and do experience outages and performance slowdowns [Lea09]. A number of examples illustrate this point. In February 2008, a popular storage cloud service suffered a three-hour outage that affected its subscribers, including Twitter and other startup companies [Dig08, Kri08, Mil08]. In June 2009, a lightning storm caused a partial outage of an IaaS cloud that affected some users for four hours [Mil09]. Similarly, in February 2008, a database cluster failure at a SaaS cloud caused an outage for several hours, and in January 2009, another brief outage occurred due to a network device failure [Fer09, Goo09a, Mod08]. In March 2009, a PaaS cloud experienced severe degradation for about 22 hours due to networking issues related to an upgrade [Cla09, Mic09].

  At a level of 99.95% reliability, 4.38 hours of downtime are to be expected in a year. Periods of scheduled maintenance are also usually excluded as a source of downtime in SLAs and able to be scheduled by the cloud provider with short notice. The level of reliability of a cloud service and its capabilities for backup and recovery need to be addressed in the organization's contingency planning to ensure the recovery and restoration of disrupted cloud services and operations, using alternate services, equipment, and locations, if required. Cloud storage services may represent a single point of failure for the applications hosted there. In such situations, the services of a second cloud provider could be used to back up data processed by the primary provider to ensure that during a prolonged disruption or serious disaster at the primary, the data remains available for immediate resumption of critical operations.

- **Prolonged and Permanent Outages.** The possibility exists for a cloud provider to experience serious problems, like bankruptcy or facility loss, which affect service for extended periods or cause a complete shutdown. For example, in April 2009, the Federal Bureau of Investigation raided computing centers in Texas and seized hundreds of servers, when investigating fraud allegations against a handful of companies that operated out of the centers [Zet09a]. The seizure disrupted service to hundreds of other businesses unrelated to the investigation, but who had the misfortune of having their computer operations collocated at the targeted centers [Zet09a]. Other examples of outages are the major data loss experienced in 2009 by a bookmark repository service, and the abrupt failure of an on-line storage-as-a-service provider, who closed without warning to its users in 2008 [Cal09, Gun08]. Changing business conditions may also cause a cloud provider to disband its services, as occurred recently with an online cloud storage service [Sto10]. The organization's contingency plan should address prolonged and permanent system disruptions through support for continuity of operations that effect the restoration of essential functions elsewhere.

- **Denial of Service.** A denial of service attack involves saturating the target with bogus requests to prevent it from responding to legitimate requests in a timely manner. An attacker typically uses multiple computers or a botnet to launch an assault. Even an

unsuccessful distributed denial of service attack can quickly consume large amounts of resources to defend against and cause charges to soar. The dynamic provisioning of a cloud in some ways simplifies the work of an attacker to cause harm. While the resources of a cloud are significant, with enough attacking computers they can become saturated [Jen09]. For example, a denial of service attack against a code hosting site operating over an IaaS cloud resulted in more than 19 hours of downtime [Bro09, Met09].

Besides attacks against publicly accessible services, denial of service attacks can occur against internally accessible services, such as those used in cloud management [Sla09]. Internally assigned non-routable addresses, used to manage resources within a cloud provider's network, may also be used as an attack vector [Kre08]. A worst-case possibility that exists is for elements of one cloud to attack those of another or to attack some of its own elements [Jen09].

- **Value Concentration.** A response to the question "Why do you do rob banks?" is often attributed to Willie Hutton, a historic and prolific bank robber [Coc97]—his answer: "because that is where the money is." In many ways, data records are the currency of the 21st century and cloud-based data stores are the bank vault, making them an increasingly preferred target due to the collective value concentrated there [Row07]. Just as economies of scale exist in robbing banks instead of individuals, a high payoff ratio also exists for successfully compromising a cloud.

  As opposed to a direct approach, finesse and circumvention was Willie's trademark. That style works as well in the digital world of cloud computing. For instance, a recent exploit involved targeting an electronic mail account of a social networking service administrator, reportedly by answering a set of security questions to gain access to the account, and using the information found there to gain access to company files stored in a PaaS cloud [Inf09, Sut09]. Similar weaknesses have been identified in public clouds [Gar07]. A registered electronic mail address and valid password for an account are all that are required to download authentication credentials from a cloud provider's management dashboard, which in turn grant access to all of the account's resources. Since lost passwords can be reset by electronic mail, an attacker controlling the mail system of an account, or passively eavesdropping on the network through which electronic mail containing a password reset would pass, could effectively take control of the account.

  Having data collocated with that of an organization with a high threat profile could also lead to a denial of service, as an unintended casualty from an attack targeted against that organization [Row07]. Similarly, side effects from a physical attack against a high profile organization's cloud-based resources are also a possibility. For example, over the years, facilities of the Internal Revenue Service have attracted their share of attention from would-be attackers [Kat10, Lab95, Lat96, Sch10].

4.9 Incident Response

As the name implies, incident response involves an organized method for dealing with the consequences of an attack against the security of a computer system. The cloud provider's role is vital in performing incident response activities, including incident verification, attack analysis, containment, data collection and preservation, problem remediation, and service restoration. Revising an organization's incident response plan to address differences between the organizational computing environment and a cloud computing environment is an important, but easy-to-overlook prerequisite to transitioning applications and data.

Collaboration between the service subscriber and provider in recognizing and responding to an incident is essential to security and privacy in cloud computing. The complexity of the service can obscure recognition and analysis of incidents. For example, it reportedly took one IaaS provider approximately eight hours to recognize and begin taking action on an apparent denial of service attack against its cloud infrastructure, after the issue was reported by a subscriber of the service [Bro09, Met09]. Understanding and negotiating the provisions and procedures for incident response should be done before entering a service contract, rather than as an afterthought. The geographic location of data is a related issue that can impede an investigation, and is a relevant subject for contract discussions.

Response to an incident should be handled in a way that limits damage and reduces recovery time and costs. Being able to convene a mixed team of representatives from the cloud provider and service subscriber quickly is an important facet to meeting this goal. Remedies may involve only a single party or require the participation of both parties. Resolution of a problem may also affect other subscribers of the cloud service. It is important that cloud providers have a transparent response process and mechanisms to share information with their subscribers during and after the incident.

4.10 Summary of Recommendations

A number of significant security and privacy issues were covered in the previous subsections. Table 1 summarizes those issues and the precautions that apply as a set of recommendations for organizations to follow when planning, reviewing, negotiating, or initiating a public cloud service outsourcing arrangement.

Table 1: Security and Privacy Issues and Precautions

Areas	Precautions
Governance	Extend organizational practices pertaining to the policies, procedures, and standards used for application development and service provisioning in the cloud, as well as the design, implementation, testing, and monitoring of deployed or engaged services. Put in place audit mechanisms and tools to ensure organizational practices are followed throughout the system lifecycle.

Areas	Precautions
Compliance	Understand the various types of laws and regulations that impose security and privacy obligations on the organization and potentially impact cloud computing initiatives, particularly those involving data location, privacy and security controls, and electronic discovery requirements. Review and assess the cloud provider's offerings with respect to the organizational requirements to be met and ensure that the contract terms adequately meet the requirements.
Trust	Incorporate mechanisms into the contract that allow visibility into the security and privacy controls and processes employed by the cloud provider, and their performance over time. Institute a risk management program that is flexible enough to adapt to the continuously evolving and shifting risk landscape.
Architecture	Understand the underlying technologies the cloud provider uses to provision services, including the implications of the technical controls involved on the security and privacy of the system, with respect to the full lifecycle of the system and for all system components.
Identity and Access Management	Ensure that adequate safeguards are in place to secure authentication, authorization, and other identity and access management functions.
Software Isolation	Understand virtualization and other software isolation techniques that the cloud provider employs, and assess the risks involved.
Data Protection	Evaluate the suitability of the cloud provider's data management solutions for the organizational data concerned.
Availability	Ensure that during an intermediate or prolonged disruption or a serious disaster, critical operations can be immediately resumed and that all operations can be eventually reinstituted in a timely and organized manner.
Incident Response	Understand and negotiate the contract provisions and procedures for incident response required by the organization.

5. Public Cloud Outsourcing

Although cloud computing is a new computing paradigm, outsourcing information technology services is not. The steps that organizations take remain basically the same for public clouds as with other, more traditional, information technology services, and existing guidelines for outsourcing generally apply as well. What does change with public cloud computing, however, is the potential for increased complexity and difficulty in providing adequate oversight to maintain accountability and control over deployed applications and systems throughout their lifecycle. This can be especially daunting when non-negotiable SLAs are involved, since responsibilities normally held by the organization are given over to the cloud provider with little recourse for the organization to address problems and resolve issues, which may arise, to its satisfaction.

Reaching agreement on the terms of service of a negotiated SLA for public cloud services can be a complicated process fraught with technical and legal issues. As discussed in the previous chapter, migrating organizational data and functions into the cloud is accompanied by a host of security and privacy issues to be addressed, many of which concern the adequacy of the cloud provider's technical controls for an organization's needs. Service arrangements defined in the terms of service must also meet existing privacy policies for information protection, dissemination and disclosure. Each cloud provider and service arrangement has distinct costs and risks associated with it. A decision based on any one issue can have major implications for the organization in other areas [Gra03].

Considering the growing number of cloud providers and range of services offered, organizations must exercise due diligence when moving functions to the cloud. Decision making about new services and service arrangements entails striking a balance between benefits in cost and productivity versus drawbacks in risk and liability.

Cloud Migration Case Study: The City of Los Angeles' initiative to move to cloud computing provides some insight to the planning involved and the issues that can arise [CSC10]. The effort involves switching the city's electronic mail and calendaring system from an on-site solution to a public SaaS cloud that provides those services, and adding capabilities to improve productivity and collaboration [CSC10, DPW10, SECS09]. User training and data migration are also part of the contracted effort.

Reports of the SaaS E-mail and Collaboration Solution (SECS) indicate that the city was able to negotiate a number of security and privacy-related items into the contract that would be of interest to most government agencies [CSC10, Ove10, Wil10]. For example, the police and fire departments expressed concern that arrest records and other sensitive criminal data they handle could be vulnerable when maintained on external servers, which resulted in a private cloud arrangement for sensitive data. As an added measure, the California Department of Justice will certify the cloud provider's employees who have access to the city's data.

Other important negotiated features include mandatory data encryption, data storage location constraints, service level requirements with monetary penalties, electronic discovery functionality, well-defined data ownership and exit rights, mandatory subcontractor flow-down, and a broad indemnification obligation with unlimited liability for

certain breaches [Ove10, Wil10]. Data is stored in encrypted form and remains permanently the property of the city [Cra10]. The cloud provider requires written approval from the city to open any files in the clear; all accesses are logged and the city has a means to self-audit accesses [Cra10]. The cloud provider is reportedly building a segregated cloud to house data owned by the City of Los Angeles and other public-sector customers.

As with nearly all software changeovers, training, integration, data migration, and other related issues exist and their impact on productivity should not be underestimated or discounted when planning cloud computing initiatives [Mic10]. For example, there are some striking differences in features between L.A.'s outgoing electronic mail services and those of the SaaS [DPW10]: the cloud provider's mail service does not support classifying outgoing electronic mail with High, Standard, and Low priority; it also does not support a feature to track replies from recipients; nor does it support the use of folders to organize email, relying instead on labels. City employees are also required to carry out important migration-related tasks that include such things as cleaning up existing mail accounts by deleting all unimportant electronic mail and canceled appointments; archiving all mail by year; and individually saving any mail attachments larger than 25 megabytes, since they will not be automatically migrated to the new system [DPW10].

The security of sensitive data from the police department and other city agencies is proving to be a more difficult requirement to satisfy than thought originally, and causing a delay in implementation [Sar10]. Because of the delay, the outgoing system must continue to operate in parallel with the replacement system longer than originally planned and at additional cost, until a solution is in place. Hundreds of police department accounts that were switched over to the new system had to be restored to the outgoing system in the interim. The Director of Operations for the cloud provider noted: "LA's move to the cloud is the first of its kind, and it's not surprising that it's taken a little longer than anticipated to identify and address all of the City's unique requirements" [Din10]. It has been reported that if this issue is not resolved by June 2011, the end of the fiscal year, city officials would consider terminating the agreement, and possibly looking into whether a breach of contract occurred [Sar10].

5.1 General Concerns

The terms of traditional information technology outsourcing contracts, particularly those involving sensitive data, can serve as guidelines for cloud computing initiatives. Three main security and privacy issues in service contracts have been identified previously and are relevant to outsourcing public cloud computing services [All88, Len03]:

- **Inadequate Policies and Practices.** The security policies and practices of the cloud provider might not be adequate or compatible with those of the organization. The same issue applies to privacy as well. This can result in undetected intrusions or violations due to insufficient auditing and monitoring policies by the cloud provider; lack of sufficient data and configuration integrity due to a mismatch between the organization's and the cloud provider's policies for separation of duty (i.e., clear assignment of roles and responsibilities) or redundancy (i.e., having sufficient checks and balances to ensure an operation is done consistently and correctly); and loss of privacy due to the cloud provider handling sensitive information less rigorously than the organization's policy dictates [All88].

- **Weak Confidentiality and Integrity Sureties.** Insufficient security controls in the cloud provider's platform could affect negatively the confidentiality and privacy, or integrity of the system. For example, use of an insecure method of remote access could allow intruders to gain unauthorized access, modify, or destroy the organization's information systems and resources; to deliberately introduce security vulnerabilities or malware into the system; or to launch attacks on other systems from the organization's network, perhaps making it liable for damages [All88].

- **Weak Availability Sureties.** Insufficient safeguards in the cloud provider's platform could negatively affect the availability of the system. Besides the applications directly affected, a loss of system availability may cause a conflict for key resources that are required for critical organizational operations. For example, if disruptive processing operations are performed by the cloud provider (e.g., load rebalancing due to site failure or emergency maintenance) at the same time as peak organizational processing occurs, a denial of service condition could arise [All88]. A denial of service attack targeted at the cloud provider could also affect the organization's applications and systems operating in the cloud or at the organization's data center.

Other noteworthy concerns, which are less directly related to security and privacy, also exist with outsourcing to public clouds. One of the most prevalent and challenging concerns is called the principal-agent problem. Another is the attenuation of an organization's technical expertise.

- **Principal-Agent Problem.** The principal-agent problem occurs when the incentives of the agent (i.e., the cloud provider) are not aligned with the interests of the principal (i.e., the organization) [Row07]. Because it can be difficult to determine the level of effort a cloud provider is exerting towards security and privacy administration and remediation, the concern is that the organization might not recognize if the service level is dropping or has dropped below the extent required. One confounding issue is that increased security efforts are not guaranteed to result in noticeable improvements (e.g., fewer incidents), in part because of the growing amounts of malware and new types of attacks [Row07].

- **Attenuation of Expertise.** Outsourced computing services can over time diminish the level of technical knowledge and expertise of the organization, since management and staff no longer need to deal with technical issues at a detailed level [Gon09]. As new advancements and improvements are made to the cloud computing environment, the knowledge and expertise gained directly benefit the cloud provider, not the organization. Unless precautions are taken, an organization can lose its ability to keep up to date with technology advances and related security and privacy considerations, which in turn can affect its ability to plan and oversee new information technology projects effectively and to maintain accountability over existing cloud-based systems.

To remain accountable and mitigate the above-mentioned security and privacy issues, an organization can carry out a number of activities at each of three distinct stages of the outsourcing lifecycle: when initiating, conducting, and concluding outsourced services [All88, Len03]. Non-negotiable SLAs generally limit the range of activities available to an organization during the lifecycle, while negotiated SLAs, which provide greater range and flexibility,

necessitate careful scrutiny and prioritization of requirements that are incorporated into the terms of service in order to be cost effective. An organization may be able to employ compensating controls to work around identified shortcomings in a public cloud service with a non-negotiable SLA. Another alternative is for the organization to employ a cloud computing environment with a more suitable deployment model, such as a private cloud, which offers greater oversight and control over security and privacy.

5.2 Preliminary Activities

The organization must perform various planning activities in preparation for issuing a contract for outsourced public cloud services. Planning helps to ensure that an organization derives full benefit from information technology spending. It also helps to ensure that the computing environment is as secure as possible and in compliance with all relevant organizational policies and that data privacy is maintained. Planning activities include the following security-related items:

- **Specify Requirements.** The organization must identify its security, privacy, and other requirements for cloud services to meet, as a criterion for the selection of a cloud provider. Common security requirements include the following items [Len03]:

  - Personnel requirements, including clearances and responsibilities
  - Access control
  - Service availability
  - Problem reporting, review, and resolution
  - Information handling and disclosure agreements and procedures
  - Physical and logical access controls
  - Network connectivity and filtering
  - Configuration and patch management
  - Backup and recovery
  - Incident reporting, handling, and response
  - Continuity of operations
  - Account and resource management
  - Certification and accreditation
  - Assurance levels
  - Independent auditing of services.

  Other requirements related to security, such as privacy, data ownership rights, records management controls, and user training should also be identified. Reviewing common outsourcing provisions in existing cloud computing contracts that cover areas such as privacy and security standards, regulatory and compliance issues, service level criteria and penalties, change management processes, continuity of service provisions, and termination rights, can be helpful in formulating requirements [Ove10].

- **Assess Security and Privacy Risks.** While outsourcing relieves operational commitment on the part of the organization, the act of engaging a cloud provider's offerings for public cloud services poses risks against which an organization needs to safeguard itself. The analysis must include factors such as the service model involved,

the purpose and scope of the service, the types and level of access needed by the provider and proposed for use between the organizational computing environment and provider services, the service duration and dependencies, and the strength of protection offered via the security controls available from the cloud provider [Len03]. Privacy controls should be assessed as well as operational risks due to the cloud provider's locations. Another consideration, if a non-negotiable SLA applies, is whether the terms of service are subject to unilateral amendment by the cloud provider.

Data sensitivity is an important and crucial factor when analyzing risk. The range of data an organization deals with is sometimes not fully appreciated at first. While data repositories containing Personally Identifiable Information (PII) or classified information are easily recognized and taken into account, pockets of other types of sensitive data with different rules for handling may also exist. They include items such as the following:

- Law enforcement and investigative unit data
- Licensed source code and libraries, used in application development
- Digital documents and materials obtained under a non-disclosure agreement or memorandum of agreement
- Laboratory and research data whose collection, storage, and sharing are regulated
- Culturally sensitive data related to resource protection and management of Indian tribal land.

If the results of a risk analysis show that the level is too high, the organization may be able to apply compensating controls to reduce risk to an acceptable level. Otherwise, it must either reject the public cloud service or accept a greater degree of risk. As an alternative to rejecting the service and not going forward, it might be possible to reduce the scope of the outsourcing effort to deal exclusively with less sensitive data.

- **Assess the Competency of the Cloud Provider.** Before the contract for outsourced services is awarded, the organization should evaluate the cloud provider's ability and commitment to deliver the services over the target timeframe and meet the security and privacy levels stipulated. The cloud provider can be asked to demonstrate its capabilities and approach to security and privacy enforcement or to undergo an independent evaluation of its installation and systems [All98]. Interviewing current subscribers of the cloud provider and assessing their level of satisfaction can also provide insight into the competency of the cloud provider. Evaluating the privacy and security levels of the services should be done thoroughly, including consideration of such items as the following ones [Len03]:

  - Experience and technical expertise of personnel
  - The vetting process personnel undergo
  - Quality and frequency of security and awareness training provided to personnel
  - The type and effectiveness of the security services provided and underlying mechanisms used

- The adoption rate of new technologies
- The cloud provider's track record and ability to meet the organization's security and privacy policy, procedures, and regulatory compliance needs.

5.3 Initiating and Coincident Activities

The organization has a number of activities to carry out when awarding a contract to a cloud provider and also throughout the duration of the contract.

- **Establish Contractual Obligations.** The organization should ensure that all contractual requirements are explicitly stated in the SLA, including privacy and security provisions [Gra03, Len03]. The agreement should include definitions of both the organization's and the cloud provider's roles and responsibilities. It should also include the following items [Gra03]:

  - A detailed description of the service environment, including facility locations and applicable security requirements
  - Policies, procedures, and standards, including vetting and management of staff
  - Predefined service levels and associated costs
  - The process for assessing the cloud provider's compliance with the service level agreement, including independent audits and testing
  - Specific remedies for noncompliance or harm caused by the cloud provider
  - The period of performance and due dates for any deliverable
  - The cloud provider's points of interface with the organization
  - The organization's responsibilities for providing relevant information and resources to the cloud provider
  - Procedures, protections, and restrictions for commingling organizational data and for handling sensitive data
  - The cloud provider's obligations upon contract termination, such as the return and purging of data.

  Before entering into the contract, it is advisable to have an experienced legal advisor review its terms. Non-negotiable SLAs are typically drafted in favor of the cloud provider and may prove impracticable for an organization. If a negotiated SLA is used, a legal advisor should be involved in the negotiation process, since complicated legal issues are likely to arise during negotiations.

- **Assess Cloud Provider Performance.** Continual assessment of performance is needed to ensure all contract obligations are being met. It allows the organization to take immediate corrective or punitive action for noted deficiencies and also provides a reference point to improve the SLA terms of service [All88, Gra03, Len03].

5.4 Concluding Activities

At the end of a project lifecycle, transition to another cloud provider, or for other reasons, the organization can decide to conclude the outsourced public cloud services and close out the

contract. Organizations should perform the following activities upon the termination of an outsourcing contract.

- **Reaffirm Contractual Obligations.** The organization should alert the cloud provider about any contractual requirements, such as nondisclosure of certain terms of the contract and purging of organizational data, which must be observed upon termination [Len03].

- **Eliminate Physical and Electronic Access Rights.** All accounts and access rights assigned to the cloud provider should be revoked in a timely manner by the organization [All88, Len03]. Physical access rights of security tokens and badges also need to be revoked, and furthermore, any personal tokens and badges used for access need to be recovered [All88].

- **Recover Organizational Resources and Data.** The organization should ensure that any resources, such as software, equipment, documentation, and data made available to the cloud provider under the SLA, are returned in a usable form, as stipulated. If the terms of service require the cloud provider to purge data programs, backup copies, and other cloud subscriber content from its environment, evidence such as system reports or logs should be obtained and verified to ensure that the information has been properly expunged [Len03].

5.5 Summary of Recommendations

Table 2 below summarizes the issues and the precautions that apply at the various stages of outsourcing. They serve as a set of recommendations, complementary to those given earlier in Table 1, for organizations to follow when pursuing a public cloud service outsourcing arrangement.

Table 2: Outsourcing Activities and Precautions

Areas	Precautions
Preliminary Activities	Identify security, privacy, and other organizational requirements for cloud services to meet, as a criterion for selecting a cloud provider. Perform a risk assessment, analyzing the security and privacy controls of a cloud provider's environment with respect to the control objectives of the organization. Evaluate the cloud provider's ability and commitment to deliver cloud services over the target timeframe and meet the security and privacy levels stipulated.
Initiating and Coincident Activities	Ensure that all contractual requirements are explicitly recorded in the SLA, including privacy and security provisions, and that they are endorsed by the cloud provider. Involve a legal advisor in the negotiation and review of the terms of service of the SLA. Continually assess the performance of the cloud provider and ensure all contract obligations are being met.

Areas	Precautions
Concluding Activities	Alert the cloud provider about any contractual requirements that must be observed upon termination. Revoke all physical and electronic access rights assigned to the cloud provider and recover physical tokens and badges in a timely manner. Ensure that resources made available to the cloud provider under the SLA are returned in a usable form, and confirm evidence that information has been properly expunged.

6. Conclusion

Cloud computing promises to have far-reaching effects on the systems and networks of federal agencies and other organizations. Emphasis on the cost and performance benefits of public cloud computing, however, tend to overshadow some of the fundamental security and privacy concerns federal agencies and organizations have with these computing environments. Many of the features that make cloud computing attractive can also be at odds with traditional security models and controls. Several critical pieces of technology, such as a solution for federated trust, are not yet fully realized, impinging on successful cloud computing deployments. Determining the security of complex computer systems composed together is also a long-standing security issue that plagues large-scale computing in general, and cloud computing in particular. Attaining high-assurance qualities in implementations has been an elusive goal of computer security researchers and practitioners and, as demonstrated in the examples given in this report, is also a work in progress for cloud computing. Nevertheless, public cloud computing is a compelling computing paradigm that agencies need to incorporate as part their information technology solution set.

Accountability for security and privacy in public clouds remains with the organization. Federal agencies must ensure that any selected public cloud computing solution is configured, deployed, and managed to meet the security, privacy, and other requirements of the organization. Organizational data must be protected in a manner consistent with policies, whether in the organization's computing center or the cloud. The organization must ensure that security and privacy controls are implemented correctly and operate as intended.

The transition to an outsourced, public cloud computing environment is in many ways an exercise in risk management. Risk management entails identifying and assessing risk, and taking the steps to reduce it to an acceptable level. Assessing and managing risk in cloud computing systems can be a challenge. Throughout the system lifecycle, risks that are identified must be carefully balanced against the security and privacy controls available and the expected benefits from their utilization. Too many controls can be inefficient and ineffective, if the benefits outweigh the costs and associated risks. Federal agencies and organizations should work to ensure an appropriate balance between the number and strength of controls and the risks associated with cloud computing solutions.

7. References

[All88] Julia Allen et al., Security for Information Technology Service Contracts, CMU/SEI-SIM-003, Software Engineering Institute, Carnegie Mellon University, January 1988, <URL: http://www.sei.cmu.edu/reports/98sim003.pdf>.

[Arm10] Michael Armbrust et al., A View of Cloud Computing, Communications of the ACM, Association for Computing Machinery, Vol. 53, No. 4, April 2010.

[Ash10] Warwick Ashford, Google Confirms Dismissal of Engineer for Breaching Privacy Rules, Computer Weekly, September 16, 2010, <URL: http://www.computerweekly.com/Articles/2010/09/16/242877/Google-confirms-dismissal-of-engineer-for-breaching-privacy.htm>.

[Avo00] Frederick M. Avolio, Best Practices in Network Security, Network Computing, March 20, 2000, <URL: http://www.networkcomputing.com/1105/1105f2.html>.

[Bin09] David Binning, Top Five Cloud Computing Security Issues, Computer Weekly, April 24, 2009, <URL: http://www.computerweekly.com/Articles/2010/01/12/235782/Top-five-cloud-computing-security-issues.htm>.

[Bra10] Simon Bradshaw, Christopher Millard, Ian Walden, Contracts for Clouds: Comparison and Analysis of the Terms and Conditions of Cloud Computing Services, Queen Mary School of Law Legal Studies, Research Paper No. 63/2010, September 2, 2010, <URL: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1662374>.

[Bro08] Jon Brodkin, Loss of Customer Data Spurs Closure of Online Storage Service 'The Linkup,' Network World, August 11, 2008, <URL: http://www.networkworld.com/news/2008/081108-linkup-failure.html?page=1>.

[Bro09] Carl Brooks, Amazon EC2 Attack Prompts Customer Support Changes, Tech Target, October 12, 2009, <URL: http://searchcloudcomputing.techtarget.com/news/article/0,289142,sid201_gci1371090,00.html>.

[Cal09] Michael Calore, Ma.gnolia Suffers Major Data Loss, Site Taken Offline, Wired Magazine, January 30, 2009, <URL: http://www.wired.com/epicenter/2009/01/magnolia-suffer/>.

[Cap09] Dawn Cappelli, Andrew Moore, Randall Trzeciak, Timothy J. Shimeall, Common Sense Guide to Prevention and Detection of Insider Threats, 3rd Edition, Version 3.1, CERT, January 2009, <URL: http://www.cert.org/archive/pdf/CSG-V3.pdf>.

[CBC04] USA Patriot Act Comes under Fire in B.C. Report, CBC News, October 30, 2004, <URL: http://www.cbc.ca/canada/story/2004/10/29/patriotact_bc041029.html>.

[Cho09] Richard Chow et al., Controlling Data in the Cloud: Outsourcing Computation without Outsourcing Control, ACM Workshop on Cloud Computing Security, Chicago, Illinois, November 2009, <URL: http://www2.parc.com/csl/members/eshi/docs/ccsw.pdf>.

[Cla09] Gavin Clarke, Microsoft's Azure Cloud Suffers First Crash, The Register, March 16, 2009, <URL: http://www.theregister.co.uk/2009/03/16/azure_cloud_crash/>.

[Coc97] Steve Cocheo, The Bank Robber, the Quote, and the Final Irony, nFront, American Bankers Association (ABA) Banking Journal, 1997, <URL: http://www.banking.com/aba/profile_0397.htm>.

[Cra10] Private phone conversation with Kevin K. Crawford, Assistant General Manager, Information Technology Agency, City of Los Angeles, December 15, 2010.

[CSC10] LA SECS Overview: SaaS E-mail and Collaboration Solution (SECS) – Implementing Google for the Los Angeles, CSC, April 15, 2010, <URL: http://assets1.csc.com/lef/downloads/LEFBriefing_CSC_LA_Google_041510.pdf>.

[Daw05] Alistair B. Dawson, Understanding Electronic Discovery and Solving Its Problems, 56th Annual Program on Oil and Gas Law, The Center for American and International Law, February 17-18, 2005, Houston, Texas, <URL: http://www.brsfirm.com/publications/docs/00037W.pdf>.

[Dig08] Larry Dignam, Amazon Explains Its S3 Outage, ZDNET, February 16, 2008, <URL: http://www.zdnet.com/blog/btl/amazon-explains-its-s3-outage/8010>.

[Din10] Jocelyn Ding, LA's Move to Google Apps Continues Apace, Official Google Enterprise Blog, August 04, 2010, <URL: http://googleenterprise.blogspot.com/2010/08/las-move-to-google-apps-continues-apace.html>.

[DoC00] Safe Harbor Privacy Principles, U.S. Department of Commerce, July 21, 2000, <URL: http://www.export.gov/safeharbor/eg_main_018247.asp>.

[DPW10] LA DPW Engineering Newsletter, No. 10-22, Los Angeles City, Department of Public Works (DPW), April 21, 2010, <URL: http://eng.lacity.org/newsletters/2010/04-21-10.pdf>.

[Dun10a] John E. Dunn, Ultra-secure Firefox Offered to UK Bank Users, Techworld, February 26, 2010, <URL: http://news.techworld.com/security/3213740/ultra-secure-firefox-offered-to-uk-bank-users/>.

[Dun10b] John E. Dunn, Virtualised USB Key Beats Keyloggers, Techworld, February 22, 2010, <URL: http://news.techworld.com/security/3213277/virtualised-usb-key-beats-keyloggers/>.

[DVA] What the VA Is Doing to Protect Your Privacy, VA Pamphlet 005-06-1, Department of Veteran Affairs, <URL: http://www.privacy.va.gov/docs/VA005-06-1_privacy_brochure.pdf>.

[Eis05] Margaret P. Eisenhauer, Privacy and Security Law Issues in Off-shore Outsourcing Transactions, Hunton & Williams LLP, The Outsourcing Institute, Legal Corner, February 15, 2005, <URL: http://www.outsourcing.com/legal_corner/pdf/Outsourcing_Privacy.pdf>.

[Fer07] Peter Ferrie, Attacks on Virtual Machine Emulators, White Paper, Symantec Corporation, January 2007, <URL: http://www.symantec.com/avcenter/reference/Virtual_Machine_Threats.pdf>.

[Fer09] Tim Ferguson, Salesforce.com Outage Hits Thousands of Businesses, CNET News, January 8, 2009, <URL: http://news.cnet.com/8301-1001_3-10136540-92.html>.

[Fer10] David S. Ferreiro, Guidance on Managing Records in Cloud Computing Environments, NARA Bulletin 2010-05, September 8, 2010, <URL: http://www.archives.gov:80/records-mgmt/bulletins/2010/2010-05.html>.

[Fre08] Stefan Frei, Thomas Duebendorfer, Gunter Ollmann, Martin May, Understanding the Web Browser Threat: Examination of vulnerable online Web browser populations and the "insecurity iceberg", ETH Zurich, Tech Report Nr. 288, 2008, <URL: http://e-collection.ethbib.ethz.ch/eserv/eth:30892/eth-30892-01.pdf>.

[Fow09] Geoffrey Fowler, Ben Worthen, The Internet Industry Is on a Cloud – Whatever That May Mean, The Wall Street Journal, March 26, 2009, <URL: http://online.wsj.com/article/SB123802623665542725.html>.

[Gaj09] Sebastian Gajek, Meiko Jensen, Lijun Liao, and Jörg Schwenk, Analysis of Signature Wrapping Attacks and Countermeasures, IEEE International Conference on Web Services, Los Angeles, California, July 2009.

[Gar05] Tal Garfinkel, Mendel Rosenblum, When Virtual Is Harder than Real: Security Challenges in Virtual Machine Based Computing Environments, HotOS'05, Santa Fe, New Mexico, June 2005, <URL: http://www.stanford.edu/~talg/papers/HOTOS05/virtual-harder-hotos05.pdf>.

[Gar07] Simson Garfinkel, An Evaluation of Amazon's Grid Computing Services: EC2, S3 and SQS, Technical Report TR-08-07, Center for Research on Computation and Society, School for Engineering and Applied Sciences, Harvard University, July 2007, <URL: http://simson.net/clips/academic/2007.Harvard.S3.pdf>.

[Gee08] Daniel E. Geer, Complexity Is the Enemy, IEEE Security and Privacy, Vol. 6, No. 6, November/December 2008.

[Gon09] Reyes Gonzalez, Jose Gasco, and Juan Llopis, Information Systems Outsourcing Reasons and Risks: An Empirical Study, International Journal of Human and Social Sciences, Vol. 4, No. 3, 2009, <URL: http://www.waset.org/journals/ijhss/v4/v4-3-24.pdf>.

[Goo09a] Dan Goodin, Salesforce.com Outage Exposes Cloud's Dark Linings, The Register, January 6, 2009, <URL: http://www.theregister.co.uk/2009/01/06/salesforce_outage/>.

[Goo09b] Dan Goodin, Webhost Hack Wipes Out Data for 100,000 Sites, The Register, June 8, 2009, <URL: http://www.theregister.co.uk/2009/06/08/webhost_attack/>.

[Goo10] Dan Goodin, Privacy Watchdog Pack Demands Facebook Close the 'App Gap', The Register, June 16, 2010, <URL: http://www.theregister.co.uk/2010/06/16/facebook_privacy/>.

[Gra03] Tim Grance et al., Guide to Information Technology Security Services, NIST Special Publication 800-35, October 2003, <URL: http://csrc.nist.gov/publications/nistpubs/800-35/NIST-SP800-35.pdf>.

[Gre09] Andy Greenberg, IBM's Blindfolded Calculator, Forbes Magazine, July 13, 2009, <URL: http://www.forbes.com/forbes/2009/0713/breakthroughs-privacy-super-secret-encryption.html>.

[Gru09] Nils Gruschka, Luigi Lo Iacono, Vulnerable Cloud: SOAP Message Security Validation Revisited, IEEE International Conference on Web Services, Los Angeles, California, July 2009.

[Gun08] Mike Gunderloy, Who Protects Your Cloud Data?, Web Worker Daily, January 13, 2008, <URL: http://webworkerdaily.com/2008/01/13/who-protects-your-cloud-data/>.

[HR2458] Federal Information Security Management Act of 2002 (FISMA), H.R. 2458, Title III—Information Security, <URL: http://csrc.nist.gov/drivers/documents/FISMA-final.pdf>.

[Inf09] Twitter Email Account Hack Highlights Cloud Dangers, Infosecurity Magazine, July 23, 2009, <URL: http://www.infosecurity-magazine.com/view/2668/twitter-email-account-hack-highlights-cloud-dangers-/>.

[Jac07] Dean Jacobs, Stefan Aulbach, Ruminations on Multi-Tenant Databases, Fachtagung für Datenbanksysteme in Business, Technologie und Web, Aachen, Germany, March 5-9, 2007, <URL: http://www.btw2007.de/paper/p514.pdf>.

[Jan08] Wayne Jansen, Karen Scarfone, Guidelines on Cell Phone and PDA Security, Special Publication (SP) 800-124, National Institute of Standards and Technology, October 2008, <URL: http://csrc.nist.gov/publications/nistpubs/800-124/SP800-124.pdf>

[Jan09] Wayne Jansen, Directions in Security Metrics Research, Interagency Report (IR) 7564, National Institute of Standards and Technology, April 2009, <URL: http://csrc.nist.gov/publications/nistir/ir7564/nistir-7564_metrics-research.pdf>.

[Jen09] Meiko Jensen, Jörg Schwenk, Nils Gruschka, Luigi Lo Iacono, On Technical Security Issues in Cloud Computing, IEEE International Conference on Cloud Computing, Bangalore, India, September 21-25, 2009.

[Jtf10] Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach, Joint Task Force Transformation Initiative, NIST Special Publication 800-37, Revision 1, <URL: http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf>.

[Kan09] Balachandra Reddy Kandukuri, Ramakrishna Paturi V, Atanu Rakshit, Cloud Security Issues, IEEE International Conference on Services Computing, Bangalore, India, September 21-25, 2009.

[Kar08] Paul A. Karger, I/O for Virtual Machine Monitors: Security and Performance Issues, IEEE Security and Privacy, September/October 2008.

[Kat10] Neil Katz, Austin Plane Crash: Pilot Joseph Andrew Stack May Have Targeted IRS Offices, Says FBI, CBS News, February 18, 2010, <URL: http://www.cbsnews.com/8301-504083_162-6220271-504083.html?tag=contentMain%3bcontentBody>.

[Kel05] Yared Keleta, J.H.P. Eloff, H.S. Venter, Proposing a Secure XACML Architecture Ensuring Privacy and Trust, Research in Progress Paper, University of Pretoria, 2005, <URL: http://icsa.cs.up.ac.za/issa/2005/Proceedings/Research/093_Article.pdf>.

[Ker10] Sean Michael Kerner, Mozilla Confirms Security Threat from Malicious Firefox Add-ons, eSecurity Planet, February 5, 2010, <URL: http://www.esecurityplanet.com/news/article.php/3863331/Mozilla-Confirms-Security-Threat-From-Malicious-Firefox-Add-Ons.htm>.

[Kin06] Samuel King, Peter Chen, Yi-Min Wang, Chad Verbowski, Helen Wang, Jacob Lorch, SubVirt: Implementing Malware with Virtual Machines, IEEE Symposium on Security and Privacy, Berkeley, California, May 2006, <URL: http://www.eecs.umich.edu/~pmchen/papers/king06.pdf>.

[Kre07] Brian Krebs, Salesforce.com Acknowledges Data Loss, Security Fix, The Washington Post, November 6, 2007, <URL: http://blog.washingtonpost.com/securityfix/2007/11/salesforcecom_acknowledges_dat.html>.

[Kre08] Brian Krebs, Amazon: Hey Spammers, Get Off My Cloud! The Washington Post, July 1, 2008, <URL: http://voices.washingtonpost.com/securityfix/2008/07/amazon_hey_spammers_get_off_my.html>.

[Kow08] Eileen Kowalski et al., Insider Threat Study: Illicit Cyber Activity in the Government Sector, U.S. Secret Service and Carnegie Mellon University, Software Engineering Institute, January 2008, <URL: http://www.cert.org/archive/pdf/insiderthreat_gov2008.pdf>.

[Kri08] Michael Krigsma, Amazon S3 Web Services Down. Bad, Bad News for Customers, ZDNET, February 15, 2008, <URL: http://blogs.zdnet.com/projectfailures/?p=602>.

[Kum08] Sushil Kumar, Oracle Database Backup in the Cloud, White Paper, Oracle Corporation, September 2008.

[Lab95] Stephen Labaton, 2 Men Held in Attempt to Bomb I.R.S. Office, New York Times, December 29, 1995, <URL: http://www.nytimes.com/1995/12/29/us/2-men-held-in-attempt-to-bomb-irs-office.html?pagewanted=1>.

[Lat96] 20-Year Term in Plot to Bomb IRS Offices, Nation In Brief, Los Angeles Times, August 10, 1996, <URL: http://articles.latimes.com/1996-08-10/news/mn-32970_1_20-year-term>.

[Lea09] Neal Leavitt. Is Cloud Computing Really Ready for Prime Time?, IEEE Computer, January 2009.

[Len03] Bee Leng, A Security Guide for Acquiring Outsourced Service, GIAC GSEC Practical (v1.4b), SANS Institute, August 19, 2003, <URL: http://www.sans.org/reading_room/whitepapers/services/a_security_guide_for_acquiring_outsourced_service_1241>.

[Mag10] James Maguire, How Cloud Computing Security Resembles the Financial Meltdown, Datamation, internet.com, April 27, 2010, <URL: http://itmanagement.earthweb.com/netsys/article.php/3878811/How-Cloud-Computing-Security-Resembles-the-Financial-Meltdown.htm>.

[Mcd10] Steve McDonald, Legal and Quasi-Legal Issues in Cloud Computing Contracts, Workshop Document, EDUCAUSE and NACUBO Workshop on Cloud Computing

and Shared Services, Tempe, Arizona, February 8–10, 2010, <URL: http://net.educause.edu/section_params/conf/CCW10/issues.pdf>.

[Mcm07] Robert McMillan, Salesforce.com Warns Customers of Phishing Scam, PC Magazine, IDG News Network, November 6, 2007, <URL: http://www.pcworld.com/businesscenter/article/139353/salesforcecom_warns_customers_of_phishing_scam.html>.

[Mcm09a] Robert McMillan, Hackers Find a Home in Amazon's EC2 Cloud, Infoworld, IDG News Network, December 10, 2009, <URL: http://www.infoworld.com/d/cloud-computing/hackers-find-home-in-amazons-ec2-cloud-742>.

[Mcm09b] Robert McMillan, Misdirected Spyware Infects Ohio Hospital, PC Magazine, IDG News Service September 17, 2009, <URL: http://www.pcworld.com/businesscenter/article/172185/misdirected_spyware_infects_ohio_hospital.html>.

[Mee09] Haroon Meer, Nick Arvanitis, Marco Slaviero, Clobbering the Cloud, Part 4 of 5, Black Hat USA Talk Write-up, SensePost SDH Labs, 2009, <URL: http://www.sensepost.com/labs/conferences/clobbering_the_cloud/amazon>.

[Mel09] Peter Mell, Tim Grance, The NIST Definition of Cloud Computing, Version 15, October 7, 2009, <URL: http://csrc.nist.gov/groups/SNS/cloud-computing>.

[Met09] Cade Metz, DDoS Attack Rains Down on Amazon Cloud, The Register, October 5, 2009, <URL: http://www.theregister.co.uk/2009/10/05/amazon_bitbucket_outage/>.

[Mic09] The Windows Azure Malfunction This Weekend, Windows Azure <Team Blog>, Microsoft Corporation, March 18, 2009, <URL: http://blogs.msdn.com/windowsazure/archive/2009/03/18/the-windows-azure-malfunction-this-weekend.aspx>.

[Mic10] Fact-Based Comparison of Hosted Services: Google vs. Microsoft, Microsoft Corporation, May 16, 2010, <URL: http://download.microsoft.com/download/0/5/F/05FF69ED-6F8F-4357-863B-12E27D6F1115/Hosted%20Services%20Comparison%20Whitepaper%20-%20Google%20vs%20Microsoft.pdf>.

[Mil08] Rich Miller, Major Outage for Amazon S3 and EC2, Data Center Knowledge, February 15, 2008, <URL: http://www.datacenterknowledge.com/archives/2008/02/15/major-outage-for-amazon-s3-and-ec2/>.

[Mil09] Rich Miller, Lightning Strike Triggers Amazon EC2 Outage, Data Center Knowledge, June 11, 2009, <URL:

http://www.datacenterknowledge.com/archives/2009/06/11/lightning-strike-triggers-amazon-ec2-outage/>.

[Mod08] Austin Modine, Downed Salesforce Systems Slow Europe and US, The Register, February 11, 2008, <URL: http://www.theregister.co.uk/2008/02/11/salesforce_outages_feb_2008/>.

[MRG10] Online Banking: Browser Security Project, Malware Research Group, Zorin Nexus Ltd., June 2010, <URL: http://malwareresearchgroup.com/wp-content/uploads/2009/01/Online-Banking-Browser-Security-Project-June-201013.zip>.

[Nav10] Eliminating the Data Security and Regulatory Concerns of Using SaaS Applications, White Paper, Navajo Systems, January 2010, <URL: http://www.navajosystems.com/media/Virtual_Private_SaaS_White_Paper.pdf>.

[Obe08a] Jon Oberheide, Evan Cooke, Farnam Jahanian, Empirical Exploitation of Live Virtual Machine Migration, Black Hat Security Conference, Washington, DC, February 2008, <URL: http://www.blackhat.com/presentations/bh-dc-08/Oberheide/Whitepaper/bh-dc-08-oberheide-WP.pdf>.

[Obe08b] Jon Oberheide, Evan Cooke, Farnam Jahanian, CloudAV: N-Version Antivirus in the Network Cloud, USENIX Security Symposium, Association, San Jose, CA, July 28-August 1, 2008, <URL: http://www.eecs.umich.edu/fjgroup/pubs/usenix08-cloudav.pdf>.

[Orm07] Tavis Ormandy, An Empirical Study into the Security Exposure to Hosts of Hostile Virtualized Environments, 2007, <URL: http://taviso.decsystem.org/virtsec.pdf>.

[Ove10] Stephanie Overby, How to Negotiate a Better Cloud Computing Contract, CIO, April 21, 2010, <URL: http://www.cio.com/article/591629/How_to_Negotiate_a_Better_Cloud_Computing_Contract>.

[Pea09] Siani Pearson, Taking Account of Privacy When Designing Cloud Computing Services, International Conference on Software Engineering (ICSE) Workshop on Software Engineering Challenges of Cloud Computing, Vancouver, Canada, May 23, 2009.

[Pon10] Larry Ponemon, Security of Cloud Computing Users, Ponemon Institute, May 12, 2010, <URL: http://www.ca.com/files/IndustryResearch/security-cloud-computing-users_235659.pdf>.

[Pro07] Niels Provos, Dean McNamee, Panayiotis Mavrommatis, Ke Wang, Nagendra Modadugu, The Ghost in the Browser: Analysis of Web-based Malware, Hot Topics

in Understanding Botnets (HotBots), April 10, 2007, Cambridge, Massachusetts, <URL: http://www.usenix.org/events/hotbots07/tech/full_papers/provos/provos.pdf>.

[Pro09] Niels Provos, Moheeb Abu Rajab, Panayiotis Mavrommatis, Cybercrime 2.0: When the Cloud Turns Dark, Communications of the ACM, April 2009.

[Rag09] Steve Ragan, New Service Offers Cloud Cracking for WPA, The Tech Herald, December 8, 2009, <URL: http://www.thetechherald.com/article.php/200950/4906/New-service-offers-cloud-cracking-for-WPA>.

[Rap09] J.R. Raphael, Facebook Privacy Change Sparks Federal Complaint, PC World, February 17, 2009, <URL: http://www.pcworld.com/article/159703/facebook.html?tk=rel_news>.

[Ref10] Security Within a Virtualized Environment: A New Layer in Layered Security, White Paper, Reflex Security, retrieved April 23, 2010, <URL: http://www.vmware.com/files/pdf/partners/security/security-virtualized-whitepaper.pdf>.

[Ris09] Thomas Ristenpart, Eran Tromer, Hovav Shacham, Stefan Savage, Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds, ACM Conference on Computer and Communications Security, November 2009, <URL: http://cseweb.ucsd.edu/~hovav/dist/cloudsec.pdf>.

[Row07] Brent R. Rowe, Will Outsourcing IT Security Lead to a Higher Social Level of Security?, Research Triangle Institute International, July 2007, <URL: http://weis2007.econinfosec.org/papers/47.pdf>.

[Sar10] David Sarno, Los Angeles Police Department Switch to Google E-mail System Hits Federal Roadblock, Los Angeles Times, November 03, 2010, <URL: http://articles.latimes.com/2010/nov/03/business/la-fi-google-la-20101103>.

[Sch00] Bruce Schneier, Crypto-Gram Newsletter, Software Complexity and Security, March 15, 2000, <URL: http://www.schneier.com/crypto-gram-0003.html#8>.

[Sch10] Jeff Schnepper, Don't Like the Tax Law? Don't Shoot the IRS, MSN, March 10, 2010, <URL: http://articles.moneycentral.msn.com/Taxes/blog/page.aspx?post=1692029&_blg=1,1619827>.

[Sha02] Adi Shamir, Cryptography: State of the Science, 2002 ACM Turing Lecture, <URL: http://awards.acm.org/citation.cfm?id=0028491&srt=year&year=2002&aw=140&ao=AMTURING&yr=2002>.

[Sha08] Amit Shah, Kernel-based Virtualization with KVM, Linux Magazine, issue 86, January 2008, <URL: http://www.linux-magazine.com/w3/issue/86/Kernel_Based_Virtualization_With_KVM.pdf>.

[Sec05] VMware Vulnerability in NAT Networking, BugTraq, SecurityFocus, December 21, 2005, <URL: http://www.securityfocus.com/archive/1/420017 and http://www.securityfocus.com/bid/15998/>.

[SECS09] Professional Services Contract, SAAS E-Mail & Collaboration Solution (SECS), City of Los Angeles, November 10, 2009, <URL: https://sites.google.com/a/lageecs.lacity.org/la-geecs-blog/home/faqs-1/C-116359_c_11-20-09.pdf?attredirects=0&d=1>

[She05] Tim Shelton, Remote Heap Overflow, ACSSEC-2005-11-25 - 0x1, <URL: http://packetstormsecurity.org/0512-advisories/ACSSEC-2005-11-25-0x1.txt>.

[Sla09] Marco Slaviero, BlackHat Presentation Demo Vids: Amazon, part 4 of 5, AMIBomb, August 8, 2009, <URL: http://www.sensepost.com/blog/3797.html>.

[Sob06] Charles H. Sobey, Laslo Orto, and Glenn Sakaguchi, Drive-Independent Data-Recovery: The Current State-of-the-Art, IEEE Transactions on Magnetics, February 2006, <URL: http://www.actionfront.com/whitepaper/Drive%20Independent%20Data%20Recovery%20TMRC2005%20Preprint.pdf>.

[Sto02] Gary Stoneburner, Alice Goguen, and Alexis Feringa, Risk Management Guide for Information Technology Systems, SP 800-30, NIST, July 2002, <URL: http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf>.

[Sto10] Jon Stokes, EMC's Atmos Shutdown Shows Why Cloud Lock-in Is Still Scary, Ars Technica, July 2010, <URL: http://arstechnica.com/business/news/2010/07/emcs-atmos-shutdown-shows-why-cloud-lock-in-is-still-scary.ars>.

[Sut09] John D. Sutter, Twitter Hack Raises Questions about 'Cloud Computing', CNN, July 16, 2009, <URL: http://edition.cnn.com/2009/TECH/07/16/twitter.hack/>.

[UCG10] Cloud Computing Use Cases White Paper, Version 4.0, Cloud Computing Use Case Discussion Group, July 2, 2010, <URL: http://opencloudmanifesto.org/Cloud_Computing_Use_Cases_Whitepaper-4_0.pdf>.

[Val08] Craig Valli, Andrew Woodward, The 2008 Australian Study of Remnant Data Contained on 2nd Hand Hard Disks: The Saga Continues, The 6th Australian Digital Forensics Conference, Perth, Western Australia, December 1-3, 2008, <URL: http://conferences.secau.org/proceedings/2008/forensics/Valli%20and%20Woodward%202008%20remnant%20Data%20saga%20continues.pdf>.

[Vaq09] Luis M. Vaquero1, Luis Rodero-Merino1, Juan Caceres, Maik Lindner, A Break in the Clouds: Towards a Cloud Definition, Computer Communication Review (CCR) Online, Short technical Notes, January 2009, <URL: http://ccr.sigcomm.org/online/files/p50-v39n1l-vaqueroA.pdf>.

[Vie09] Kleber Vieira, Alexandre Schulter, Carlos Westphall, Carla Westphall, Intrusion Detection Techniques in Grid and Cloud Computing Environment, IT Professional, IEEE Computer Society, August 26, 2009.

[Vmw09] VMware Hosted Products and Patches for ESX and ESXi Resolve a Critical Security Vulnerability, VMware Security Advisory, VMSA-2009-0006, <URL: http://www.vmware.com/security/advisories/VMSA-2009-0006.html>.

[Vmw10] VMware vShield: Virtualization-Aware Security for the Cloud, product brochure, 2010, <URL: http://www.vmware.com/files/pdf/vmware-vshield_br-en.pdf>.

[Wai08] Phil Wainewright. Many Degrees of Multi-tenancy, ZDNET News and Blogs, June 16, 2008, <URL: http://blogs.zdnet.com/SAAS/?p=533>.

[Wal10] Hannah Wald, Cloud Computing for the Federal Community, IAnewsletter, Vol. 13, No. 2, Information Assurance Technology Analysis Center, Spring 2010.

[Wei09] Jinpeng Wei, Xiaolan Zhang, Glenn Ammons, Vasanth Bala, Peng Ning, Managing Security of Virtual Machine Images in a Cloud Environment, ACM Cloud Computing Security Workshop (CCSW'09) , Chicago, Illinois, November 13, 2009.

[Whi09] Lance Whitney, Amazon EC2 Cloud Service Hit by Botnet, Outage, December 11, 2009, CNET News, <URL: http://news.cnet.com/8301-1009_3-10413951-83.html>.

[Wil10] Matt Williams, All Eyes are on Los Angeles as City Deploys Cloud-Based E-Mail, Government Technology, February 10, 2010, <URL: http://www.govtech.com/gt/744804?id=744804&full=1&story_pg=1>.

[Xen08] Xen Architecture Overview, Version 1.2, Xen Wiki Whitepaper, February 13, 2008, <URL: http://wiki.xensource.com/xenwiki/XenArchitecture?action=AttachFile&do=get&target=Xen+Architecture_Q1+2008.pdf>.

[You07] Greg Young, Neil MacDonald, John Pescatore, Limited Choices are Available for Network Firewalls in Virtualized Servers, Gartner, Inc., ID Number: G00154065, December 20, 2007, <URL: http://www.reflexsystems.com/Content/News/20071220-GartnerVirtualSecurityReport.pdf>.

[You08] Lamia Youseff, Maria Butrico, Dilma Da Silva, Toward a Unified Ontology of Cloud Computing, Grid Computing Environments Workshop (GCE08), held in conjunction

with SC08, November 2008, <URL: http://www.cs.ucsb.edu/~lyouseff/CCOntology/CloudOntology.pdf>.

[Zet09a] Kim Zetter, FBI Defends Disruptive Raids on Texas Data Centers, Wired Magazine, April 7, 2009, <URL: http://www.wired.com/threatlevel/2009/04/data-centers-ra/>.

[Zet09b] Kim Zetter, Bank Sends Sensitive E-mail to Wrong Gmail Address, Sues Google, Wired Magazine, September 21, 2009, <URL: http://www.wired.com/threatlevel/2009/09/bank-sues-google/>.

Appendix A—Acronyms

CAPTCHA	Completely Automated Public Turing test to tell Computers and Humans Apart
CRM	Customer Relationship Management
FISMA	Federal Information Security Management Act
FOIA	Freedom of Information Act
FTP	File Transfer Protocol
HIPPA	Health Insurance Portability and Accountability Act
HVAC	Heating, Ventilation, and Air Conditioning
IA	Information Assurance
IaaS	Infrastructure as a Service
MX	Mail eXchange
NARA	National Archives and Records Administration
NAT	Network Address Translation
OMB	Office of Management and Budget
PaaS	Platform as a Service
PCI DSS	Payment Card Industry Data Security Standard
PII	Personally Identifiable Information
SaaS	Software as a Service
SAS	Statement on Auditing Standards
SECS	SaaS E-mail and Collaboration Solution
SAML	Security Assertion Markup Language
SLA	Service Level Agreement
SOAP	Simple Object Access Protocol
WPA	WiFi Protected Access
XACML	eXtensible Access Control Markup Language
XML	eXtensible Markup Language

Appendix B—Online Resources

The table below contains a list of online resources that may be helpful to security professionals and other readers of this publication in gaining a greater understanding of cloud computing security and privacy issues.

Resource Description	URL
Top Threats to Cloud Computing, V1.0, Cloud Security Alliance, March 2010	http://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf
Security Guidance For Critical Areas of Focus in Cloud Computing, V2.1, Cloud Security Alliance, December 2009	http://www.cloudsecurityalliance.org/csaguide.pdf
Cloud Computing Risk Assessment, European Network and Information Security Agency, November 2009	http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport
The Future of Cloud Computing, Version 1.0, Commission of the European Communities, Expert Group on Cloud Computing, January 2010	http://cordis.europa.eu/fp7/ict/ssai/docs/cloud-report-final.pdf