

ITL BULLETIN FOR JUNE 2012

CLOUD COMPUTING: A REVIEW OF FEATURES, BENEFITS, AND RISKS, AND RECOMMENDATIONS FOR SECURE, EFFICIENT IMPLEMENTATIONS

Shirley Radack, Editor
Computer Security Division
Information Technology Laboratory
National Institute of Standards and Technology
Department of Commerce

Cloud computing is an emerging area of distributed computing that offers many potential benefits to organizations by making information technology (IT) services available as a commodity. When they contract for cloud services, such as applications, software, data storage, and processing capabilities, organizations can improve their efficiency and their ability to respond more quickly and reliably to their customers' needs. At the same time, there are risks to be considered, including maintaining the security and privacy of systems and information, and assuring the wise expenditure of IT resources.

Cloud computing is not a single type of system, but it encompasses a range of underlying technologies and configuration options. The strengths and weaknesses of the different cloud technologies, configurations, service models, and deployment methods should be considered by organizations evaluating services to meet their requirements.

The U.S. Office of Management and Budget (OMB) has identified cloud computing as a tool for helping federal government agencies provide reliable, innovative, and timely services, especially when resources are constrained. The National Institute of Standards and Technology (NIST) has been working in collaboration with government, industry, and standards bodies to accelerate the federal government's secure adoption of cloud computing. NIST's goals are to foster cloud computing systems and practices; support interoperability, portability, and security requirements; and cooperate in the development of needed standards and guidelines.

NIST Special Publication (SP) 800-146, Cloud Computing Synopsis and Recommendations: Recommendations of the National Institute of Standards and Technology

The Information Technology Laboratory (ITL) at NIST recently issued a new publication that explains the different cloud computing technologies and configurations, and recommends methods and approaches that organizations should consider when making decisions about implementing cloud computing. NIST SP 800-146, which was written by Lee Badger and Tim Grance of NIST, Robert Patt-Corner of Global Tech, Inc., and Jeff Voas of NIST, also identifies areas that need further analysis.

The publication includes an explanation of the commercial terms that are frequently used in service agreements between subscribers and providers of cloud computing systems. One section is devoted to a discussion of how cloud computing may be deployed and the general issues associated with each deployment method. Other sections of the publication explain the various service models for cloud computing and their different strengths and weaknesses.

A section of the new publication details the open issues that remain to be resolved for cloud computing: computing performance; cloud reliability; economic goals; compliance; and data and application security. General recommendations for the implementation of cloud computing cover the areas of management, data governance, security and reliability, virtual machines, and software and applications. The appendices to SP 800-146 include an example of the different costs and benefits that can be associated with cloud computing; a discussion of roles and responsibilities of providers and subscribers of cloud services in protecting the security and privacy of information systems; a list of acronyms and abbreviations used; a glossary of terms used; and a list of references.

NIST SP 800-146 is available [here](#).

Characteristics of Cloud Computing

NIST has defined cloud computing 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 service provider interaction.” For a discussion of this definition, see NIST SP 800-145, *The NIST Definition of Cloud Computing*, listed in the For More Information section below.

With cloud computing, organizations can have **on-demand self-service** for computing capabilities, such as server time and network storage when needed, and through a single provider.

Capabilities for different platforms, such as mobile phones, laptops computers, and personal digital assistants, are available **through broad network access**.

The provider’s **computing resources are pooled** to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. While the location of the resources, such as storage, processing, memory, network bandwidth, and virtual machines, is not controlled by the subscriber, it may be possible for the subscriber to specify the country, state, or data center that provides the cloud services.

Cloud capabilities can be provided to the subscriber **rapidly and elastically**, allowing the subscriber to either increase or decrease services. The capabilities available often appear to be unlimited to the subscriber and can be purchased in any quantity at any time.

Cloud systems automatically control and optimize resource use through a **measured service** capability that is appropriate for the type of service provided. Resource usage can be monitored, controlled, and reported providing transparency for both the provider and the consumer of the utilized service.

Issues to be Considered in the Deployment of Cloud Computing

Cloud computing allows computer users to conveniently rent access to fully featured applications, to software development and deployment environments, and to computing infrastructure assets such as network-accessible data storage and processing.

When considering the move to cloud computing, organizations should evaluate the different technologies and configurations, and determine the specific parts of the cloud computing spectrum that meet their needs. The factors to be considered include:

Deployment Models. Depending on the kind of cloud deployment, the cloud may have limited private computing resources, or may have access to large quantities of remotely accessed resources. The following deployment models present a number of trade-offs in how customers can control their resources, and the scale, cost, and availability of resources.

- **Private cloud.** The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise.
- **Community cloud.** The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on premise or off premise.
- **Public cloud.** The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.
- **Hybrid cloud.** The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but that are bound together by standardized or proprietary technology enabling data and application portability.
- **On-site private cloud.** The security perimeter for this deployment model extends around both the subscriber's on-site resources and the private cloud's resources. The private cloud may be centralized at a single subscriber site or may be distributed over several subscriber sites. The subscriber implements the security perimeter, which will not guarantee control over the private

cloud's resources, but will enable the subscriber to exercise control over resources entrusted to the on-site private cloud.

Service Models. The following service models have different strengths and are suitable for different customers and business objectives. In general, interoperability and portability of customer workloads are more achievable in the Infrastructure as a Service (IaaS) service model because the building blocks of this service are relatively well-defined.

- **Cloud Software as a Service (SaaS).** The subscriber uses the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a Web browser. The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, storage, or individual application capabilities. It might be possible for the subscriber to specify application configuration settings.

- **Cloud Platform as a Service (PaaS).** This service allows the subscriber to deploy onto the cloud infrastructure applications that the subscriber created or acquired using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

- **Cloud Infrastructure as a Service (IaaS).** This service enables the subscriber to use processing, storage, networks, and other fundamental computing resources, and to deploy and run other software, including operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components, such as host firewalls.

Economic Considerations. In outsourced and public deployment models, cloud computing provides convenient rental of computing resources: users pay service charges while using a service but need not pay large up-front acquisition costs to build a computing infrastructure. The reduction of up-front costs reduces the risks for pilot projects and experimental efforts, and enhances organizational flexibility. In outsourced and public deployment models, cloud computing allows the customer to request, receive, and later release as many resources as needed. This elasticity can enable the customer to avoid excessive costs from over-provisioning capacity to meet peak demand but not using the capacity in nonpeak periods. A careful analysis of the costs of operation, compliance, and security, including costs to migrate to and, if necessary, migrate from a cloud, is necessary to determine overall costs.

Operational Characteristics. Cloud systems generally depend on networking and any limitations on networking, such as data import/export bottlenecks or service disruptions, reduce cloud utility, especially for applications that are not tolerant of disruptions.

Service-Level Agreements (SLAs). Organizations should understand the terms of the service agreements that define the legal relationships between cloud customers and cloud providers. Organization should also understand customer responsibilities, and those of the service provider, before using a cloud service.

Security. Cloud computing systems are complex networked systems that are affected by traditional computer and network security issues, such as the need to provide data confidentiality, data integrity, and system availability. By imposing uniform management practices, cloud providers may be able to improve on some security update and response issues. Clouds, however, have the potential to aggregate private, sensitive information about customers in cloud data centers. Cloud providers must assure the subscriber that they can keep customer data isolated and protected. Since cloud users and administrators rely heavily on Web browsers, browser security failures can lead to cloud security breaches. The privacy and security of cloud computing depend primarily on whether the cloud service provider has implemented robust security controls and a sound privacy policy required by their customers. Customers need confidence and transparency about the performance of the cloud system and how well it is managed.

The move to cloud computing is a business decision that takes into account the readiness of existing applications for cloud deployment, the transition and life-cycle costs, the maturity of service orientation in the existing infrastructure, and the organization's security and privacy requirements.

Open Issues

Cloud computing may not be a solution for all organizations, nor is it appropriate for all applications. Many of the open issues with respect to deploying cloud computing are similar to concerns that apply to other IT hosted services. Complex computing and software systems may contain flaws, fail, or compromise security. As a result, techniques for detecting failures, understanding their consequences, isolating their effects, and remediating them, are central to the wide-scale adoption of clouds.

Cloud computing has potential to foster more efficient markets through swift leasing of computing resources. Consumers may be able to forgo capital expenses in exchange for variable service fees. Providers of cloud computing can leverage capital expenses to serve many clients. These economic issues become mixed with the complexities of network and system configurations as well as the risks of exposing data and software assets to an external party.

The technical means used to provide the quality of service promised by cloud providers are usually not disclosed to the consumer, thus raising questions about how consumers can verify that the promised quality of service has been provided. Additionally, efficient

markets rely on consumers' ability to compare service offerings. This is difficult since service agreements may not adhere to standard metrics, terminology, and vocabularies.

Different types of applications require differing levels of **computing system performance**. For example, email services may tolerate short service interruptions, but industrial automation and real-time processing generally require both high performance and a high degree of predictability. These performance issues are similar to performance issues for other forms of distributed computing.

Reliability refers to the probability that a system will offer failure-free service for a specified period of time within the bounds of a specified environment. **Cloud reliability** is a function of the reliability of four individual components: the hardware and software facilities offered by providers; the provider's personnel; connectivity to the subscribed services; and the consumer's personnel.

Cloud computing offers an opportunity for consumers to meet **economic goals** by using computing resources with small or modest up-front costs; also cloud computing promotes business agility by reducing the costs of pilot efforts, and may reduce costs to consumers through economies of scale. Although the benefits can be substantial, a number of economic risks must be considered as well.

When data or processing is moved to a cloud, the consumer retains the ultimate responsibility for **compliance** with established policies and regulations but the provider who has direct access to the data may be in the best position to enforce compliance rules. The issues of compliance should be addressed in the contract for cloud services. The General Services Administration (GSA) through the Federal Risk and Authorization Management Program (FedRAMP) is working towards a more consolidated set of compliance requirements and efficient reuse of compliance information. This governmentwide program provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services.

Information security involves protecting the confidentiality and integrity of data and ensuring data availability. An organization that owns and runs its IT operations will apply organizational and administrative controls, specifying who can perform data-related operations such as creation, access, disclosure, transport, and destruction of data; physical controls to protect storage media and devices; and technical controls for identity and access management, data encryption, and other data protection measures.

When an organization subscribes to cloud services, all of the data generated and processed will physically reside in systems that are owned and operated by a provider. An open issue is whether the consumer can obtain assurance that the provider is implementing the same or equivalent controls as the consumer would have implemented. Other considerations include the quality of a cloud's implementation, the vulnerability of the cloud to attack, system complexity, and the expertise level of cloud administrators.

General Recommendations

Under the Federal Information Security Management Act (FISMA) of 2002, federal agencies are required to apply a risk-based policy to achieve cost-effective results for the security of their information and information systems. Standards and guidelines developed by NIST help agencies to carry out effective information security programs based on the management of risk. SP 800-146 includes general recommendations that supplement existing NIST standards and recommendations for cloud systems. See Section 9 of the publication for recommendations in the areas of management, data governance, security and reliability, virtual machines, and software and applications.

For More Information

Federal Information Processing Standards (FIPS) and NIST Special Publications (SPs), including the following, are referenced in NIST SP 800-146:

FIPS 199, *Standards for Security Categorization of Federal Information and Information Systems*

FIPS 200, *Minimum Security Requirements for Federal Information and Information Systems*

NIST SP 800-41, Revision 1, *Guidelines on Firewalls and Firewall Policy*

NIST SP 800-53, Revision 3, *Recommended Security Controls for Federal Information Systems and Organizations*

NIST SP 800-144, *Guidelines on Security and Privacy in Public Cloud Computing*

NIST SP 800-145, *The NIST Definition of Cloud Computing*

For information about these NIST standards and recommendations, as well as other security-related publications, see [here](#).

Information about FEDRAMP is [here](#).

ITL Bulletin Publisher: Elizabeth Lennon
Writer/Editor
Information Technology Laboratory
National Institute of Standards and Technology
Email [here](#).

Disclaimer

Any mention of commercial products or reference to commercial organizations is for information only; it does not imply recommendation or endorsement by NIST, nor does it imply that the products mentioned are necessarily the best available for the purpose.