



Section 4: Cloud Computing

This section contains one module:

- Module 13 : Cloud Computing

Welcome to Section 4: Cloud Computing. This section covers various aspects of cloud computing.



Module 13: Cloud Computing

Upon completion of this module, you should be able to:

- Explain the characteristics of cloud computing
- Describe cloud services and deployment models
- Describe cloud computing infrastructure
- Discuss the challenges of cloud computing
- Discuss cloud adoption considerations

This module focuses on the essential characteristics of cloud computing, cloud services and deployment models, and cloud computing infrastructure. It also focuses on the challenges of cloud computing and cloud adoption considerations.

Module 13: Cloud Computing

Lesson 1: Cloud Computing Overview

During this lesson the following topics are covered:

- Definition of cloud computing
- Essential characteristics of cloud computing
- Benefits of cloud computing
- Cloud enabling technologies

This lesson covers definition and essential characteristics of cloud computing. It also covers benefits of cloud computing and cloud enabling technologies.

Drivers for Cloud Computing

- Business requirements
 - ▶ Transformation of IT processes to achieve more with less
 - ▶ Better agility and higher availability at reduced expenditure
 - ▶ Reduced time-to-market
 - ▶ Accelerated pace of innovation
- IT challenges to meet business requirements are:
 - ▶ Serving customers worldwide round the clock, refreshing technology quickly, faster provisioning of IT resources – all at reduced cost
- These challenges are addressed with the emergence of cloud computing

In today's competitive environment, organizations are under increasing pressure to improve efficiency and transform their IT processes to achieve more with less. Businesses need reduced time-to-market, better agility, higher availability, and reduced expenditures to meet the changing business requirements and accelerated pace of innovation. These business requirements are posing several challenges to IT teams. Some of the key challenges are serving customers worldwide round the clock, refreshing technology quickly, and faster provisioning of IT resources—all at reduced costs.

These long-standing challenges are addressed with the emergence of a new computing style, called cloud computing, which enables organizations and individuals to obtain and provision IT resources as a service.

What is Cloud Computing?


Cloud Computing

A model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., servers, storage, networks, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

– NIST

- Essential Cloud characteristics

- ▶ On-demand self-service
- ▶ Broad network access
- ▶ Resource pooling
- ▶ Rapid elasticity
- ▶ Measured service



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 5

With cloud computing users can browse and select relevant cloud services, such as compute, software, information storage, or a combination of these resources. via a portal. Cloud computing automates delivery of selected cloud services to the users. It helps organizations and individuals deploy IT resources at reduced total cost of ownership with faster provisioning and compliance adherence. A widely adopted definition of cloud computing comes from the U.S. National Institute of Standards and Technology (NIST Special Publication 800-145) as provided in the slide.

A computing infrastructure used for cloud services must meet certain capabilities or characteristics. According to NIST, the cloud infrastructure should have five essential characteristics:

- On-Demand Self-Service
- Broad Network Access
- Resource Pooling
- Rapid Elasticity
- Measured Service

On-demand Self-service



- Enables consumers to unilaterally provision computing capabilities (examples: server time and storage capacity) as needed automatically
- Consumers view service catalogue via a Web-based user interface and use it to request for a service

A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

A cloud service provider publishes a service catalogue, which contains information about all cloud services available to consumers. The service catalogue includes information about service attributes, prices, and request processes. Consumers view the service catalogue via a web-based user interface and use it to request for a service. Consumers can either leverage the “ready-to-use” services or change a few service parameters to customize the services.

Broad Network Access



- Computing capabilities are available over the network
- Computing capabilities are accessed from a broad range of client platforms such as:
 - ▶ Desktop computer
 - ▶ Laptop
 - ▶ Tablet
 - ▶ Mobile device

Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (example, mobile phones, tablets, laptops, and workstations)

Resource Pooling



- Provider's computing resources are pooled to serve multiple consumers using a multitenant model
- Resources are assigned from the pool according to consumer demand
- Consumers have no control or knowledge over the exact location of the provided resources

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. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (example, country, state, or data center). Examples of resources include storage, processing, memory, and network bandwidth.

Note:

Multitenancy refers to an architecture in which multiple independent consumers (tenants) are serviced using a single set of resources. This lowers the cost of services for consumers. Virtualization enables resource pooling and multitenancy in the cloud. For example, multiple virtual machines from different consumers can run simultaneously on the same physical server that runs the hypervisor.

Rapid Elasticity

- Computing capabilities can be elastically provisioned and released
- Computing capabilities are scaled rapidly, commensurate with consumer's demand
 - ▶ Provides a sense of unlimited scalability



Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

Consumers can leverage rapid elasticity of the cloud when they have a fluctuation in their IT resource requirements. For example, an organization might require doubling the number of web and application servers for a specific duration to accomplish a specific task. For the remaining period, they might want to release idle server resources to cut down the expenses. The cloud enables consumers to grow and shrink the demand for resources dynamically.

Measured Service

- Cloud computing provides a metering system that continuously monitors resource consumption and generates reports
 - ▶ Helps to control and optimize resource use
 - ▶ Helps to generate billing and chargeback reports



Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (example, storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

Benefits of Cloud Computing

Benefits	Description
Reduced IT cost	• Reduces the up-front capital expenditure (CAPEX)
Business agility	• Provides the ability to deploy new resources quickly • Enables businesses to reduce time-to-market
Flexible scaling	• Enables consumers to scale up, scale down, scale out, or scale in the demand for computing resources easily • Consumers can unilaterally and automatically scale computing resources
High availability	• Ensures resource availability at varying levels, depending on consumer's policy and priority

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 11

Cloud computing offers the following key benefits:

- **Reduced IT cost:** Cloud services can be purchased based on pay-per-use or subscription pricing. This reduces or eliminates consumer's IT capital expenditure (CAPEX).
- **Business agility:** Cloud computing provides the capability to allocate and scale computing capacity quickly. Cloud can reduce the time required to provision and deploy new applications and services from months to minutes. This enables businesses to respond more quickly to market changes and reduce time-to-market.
- **Flexible scaling:** Cloud computing enables consumers to scale up, scale down, scale out, or scale in the demand for computing resources easily. Consumers can unilaterally and automatically scale computing resources without any interaction with cloud service providers. The flexible service provisioning capability of cloud often provides a sense of unlimited scalability to the cloud service consumers.
- **High availability:** Cloud computing has the ability to ensure resource availability at varying levels depending on the consumer's policy and priority. Redundant infrastructure components (servers, network paths, and storage equipments, along with clustered software) enable fault tolerance for cloud deployments. These techniques can encompass multiple datacenters located in different geographic regions, which prevents data unavailability due to regional failures.

Cloud Enabling Technologies

Technologies	Description
Grid computing	• Form of distributed computing • Enables resources of numerous computers in a network to work on a single task at the same time
Utility computing	• Service provisioning model that offers computing resources as a metered service
Virtualization	• Abstracts physical characteristics of IT resources from resource users • Enables resource pooling and creating virtual resources from pooled resources
Service-oriented architecture (SOA)	• Provides a set of services that can communicate with each other

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 12

Grid computing, utility computing, virtualization, and service-oriented architecture are enabling technologies of cloud computing.

Grid computing is a form of distributed computing that enables the resources of numerous heterogeneous computers in a network to work together on a single task at the same time. Grid computing enables parallel computing and is best for large workloads.

Utility computing is a service-provisioning model in which a service provider makes computing resources available to customers, as required, and charges them based on usage. This is analogous to other utility services, such as electricity, where charges are based on the consumption.

Virtualization is a technique that abstracts the physical characteristics of IT resources from resource users. It enables the resources to be viewed and managed as a pool and lets users create virtual resources from the pool. Virtualization provides better flexibility for provisioning of IT resources compared to provisioning in a non-virtualized environment. It helps optimizing resource utilization and delivering resources more efficiently.

Service-oriented architecture (SOA) provides a set of services that can communicate with each other. These services work together to perform some activity or simply passes data among services.

Module 13: Cloud Computing

Lesson 2: Cloud Service and Deployment Models


During this lesson the following topics are covered:

- Cloud service models
- Cloud deployment models

This lesson covers three primary cloud service models – Infrastructure-as-a-Service(IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS). This lesson also covers cloud deployment models – Public, Private, Community, and Hybrid.

Infrastructure-as-a-Service

- Consumers deploy their software, including OS and application on provider's infrastructure
 - ▶ Computing resources such as processing power, memory, storage, and networking components are offered as service
 - ▶ Example: Amazon Elastic Compute Cloud
- Consumers have control over the OSs and deployed applications



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 14


The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems and deployed applications; and possibly limited control of select networking components (example, host firewalls).

IaaS is the base layer of the cloud services stack. It serves as the foundation for both the SaaS and PaaS.

Amazon Elastic Compute Cloud (Amazon EC2) is an example of IaaS that provides scalable compute capacity, on-demand, in the cloud. It enables consumers to leverage Amazon's massive computing infrastructure with no up-front capital investment.

Platform-as-a-Service

- Consumers deploy consumer-created or acquired applications onto provider's computing platform
 - ▶ Computing platform is offered as a service
 - ▶ Example: Google App Engine and Microsoft Windows Azure Platform
- Consumer has control over deployed applications



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.


Section 4: Cloud Computing 15

The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, 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 configuration settings for the application-hosting environment.

PaaS is also used as an application development environment, offered as a service by the cloud service provider. The consumer may use these platforms to code their applications and then deploy the applications on the cloud. Because the workload to the deployed applications varies, the scalability of computing resources is usually guaranteed by the computing platform, transparently. Google App Engine and Microsoft Windows Azure Platform are examples of PaaS.

Software-as-a-Service

- Consumers use provider's applications running on the cloud infrastructure
 - ▶ Applications are offered as a service
 - ▶ Examples: EMC Mozy and Salesforce.com
- Service providers exclusively manage computing infrastructure and software to support services



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 16

The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (example, web-based e-mail), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

In a SaaS model, applications, such as Customer Relationship Management (CRM), e-mail, and Instant Messaging (IM), are offered as a service by the cloud service providers. The cloud service providers exclusively manage the required computing infrastructure and software to support these services. The consumers may be allowed to change a few application configuration settings to customize the applications.


EMC Mozy is an example of Software-as-a-Service. Consumers can leverage the Mozy console to perform automatic, secured, online backup and recovery of their data with ease. Salesforce.com is a provider of SaaS-based CRM applications, such as Sales Cloud and Service Cloud.

Cloud Deployment Models

- Public
- Private
- Community
- Hybrid

According to NIST, cloud computing is classified into four deployment models—public, private, community, and hybrid— which provide the basis for how cloud infrastructures are constructed and consumed.

Public Cloud




EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 18

In a public cloud model, the cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

Consumers use the cloud services offered by the providers via the Internet and pay metered usage charges or subscription fees. An advantage of the public cloud is its low capital cost with enormous scalability. However, for consumers, these benefits come with certain risks: no control over the resources in the cloud, the security of confidential data, network performance, and interoperability issues. Popular public cloud service providers are Amazon, Google, and Salesforce.com. Figure in the slide shows a public cloud that provides cloud services to organizations and individuals.

Private Cloud




EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 19

In a private cloud model, the cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (example, business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises. Following are two variations to the private cloud model:

- **On-premise private cloud:** The on-premise private cloud, also known as internal cloud, is hosted by an organization within its own data centers. This model enables organizations to standardize their cloud service management processes and security, although this model has limitations in terms of size and resource scalability. Organizations would also need to incur the capital and operational costs for the physical resources. This is best suited for organizations that require complete control over their applications, infrastructure configurations, and security mechanisms.
- **Externally hosted private cloud:** This type of private cloud is hosted external to an organization and is managed by a third-party organization. The third-party organization facilitates an exclusive cloud environment for a specific organization with full guarantee of privacy and confidentiality.

Community Cloud




EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 20

In a community cloud model, the cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (example, mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

In a community cloud, the costs spread over to fewer consumers than a public cloud. Hence, this option is more expensive but might offer a higher level of privacy, security, and compliance. The community cloud also offers organizations access to a vast pool of resources compared to the private cloud. An example in which a community cloud could be useful is government agencies. If various agencies within the government operate under similar guidelines, they could all share the same infrastructure and lower their individual agency's investment.

Hybrid Cloud



In a hybrid cloud model, the cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (example, cloud bursting for load balancing between clouds).

The hybrid model allows an organization to deploy less critical applications and data to the public cloud, leveraging the scalability and cost-effectiveness of the public cloud. The organization's mission-critical applications and data remain on the private cloud that provides greater security. Figure in the slide shows an example of a hybrid cloud.

Module 13: Cloud Computing


Lesson 3: Cloud Infrastructure, Challenges, and Considerations

During this lesson the following topics are covered:

- Cloud infrastructure
- Challenges of cloud computing
- Cloud adoption considerations

This lesson covers the cloud computing infrastructure, challenges of cloud computing, and cloud adoption considerations.

Cloud Infrastructure Framework



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 23

A cloud computing infrastructure is the collection of hardware and software that enables the five essential characteristics of cloud computing. Cloud computing infrastructure usually consists of the following layers:

- Physical infrastructure
- Virtual infrastructure
- Applications and platform software
- Cloud management and service creation tools

The resources of these layers are aggregated and coordinated to provide cloud services to the consumers.

Physical Infrastructure

- Physical infrastructure includes physical IT resources
 - ▶ Physical servers
 - ▶ Storage systems
 - ▶ Networks
- Physical servers are connected to each other, to the storage systems, and to clients via networks
- Physical resources may be located in a single data center or distributed across multiple data centers

The physical infrastructure consists of physical computing resources, which include physical servers, storage systems, and networks. Physical servers are connected to each other, to the storage systems, and to the clients via networks, such as IP, FC SAN, IP SAN, or FCoE networks.

Cloud service providers may use physical computing resources from one or more data centers to provide services. If the computing resources are distributed across multiple data centers, connectivity must be established among them. The connectivity enables the data centers in different locations to work as a single large data center. This enables migration of business applications and data across data centers and provisioning cloud services using the resources from multiple data centers.

Virtual Infrastructure

- Virtual infrastructure consists of:
 - ▶ Resource pools
 - ▶▶ CPU, memory, network bandwidth, storage pools
 - ▶ Identity pools
 - ▶▶ VLAN ID and VSAN ID pools
 - ▶ Virtual IT resources
 - ▶▶ Virtual Machines (VMs), virtual storage volumes, virtual networks (VLAN and VSAN)
- Virtual IT resources obtain capacity and identity from resource and identity pools respectively

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 25

Cloud service providers employ virtualization technologies to build a virtual infrastructure layer on the top of the physical infrastructure. Virtualization enables fulfilling some of the cloud characteristics, such as resource pooling and rapid elasticity. It also helps reduce the cost of providing the cloud services. Some cloud service providers may not have completely virtualized their physical infrastructure yet, but they are adopting virtualization for better efficiency and optimization.

Virtualization abstracts physical computing resources and provides a consolidated view of the resource capacity. The consolidated resources are managed as a single entity called a resource pool. For example, a resource pool might group CPUs of physical servers within a cluster. The capacity of the resource pool is the sum of the power of all CPUs (for example, 10,000 megahertz) available in the cluster. In addition to the CPU pool, the virtual infrastructure includes other types of resource pools, such as memory pool, network pool, and storage pool. Apart from resource pools, the virtual infrastructure also includes identity pools, such as VLAN ID pools and VSAN ID pools. The number of each type of pool and the pool capacity depend on the cloud service provider's requirement to create different cloud services.

Virtual infrastructure also includes virtual computing resources, such as virtual machines, virtual storage volumes, and virtual networks. These resources obtain capacities, such as CPU power, memory, network bandwidth, and storage space from the resource pools. The capacity is allocated to the virtual computing resources easily and flexibly based on the service requirement. Virtual networks are created using network identifiers, such as VLAN IDs and VSAN IDs from the respective identity pools. Virtual computing resources are used for creating cloud infrastructure services.

Applications and Platform Software

- Suite of software that may include:
 - ▶ Business applications
 - ▶ Platform software such as OS and database
 - ▶▶ Provide environments for applications to run
- Applications and platform software are hosted on VMs
 - ▶ To create software-as-a-service (SaaS) and platform-as-a-service (PaaS)

This layer includes a suite of business applications and platform software such as the OS and database. Platform software provides the environment on which business applications to run. Applications and platform software are hosted on virtual machines to create SaaS and PaaS. For SaaS, both the application and platform software are provided by cloud service providers. In the case of PaaS, only the platform software is provided by cloud service providers; consumers export their applications to the cloud.

Cloud Management and Service Creation Tools

- Include three types of software:
 - ▶ Physical and virtual infrastructure management software
 - ▶ Unified management software
 - ▶ User-access management software
- These software interact among themselves to automate provisioning of cloud services

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 27

The cloud management and service creation tools layer includes three types of software:

- Physical and virtual infrastructure management software
- Unified management software
- User-access management software

This classification is based on the different functions performed by these software. These software interact with each other to automate provisioning of cloud services.

The physical and virtual infrastructure management software is offered by the vendors of various infrastructure resources and third-party organizations. For example, a storage array has its own management software. Similarly, network and physical servers are managed independently using network and compute management software respectively. These software provide interfaces to construct a virtual infrastructure from the underlying physical infrastructure.

Unified management software interacts with all standalone physical and virtual infrastructure management software. It collects information on the existing physical and virtual infrastructure configurations, connectivity, and utilization. Unified management software compiles this information and provides a consolidated view of infrastructure resources scattered across one or more data centers. It allows an administrator to monitor performance, capacity, and availability of physical and virtual resources centrally. Unified management software also provides a single management interface to configure physical and virtual infrastructure and integrate the compute (both CPU and memory), network, and storage pools. The integration allows a group of compute pools to use the storage and network pools for storing and transferring data respectively.

Cont..

The unified management software passes configuration commands to respective physical and virtual infrastructure management software, which executes the instructions. This eliminates the administration of compute, storage, and network resources separately using native management software.

The key function of the unified management software is to automate the creation of cloud services. It enables administrators to define service attributes such as CPU power, memory, network bandwidth, storage capacity, name and description of applications and platform software, resource location, and backup policy. When the unified management software receives consumer requests for cloud services, it creates the service based on predefined service attributes.

The user-access management software provides a web-based user interface to consumers. Consumers can use the interface to browse the service catalogue and request cloud services. The user-access management software authenticates users before forwarding their request to the unified management software. It also monitors allocation or usage of resources associated to the cloud service instances. Based on the allocation or usage of resources, it generates a chargeback report. The chargeback report is visible to consumers and provides transparency between consumers and providers.

Cloud-optimized Storage

- Provides rapid elasticity, global access, and storage capacity on-demand
- Leverages object-based storage technology
- Enables self-service and fully metered access to storage resources
- Key characteristics of cloud-optimized storage solution are:
 - ▶ Massively scalable
 - ▶ Unified namespace
 - ▶ Metadata and policy-based information management
 - ▶ Secure multitenancy
 - ▶ Multiple access mechanisms (through REST and SOAP web service APIs and file-based access)

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 29

Content-rich applications combined with the growth of unstructured data are challenging to manage with traditional approach of storing data at scale. This combination of massive growth, new information types, and the need to serve multiple locations and users around the world, has led to requirements for information storage and management at a global scale. Cloud-optimized storage is a solution to meet these requirements. It delivers scalable and flexible architecture that provides rapid elasticity, global access and storage capacity on-demand. It also addresses the constraints of rigid, mount-point based interaction between storage and consumer by presenting a singular access point to the entire storage infrastructure. It leverages a built-in multi-tenancy model and enables self-service, fully metered access to storage resources thereby delivers storage-as-a-service on a shared infrastructure. Cloud-optimized storage typically leverages object-based storage technology that uses customizable, value-driven metadata to drive storage placement, protection and lifecycle policies. Key characteristics of cloud-optimized storage solution are:

- Massively scalable infrastructure that supports large number of objects across a globally distributed infrastructure
- Unified namespace that eliminates capacity, location, and other file system limitations
- Metadata and policy-based information management capabilities that optimizes data protection, availability and cost, based on service levels
- Secure multitenancy that enables multiple applications to be securely served from the same infrastructure. Each application is securely partitioned and data is neither co-mingled nor accessible by other tenants
- Provide access through REST and SOAP web service APIs and file-based access using variety of client devices

Cloud Challenges – Consumer’s Perspective

- Security and regulation
 - Consumers are indecisive to transfer control of sensitive data
 - Regulation may prevent organizations to use cloud services
- Network latency
 - Real time applications may suffer due to network latency and limited bandwidth
- Supportability
 - Service provider might not support proprietary environments
 - Incompatible hypervisors could impact VM migration
- Vendor lock-in
 - Restricts consumers from changing their cloud service providers
 - Lack of standardization across cloud-based platforms

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 30

Business critical data requires protection and continuous monitoring of its access. If the data moves to a cloud model other than an on-premise private cloud, consumers could lose absolute control of their sensitive data. Although most of the cloud service providers offer enhanced data security, consumers might not be willing to transfer control of their business-critical data to the cloud.

Cloud service providers might use multiple data centers located in different countries to provide cloud services. They might replicate or move data across these data centers to ensure high availability and load distribution. Consumers may or may not know in which country their data is stored. Some cloud service providers allow consumers to select the location for storing their data. Data privacy concerns and regulatory compliance requirements, such as the EU Data Protection Directive and the U.S. Safe Harbor program, create challenges for the consumers in adopting cloud computing.

Cloud services can be accessed from anywhere via a network. However, network latency increases when the cloud infrastructure is not close to the access point. A high network latency can either increase the application response time or cause the application to timeout. This can be addressed by implementing stringent Service Level Agreements (SLAs) with the cloud service providers.

Cont..

Another challenge is cloud platform services may not support consumer's desired applications. For example, service provider might not be able to support highly specialized or proprietary environments, such as compatible OSs and preferred programming languages, required to develop and run the consumer's application. Also a mismatch between hypervisors could impact migration of virtual machines into or between clouds.

Another challenge is vendor lock-in: the difficulty for consumers to change their cloud service provider. A lack of interoperability between the APIs of different cloud service providers could also create complexity and high migration costs when moving from one service provider to another.

Cloud Challenges – Provider's Perspective

- Service warranty and service cost
 - ▶ Resources must be kept ready to meet unpredictable demand
 - ▶ Hefty penalty, if SLAs are not fulfilled
- Complexity in deploying vendor software in the cloud
 - ▶ Many vendors do not provide cloud-ready software licenses
 - ▶ Higher cost of cloud-ready software licenses
- No standard cloud access interface
 - ▶ Cloud consumers want open APIs
 - ▶ Need agreement among cloud providers for standardization

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 32

Cloud service providers usually publish a service-level agreement (SLA) so that their consumers know about the availability of service, quality of service, downtime compensation, and legal and regulatory clauses. Alternatively, customer-specific SLAs may be signed between a cloud service provider and a consumer. SLAs typically mention a penalty amount if cloud service providers fail to provide the services levels. Therefore, cloud service providers must ensure that they have adequate resources to provide the required levels of services. Because the cloud resources are distributed and services demand fluctuates, it is a challenge for cloud service providers to provision physical resources for peak demand of all consumers and estimate the actual cost of providing the services.


Many software vendors do not have a cloud-ready software licensing model. Some of the software vendors offer standardized cloud licenses at a higher price compared to traditional licensing models. The cloud software licensing complexity has been causing challenges in deploying vendor software in the cloud. This is also a challenge to the consumer.

Cloud service providers usually offer proprietary APIs to access their cloud. However, consumers might want open APIs or standard APIs to become the tenant of multiple clouds. This is a challenge for cloud service providers because this requires agreement among cloud service providers.

Cloud Adoption Considerations

- CIOs/IT Managers seeking to move to the cloud face several questions:
 - ▶ Which deployment model fits organization's requirements?
 - ▶▶ Private, public, hybrid
 - ▶ Which are the applications suitable for cloud?
 - ▶ How do I choose the cloud service provider?
 - ▶ Is the cloud infrastructure capable of providing the required Quality of Service (QoS)?
 - ▶▶ Performance, availability, and security
 - ▶ What is the financial benefit in adopting cloud?

Organizations that decide to adopt cloud computing always face this question “How does the cloud fit the organization’s environment?” Most organizations are not ready to abandon their existing IT investments to move all of their business processes to the cloud at once. Instead, they need to consider various factors before moving their business processes to the cloud. Even individuals seeking to use cloud services need to understand some cloud adoption considerations.



Risk versus convenience is a key consideration for deciding on a cloud adoption strategy. This consideration also forms the basis for choosing the right cloud deployment model. A public cloud is usually preferred by individuals and start-up businesses. For them, the cost reduction offered by the public cloud outweighs the security or availability risks in the cloud. Small- and medium-sized businesses (SMBs) have a moderate customer base, and any anomaly in customer data and service levels might impact their business. Therefore, they may not be willing to deploy their tier 1 applications, such as Online Transaction Processing (OLTP), in the public cloud. A hybrid cloud model fits in this case. The tier 1 applications should run on the private cloud, whereas less critical applications such as backup, archive, and testing can be deployed in the public cloud. Enterprises typically have a strong customer base worldwide. They usually enforce strict security policies to safeguard critical customer data. Because they are financially capable, they might prefer building their own private clouds.

Choosing Applications for Public Cloud

- Some key questions to ask before migrating a consumer application to the public cloud:
 - ▶ Is the application compatible to cloud platform software? Is it a legacy application?
 - ▶ Is the application proprietary and mission-critical? Does the application provide competitive advantage?
 - ▶ Is the application workload network traffic intensive? Will application performance be impacted by network latency and limited network bandwidth?
 - ▶ Does the application communicate with other data center resources or applications?

Not all applications are good candidates for a public cloud. This may be due to the incompatibility between the cloud platform software and the consumer applications, or maybe the organization plans to move a legacy application to the cloud. Proprietary and mission-critical applications are core and essential to the business. They are usually designed, developed, and maintained in-house. These applications often provide competitive advantages. Due to high security risk, organizations are unlikely to move these applications to the public cloud. These applications are good candidate for an on-premise private cloud. Nonproprietary and nonmission critical applications are suitable for deployment in the public cloud. If an application workload is network traffic-intensive, its performance might not be optimal if deployed in the public cloud. Also if the application communicates with other data center resources or applications, it might experience performance issues.

Financial Advantage

- Require analysis of financial benefits in adopting cloud
- Consider CAPEX and OPEX to deploy and maintain own infrastructure versus cloud-adoption cost

Cost of Owning Infrastructure		Cloud Adoption Cost
CAPEX	OPEX	OPEX
• Servers • Storage • Operating system (OS) • Application • Network equipments • Real estate	• Power and cooling • Personnel • Bandwidth • Maintenance • Support • Backup	• Migration • Compliance and security • Subscription fee

A careful analysis of financial benefits provides a clear picture about the cost-savings in adopting the cloud. The analysis should compare both the Total Cost of Ownership (TCO) and the Return on Investment (ROI) in the cloud and non cloud environment and identify the potential cost benefit. While calculating TCO and ROI, organizations and individuals should consider the expenditure to deploy and maintain their own infrastructure versus cloud-adoption costs. While calculating the expenditures for owning infrastructure resources, organizations should include both the capital expenditure (CAPEX) and operation expenditure (OPEX). The CAPEX includes the cost of servers, storage, OS, application, network equipment, real estate, and so on. The OPEX includes the cost incurred for power and cooling, personnel, maintenance, backup, and so on. These expenditures should be compared with the operation cost incurred in adopting cloud computing. The cloud adoption cost includes the cost of migrating to the cloud, cost to ensure compliance and security, and usage or subscription fees. Moving applications to the cloud reduces CAPEX, except when the cloud is built on-premise.

Selecting a Public Cloud Service Provider

- Some key questions to ask before selecting a provider:
 - ▶ How long and how well has the provider been delivering the services?
 - ▶ How well does the provider meet the organization's current and future requirements?
 - ▶ How easy is it to add or remove services?
 - ▶ How easy is it to move to another provider, when required?
 - ▶ What happens when the provider upgrades their software? Is it forced on everyone? Can you upgrade on your own schedule?
 - ▶ Does the provider offer the required security services?
 - ▶ Does the provider meet your legal and privacy requirements?
 - ▶ Does the provider have good customer service support?

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 4: Cloud Computing 37

The selection of the provider is important for a public cloud. Consumers need to find out how long and how well the provider has been delivering the services. They also need to determine how easy it is to add or terminate cloud services with the service provider. The consumer should know how easy it is to move to another provider, when required. They must assess how the provider fulfills the security, legal, and privacy requirements. They should also check whether the provider offers good customer service support.

QoS Considerations

- Consumers should check whether the QoS attributes meet their requirements
- SLA is a contract between the cloud service provider and consumers that defines QoS attributes
 - ▶ Attributes examples: throughput, uptime, and so on

Cloud service providers typically mention quality of service (QoS) attributes such as throughput and uptime, along with cloud services. The QoS attributes are generally part of an SLA, which is the service contract between the provider and the consumers. The SLA serves as the foundation for the expected level of service between the consumer and the provider. Before adopting the cloud services, consumers should check whether the QoS attributes meet their requirements.

Module 13: Cloud Computing


Concept in Practice:

- Vblock

The Concept in Practice section covers Vblock.

Vblock

- Integrated cloud infrastructure package
 - ▶ Includes compute, storage, network, and virtualization products
 - ▶ Delivered by EMC, VMware, and Cisco
- Enables building virtualized data center and cloud infrastructure
- Pre-architected, preconfigured, pretested, and ready to be deployed
 - ▶ Saves cost and deployment time



Vblock Infrastructure Package

Vblock is completely integrated cloud infrastructure offering that includes compute, storage, network, and virtualization products. These products are provided by EMC, VMware, and Cisco, who have formed a coalition to deliver Vblocks.

Vblocks enables organizations to build virtualized data centers and cloud infrastructures. Vblocks are pre-architected, preconfigured, pretested and have defined performance and availability attributes. Rather than customers buying and assembling individual cloud infrastructure components, Vblock provides a validated cloud infrastructure solution and is factory-ready for deployment and production. This saves significant cost and deployment time.

EMC Unified Infrastructure Manager (UIM) is the unified management solution for Vblocks. UIM provides a single point of management for Vblocks and manages multiple Vblocks. With UIM, cloud infrastructure services can be provisioned automatically and based on provisioning best practices.

Module 13: Summary

Key points covered in this module:

- Characteristics of cloud computing
- Cloud services and deployment models
- Cloud computing infrastructure
- Challenges of cloud computing
- Cloud adoption considerations

This module covered the essential characteristics of cloud computing such as on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service.

It also covered various cloud service models such as Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS).

This module also covered cloud deployment models such as public, private, community, and hybrid.

It also covered cloud computing infrastructure that includes physical infrastructure, virtual infrastructure, applications and platform software, and cloud management and service creation tools.

Finally, it covered the challenges of cloud computing and cloud adoption considerations.

Check Your Knowledge – 1

- Which capability is provided by cloud computing?
 - A. Unilateral provisioning
 - B. Human interaction with service provider
 - C. Increased time-to-market
 - D. Vendor lock-in
- According to NIST, which is an essential cloud characteristic?
 - A. Open API
 - B. Policy-based monitoring
 - C. Measured service
 - D. Software-as-a-service

Check Your Knowledge – 2

- What is offered in infrastructure-as-a-service?
 - A. Database Management System
 - B. Operating system
 - C. Application
 - D. Storage
- Which is a component of virtual infrastructure in cloud?
 - A. Management software
 - B. Storage array
 - C. Network identity pool
 - D. Service catalogue

Check Your Knowledge – 3

- Which is a characteristic of cloud-optimized storage?
 - A. Server-centric
 - B. Location dependent
 - C. Secure multitenancy
 - D. Single access mechanism

Section 4: Summary

Key points covered in this section:

- Cloud Computing

This section covered various aspects of cloud computing.