Chapter 1:

Cloud Computing Concepts

Cloud Computing Concepts

1.1 Learning Objectives

After completing this chapter, learners will be able to:

1. Define cloud technology.
2. Articulate the cloud service model characteristics and capabilities.
3. Identify the components that make up cloud technology.
4. Explain what considerations need to be made for cloud service implementation.
5. Distinguish different cloud deployment models.

1.2 Overview

In many countries, such as Australia, government agencies have an explicit obligation to consider cloud services when procuring new information and communication technology (ICT) requirements for their test and development needs, and to migrate public-facing websites to public cloud services. The agencies must choose cloud services when they represent the best value and adequate risk management compared to other available options.

As defined by the US National Institute of Standards and Technology (NIST), cloud computing is a "model for enabling convenient, on-demand network access to a shared pool of configurable computing resources, such as, networks, servers, storage, applications, and services, that can be rapidly provisioned and released with minimal management effort or service provider interaction."[1]

Adding computer virtualization to the NIST definition (on-demand computer resources requiring minimal management effort) creates a cloud computing model that offers enterprises virtual processing power in a variety of possible implementations.

While the models and types of cloud services vary significantly, they share five essential characteristics:

- **On-demand self-service**—Computing capabilities can be provisioned without human interaction from the service provider.
- **Broad network access**—Computing capabilities are available over the network and can be accessed by diverse client platforms.
- **Resource pooling**—Computer resources are pooled to support a multi-tenant model.
- **Rapid elasticity**—Resources can scale up or down rapidly and, in some cases, automatically, in response to business demands.
- **Measured service**—Resource utilization can be optimized by leveraging charge-per-use capabilities.

While cloud computing is often categorized as exclusively an IT issue, such thinking underplays its potential impact. It also represents a weighty business decision for an organization, one that requires consideration and input from many stakeholders. These stakeholders and the enterprises they support are often drawn to the potential benefits of cloud computing, summarized in **figure 1.1**.

Figure 1.1—Cloud Benefits

Benefit	Description
Tangible	
Cost reduction	Computing cost is shifted from capital expenditure to operational cost because the cloud provider supplies the underlying infrastructure as part of the service bundle. In addition, the cloud promises cost reduction in the following areas: • Labor (IT system administration hours/headcount) • Application software (Software as a Service [SaaS] only) • Licensing purchase and maintenance • Technical support and user support • Maintenance (upgrades, updates, patches, etc.) • Hosting (physical building, power, cooling, etc.)
Enhanced productivity	User mobility and ubiquitous access can increase productivity. Collaborative applications increase productivity and reduce rework.
Optimized resource utilization	Enterprises use only the computing resources they need, thus reducing system idle time waste.
Improved security/ compliance	Cloud providers may offer robust security controls as a market differentiation.
Access to skills and capabilities	Customers benefit from top-notch skills and capabilities while avoiding employment costs (recruiting, salary, benefits, training, etc.)
Scalability	On-demand provisions or computing resources eliminate the cost of capacity planning.
Agility	Agility contributes to cost reduction and productivity enhancement due to faster provisioning of systems: • Faster application deployment (SaaS) • Faster application development/testing (Platform as a Service [PaaS])
Customer satisfaction	Effective utilization of cloud applications can increase collaboration between the enterprise and its customers or reduce response time to customer inquiries.
Reliability	Cloud providers have redundant sites that can address business continuity and disaster recovery in a more efficient manner.

Performance	Better performance and uptime can result from continuous and consistent operations monitoring by the cloud provider.
Intangible	
New business opportunities	A cloud application (SaaS) may be the critical element to land a new business or expand into new markets.
Focus on core business	IT resources can be allocated to support core business functions.
Employee satisfaction/ innovation	Mobility and faster performance can improve employee satisfaction and boost innovation.
Collaboration	Real-time collaboration can increase quality and innovation.
Risk transfer	Some risk (e.g., security breaches, data loss, disaster recovery) can be transferred to the cloud service provider (CSP); this could represent a tangible or intangible benefit.

Source: ISACA, *Controls and Assurance in the Cloud: Using COBIT® 5*, USA, 2014, https://www.isaca.org/bookstore/cobit-5/cb5ca

Depending on business needs, any or all of these benefits could be a sufficient reason to consider a cloud computing solution. However, cloud services can also generate significant risk, and any enterprise contemplating cloud usage must be prepared to address it. If not handled well, the risk can quickly turn the cloud experience into an information security management nightmare derived from the loss of controls over physical and logical assets.

1.3 Components of the Cloud

Virtualization. Hypervisors. Containers. These are just some of the terms encountered when entering the cloud computing space. Because there is overlap among them, it is easy to assume they are all basically the same thing, performing the same function. However, that is an oversimplification.

The ISACA Glossary defines virtualization as "[t]he process of adding a 'guest application' and data onto a 'virtual server,' recognizing that the guest application will ultimately part company from this physical server." Other definitions cite its use of software on a single, physical hardware system to create a virtual, rather than actual, version of something, such as an operating system, storage, a server or a network. "Virtual" is the opposite of "physical," not of "real." It is a technology, not an environment.

In enterprise information technology (IT), virtualization alters the technical architecture because it allows different resources to be executed in a single (or multilayer) environment. In general, it turns one piece of hardware into the host for many other pieces and, consequently, over time, reduces the enterprise's capital expense, administration cost and other financial costs on the profit and loss statement.

Virtualization, as a term and a concept, has broad utility and can be applied to several areas: virtualized servers, virtualized storage, virtualized processors, virtual memory, virtual desktops, virtualized network, etc. Because of its extensive applications and cost savings, it is being evaluated by chief information officers (CIOs) worldwide as they strategize how to provide agility and computing power to meet their enterprise needs. In addition, because organizations today require a quick and reliable way to provision technical resources that enable a faster time to market, virtualization is on the C-level agenda and is already enhancing the effectiveness of many enterprises around the globe.

Originally virtualization was used primarily to facilitate server consolidation, but now many other approaches present themselves. Virtualization starts during the technical environment design phase, when the design team considers how to support the business processes and identifies assets needed to convert the plan to reality. During this phase, enterprises often realize that the life cycle of provisioning the right hardware and other equipment can go on longer than expected and there is a faster way to deploy them: by building in some abstraction from the physical world and hosting different virtual resources within the boundaries of a unique physical resource.

There are many benefits of implementing virtualization within enterprise IT, as itemized below and summarized in **figure 1.2**.

- **Cost reduction**—By consolidating many instances of (virtualized) servers on a physical one, enterprises lower their hardware expenditures. In addition to lower capital expenditures, virtualized environments enable enterprises to save on maintenance and energy, often resulting in a reduced total cost of ownership.
- **Automation**—Technology allows some virtualized environments to be provisioned as needed and on the fly, thus facilitating automation of business processes and eliminating the need to continually resource and manage portions of the technical environment that support sporadic business needs. Some virtualization technology facilitates the automatic allocation of a process for its optimal performance within a pool of virtualized environments.
- **Responsiveness**—Since the virtual environment can provision itself to get the best out of available resources, response times are faster and downtimes can be reduced to near zero, improving agility and performance.
- **Decoupling**—Processes that once needed to exist within the same physical machine can now be easily separated while still maintaining the robustness and security required. The different virtualized worlds (network, operating system [OS], database, application, etc.) can be decoupled (even distributed in different geographies) without threatening integrity in the process.
- **Flexibility**—The relatively easy creation or preparation of the right environment for the right application enables enterprises to provide flexibility to the infrastructure, not only in the test or preproduction phases but also in the production area. When a new procedure or technical/business requirement arises, virtualization's ability to enable rapid creation of the environment allows the business to test the environment without having to wait for the regular provisioning process to be executed and delivered.
- **Agility**—Agility facilitates quick adaptation to business needs—for example, if orders peak and additional computing power is needed. An enterprise may even choose to overcommit the resources of a physical machine, since virtualization facilitates rapid movement of the different resources that "live" in one physical machine to other virtual machines. In this way, virtualization supports alignment with business needs.
- **Workload balancing**—Deploying several virtual environments guarantees the good practices of high availability, redundancy and failover since workloads can go where they are more efficient. Thus, virtualization focuses not only on effectiveness (doing the right things) but also efficiency (doing things in a faster, cheaper and more reliable way).
- **Simplification**—Virtual IT is still IT, so some of the typical IT difficulties exist even within a virtual environment. However, reducing the number of physical servers significantly reduces the probability of failure and the cost of management and results in simplification—one of the promises of virtualization.
- **Space utilization**—Server consolidation saves space in the data center and facilitates scalability since many servers exist within one server.
- **Sustainability**—Virtualized environments use less environmental resources. Energy consumption in data centers is often wasted on machines that are consistently underutilized. Since virtualization allows for many virtual machines to run on one physical machine, less energy is needed to power and cool devices.

Figure 1.2—Five Reasons to Virtualize

Outcome	How It Is Achieved
1. Reduce IT complexity.	Applications and their operations systems are encapsulated in virtual machines that are defined in software, making them easy to provision and manage.
2. Enable standardization.	Since applications are decoupled from hardware, the data center may converge on a narrower range of hardware devices.
3. Improve agility.	Applications and virtual machines can be copied and moved in real time—and in the cloud—in response to changing business conditions.
4. Improve cost efficiency.	Virtual machines can easily be moved to consume spare capacity wherever it exists, thus generating more work from less hardware.
5. Facilitate automation.	Virtual infrastructure is easily provisioned and orchestrated by software-driven processes, especially when the underlying hardware is standardized.

As with any technology, there is risk associated with virtualization. The risk can be categorized into three groups:

- **Attacks on virtualization infrastructure**—There are two primary types of attacks on virtualization infrastructure: hyperjacking and virtual machine (VM) jumping (or guest-hopping).
- **Attacks on virtualization features**—Although there are multiple features of virtualization that can be targeted for exploitation, the more common targets include VM migration and virtual networking functions.
- **Compliance and management challenges**—Compliance auditing and enforcement, as well as day-to-day system management, are challenging issues when dealing with virtualized systems. VM sprawl introduces a challenge to the enterprise: Because VMs are much easier to provision and deploy than physical systems, the number and types of VMs can easily get out of hand. Furthermore, VM provisioning is often administered by different groups within an organization, making it difficult for the IT function to control what applications, OSs and data are deployed to them.

A hypervisor is computer software that allows various software applications running on different operating systems to coexist on the same server at the same time. The hypervisor and the guests it supports are software and, as such, need to be regularly patched and hardened. Cloud infrastructures are based on hypervisors and are controlled through a central hypervisor manager or client. They provide the link between virtual machines and the underlying physical resources required to run the machines by using hypercalls (similar to system calls, but for virtualized systems). As such, they are vital for cloud virtualization.

A container is a "virtual runtime environment that runs on top of a single operating system (OS) kernel and emulates an operating system rather than the underlying hardware."[2] Virtualization provides an abstract machine; a container provides an abstract operating system. Therefore, applications that are running in a container share an operating system, while VM systems can run different operating systems.

In summary, these technologies constitute a kind of "evolution" of the cloud. Cloud computing (as it is known today) sprung from these technologies and their availability. Virtualization of hardware paved the way for a single system being used for multiple tasks and purposes. Hypervisors are the software that can be used to virtualize depending on the application. Containers take it a step further and allow for minimal resource use to run full services. Together, they are used with automation, scaling and redundancy to provide the technical cogs of the cloud.

1.4 Cloud Service Implementation Considerations

In preparing to implement cloud solutions, an enterprise must consider some of the challenges that may arise. For example, the cloud computing implementation may conflict with the enterprise's culture. The cloud solution may struggle to accommodate changes in requirements identified after implementation. The implementation may have failed to recognize and establish controls to offset known and future cloud risk. Enterprise and stakeholder expectations for the implementation may not have been thoroughly understood or communicated.

An effective implementation may require the establishment of an implementation team consisting of appropriate members, creating a trusted environment and establishing common goals and effectiveness measures. The team members—and anyone with a role to play in implementation—must be empowered to do their best work by ensuring that accountabilities are assigned, training is provided, and organizational structures and human resource (HR) processes are aligned. The team should be encouraged to identify and communicate short-term wins that maintain motivation and support change enablement. While the team members cannot be expected to have the answers and solutions to all issues that arise before and during implementation, among the many questions they must consider are:

- Which security policies apply within the enterprise? Which regulatory restrictions apply to the enterprise and to any locations where a CSP might reside?
- How will moving to the cloud influence the enterprise's processes? Which processes depend on assets that could move to the cloud? Are these processes considered to be critical for the business?
- How will the relationship with the CSP be managed? How are roles and responsibilities defined?
- How will change within the enterprise be managed? How can an information culture be imposed upon the CSP?
- Which assets are considered for cloud computing?
- Which service capabilities are expected of the CSP? How will performance be measured? How will issues be reported?
- Which skills and competencies are required to manage the enterprise assets? Does the enterprise wish to keep these in-house after a decision has been made to move to the cloud?

Because the amount of work and the complexity involved in preparing for and carrying out a successful implementation, some enterprises may choose to engage professional services to assist the in-house team in managing the transition to the cloud.

Migrating the data is a significant aspect of cloud implementation. Consideration of migration issues should be addressed as early as the drafting of the service contract terms. The contract should provide for continuity of accessibility, usability and preservation of all enterprise data, regardless of any migration of data to other formats during the contract. Terms should provide for appropriate testing to ensure data integrity prior to any migration. It should also include provisions relating to the migration of data to new formats, when appropriate, and the provision of proper documentation about migration activities to the enterprise. Because even the best of business relationships sometimes come to an end, the contract should also define provisions relating to the migration of data on termination of the contract.

Successful migration begins with a plan that aligns with the enterprise's program and project portfolios. Close coordination among all facets of migration will help ensure that value is delivered and the required resources are available when needed to perform the necessary work. Preparation for migration should be an integral part of the enterprise's development methods and may include audit trails and a recovery plan should the migration fail.

The impact of migration depends on the cloud service model and deployment model being used, as they will indicate an appropriate balance for protection of organizational assets during migration. The enterprise should classify its assets into categories for an optimal selection of cloud arrangements. Generally, data can be classified as public, restricted, for internal use, secret and top-secret. A data life-cycle process can also be defined.

Two good rules of thumb, regardless of service model and deployment selected, are:

- Prior to migration, encrypt all sensitive assets being migrated to the CSP to prevent disclosure.
- Test changes independently in accordance with a defined test plan prior to migration to the live operational environment.

A post-implementation review should be conducted to confirm the outcome and results of the cloud implementation, identify lessons learned and develop action plans for future implementations. The review should also evaluate and check the actual performance and outcomes of the cloud solution against the predicted performance and outcomes and identify opportunities for improvement.

1.5 Cloud Deployment Models

Cloud customers have their choice of four cloud deployment models (see **figure 1.3**):

- **Private cloud**—One client (one individual or one enterprise). Several departments or divisions may be represented, but all exist within the same enterprise. Private clouds often employ virtualization within an enterprise's existing computer servers to improve computer utilization. A private cloud also typically involves provisioning and metering components, enabling rapid deployment and chargeback where appropriate. This model is most closely related to traditional IT outsourcing models in the marketplace, but it can also be an enterprise's internal delivery model. Because of this, the private model tends to offer the least amount of new risk and security challenge.
- **Public cloud**—An offering from one CSP to many clients who share the cloud processing power and storage capacity concurrently. Public cloud clients share applications, processing power and data storage space communally. Client data are commingled, but segregated, for example, through the use of metatags. This model tends to present the greatest assurance challenges for security and risk managers.
- **Community cloud**—A private-public cloud with clients who have a common connection or affiliation, such as a trade association, industry or locality. The community cloud business model allows a CSP to provide cloud tools and applications that are specific to the needs of the community.
- **Hybrid cloud**—A combination of two or more of the previously defined deployment models. Each of the three cloud deployment models has specific advantages and disadvantages that are relative to the other deployment models. A hybrid cloud leverages the advantage of the other cloud models, providing a more optimal user experience.

Figure 1.3—Cloud Deployment

Deployment	Description	Examples
Private cloud	• Operated solely for one enterprise • May be managed by the enterprise or a third party • May exist on- or off-premise	• Hewlett Packard enterprise (HPE) • VMWare • Oracle
Public cloud	• Made available to the general public or a large industry group • Owned by an organization selling cloud services	• Amazon Web Services (AWS) • Microsoft (Azure) • Google Cloud Platform
Community cloud	• Shared by several enterprises • Supports a specific community that has a shared mission or interest • May be managed by the enterprises or a third party	• eCommerce • HR • Manufacturing
Hybrid cloud	A combination of two or more cloud deployment models (private, community or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability, e.g., cloud bursting for load balancing between clouds	

Source: ISACA, *Controls and Assurance in the Cloud: Using COBIT® 5*, USA, 2014, https://www.isaca.org/bookstore/cobit-5/cb5ca

The costs of services vary across the different deployment models. Private cloud deployments are often the costliest option; public clouds, the least. The hybrid may offer a reasonable alternative to either option.

Levels of information security also vary among the private, public and community deployment models, with private clouds having the most restricted user access and, most likely, the least exposure to threats.

The hybrid cloud offers a combination of two or more deployment models with varying levels of security, as needed. Users may choose to leverage private or community clouds for their most business-critical data while using the public cloud for data that are already publicly available or for other non-classified data or applications. Some enterprises may just accept the risk and choose to use the public cloud regardless of data classification.

Each enterprise must make its own choice based on its unique situation and needs. Some pertinent questions to ask to determine which deployment model may be the most appropriate include:

- Is the business process to be used in the cloud a nonstandard solution?

- What type and number of interdependencies are there among the business processes involved?
- Is there a difference between the cloud solution and a standard IT-based solution?
- Is the application custom?
- Is the hardware custom?
- Is the OS custom?
- Is the business driver cloud-compatible?

It is important to note that the offerings of CSPs are not necessarily standardized. Prospective cloud clients must use care in determining the amount of security provisioning they will require depending on the type of application and the security classifications of the data they plan to promote to the cloud.

1.6 Chapter 1 Knowledge Check

REVIEW QUESTIONS

1. One of the **MAIN** characteristics of cloud technology service models is its:
   A. On-demand service
   B. Physical device ownership
   C. Measured elasticity
   D. Limited network access

2. Cloud models provide the customer the ability to _____ risk:
   A. Accept
   B. Mitigate
   C. Transfer
   D. Ignore

3. The ability to run a guest operating system is a characteristic of:
   A. Simplification
   B. Elasticity
   C. Resource pooling
   D. Virtualization

4. The difference between a hypervisor and a container is that a container runs on top of a single OS and emulates a guest OS.
   A. True
   B. False

5. When considering a cloud vendor, special attention should be paid to:
   A. The cloud provider's culture
   B. The skill set of your internal IT team
   C. Data migration capabilities
   D. Contract usability

6. After data migration into the cloud, a _____ should be conducted.
   A. Thorough investigation
   B. Post-mortem
   C. Implementation review
   D. Post-implementation review

7. The **BEST** cloud deployment model for an on-premise single enterprise would be:
   A. Private
   B. Public
   C. Hybrid
   D. Community

8. Cloud deployment models are standardized throughout the industry.
   A. True
   B. False

Answers on page 21

Chapter 1 ANSWER KEY

Review Questions

1. **A. On-demand service. Refer to page 12.**
   B. Physical device ownership
   C. Measured elasticity
   D. Limited network access

2. A. Accept
   B. Mitigate
   C. Transfer. Refer to page 13.
   D. Ignore

3. A. Simplification
   B. Elasticity
   C. Resource pooling
   D. Virtualization. Refer to page 14.

4. A. True Refer to pages 15-16.
   B. False.

5. A. The cloud provider's culture
   B. The skill set of your internal IT team
   C. Data migration capabilities. Refer to page 17.
   D. Contract usability

6. A. Thorough investigation
   B. Post-mortem
   C. Implementation review
   D. Post-implementation review. Refer to page 17.

7. **A. Private. Refer to page 17.**
   B. Public
   C. Hybrid
   D. Community

8. **A. True. Refer to page 18.**
   B. False

---

[1] Mell, P.; T. Grance; Special Publication (SP) 800-145 (Draft), "The NIST Definition of Cloud Computing," National Institute of Standards and Technology (NIST), USA, 2011, p. 3, faculty.winthrop.edu/domanm/csci411/Handouts/NIST.pdf

[2] Firesmith, D.; "Virtualization via Containers," Carnegie Mellon University, Software Engineering Institute, 25 September 2017, insights.sei.cmu.edu/sei_blog/2017/09/virtualization-via-containers.html#:~:text=Note%20that%20virtualization%20via%20containers,rather%20than%20the%20underlying%20hardware