JEPPIAAR INSTITUTE OF TECHNOLOGY

"Self Belief | Self Discipline | Self Respect"

DEPARTMENT OF

COMPUTER SCIENCE AND ENGINEERING

LECTURE NOTES

CS8791 / CLOUD COMPUTING

(2017 Regulation)

Year/Semester: IV / VII

Prepared by

Dr. K. Tamilarasi, Professor / Dept. of CSE.

UNIT IV RESOURCE MANAGEMENT AND SECURITY IN CLOUD

Inter Cloud Resource Management –Resource Provisioning and Resource Provisioning Methods –Global Exchange of Cloud Resources –Security Overview –Cloud Security Challenges –Software-as-a-Service Security –Security Governance –Virtual Machine Security – IAM –Security Standards.

4.1 Inter Cloud Resource Management

Cloud application (SaaS)
Cloud Software environment (PaaS)
Cloud software infrastructure (IaaS, DaaS, CaaS)
Collocation cloud services (LaaS)
Network cloud services (NaaS)
Hardware or Virtualization cloud Services (HaaS)

Figure 4.1 A stack of six layers of cloud services

- Figure 4.1 shows six layers of cloud services, ranging from hardware, network, and collocation to infrastructure, platform and software applications.
- The cloud platform provides PaaS, which sits on top of the IaaS infrastructure. The top layer offers SaaS.
- The bottom three layers are more related to physical requirements.
- The bottommost layer provides Hardware as a Service (HaaS).

- The next layer is for interconnecting all the hardware components and it is simply called Network as a Service (NaaS).

- Virtual LANs fall within the scope of NaaS.

- The next layer up offers Location as a Service (LaaS), which provides a collocation service to house, power and secure all the physical hardware as well as network resources.

- Some authors say this layer provides Security as a Service (SaaS).

- The cloud infrastructure layer can be further subdivided as Data as a Service (DaaS) and Communication as a Service (CaaS) in addition to compute and storage in IaaS.

- The three cloud models as viewed by different players.

- From the software vendor perspective, application performance on a given cloud platform is most important.

- From the provider perspective, cloud infrastructure performance is the primary concern.

- From the end users perspective, the quality of services, including security, is the most important.

- CRM offered the first SaaS on the cloud successfully.

- The approach is to widen market coverage by investigating customer behaviors and revealing opportunities by statistical analysis.

- SaaS tools also apply to distributed collaboration, financial and human resources management. These cloud services have been growing rapidly in recent years.

- PaaS is provided by Google, Salesforce.com, Facebook, and so on.

- IaaS is provided by Amazon, Windows Azure, RackRack, and so on.

- Based on the observations of some typical cloud computing instances, such as Google, Microsoft, and Yahoo!, the overall software stack structure of cloud computing software can be viewed as layers.

- Each layer has its own purpose and provides the interface for the upper layers just as the traditional so□ware stack does. However, the lower layers are not completely transparent to the upper layers.

- The platform for running cloud computing services can be either physical servers or virtual servers.

- By using VMs, the platform can be flexible; It means the running services are not bound to specific hardware platforms.

- The software layer on top of the platform is the layer for storing massive amounts of data.

- This layer acts like the file system in a traditional single machine. Other layers running on top of the file system are the layers for executing cloud computing applications.

- The next layers are the components in the software stack.

4.1.1 Runtime Support Services

- As in a cluster environment, there are also some runtime supporting services in the cloud computing environment.

- Cluster monitoring is used to collect the runtime status of the entire cluster.

- The scheduler queues the tasks submitted to the whole cluster and assigns the tasks to the processing nodes according to node availability.

- The distributed scheduler for the cloud application has special characteristics that can support cloud applications, such as scheduling the programs written in MapReduce style.

- The runtime support system keeps the cloud cluster working properly with high efficiency.

- Runtime support is software needed in browser initiated applications applied by thousands of cloud customers.

- The SaaS model provides the software applications as a service, rather than lifting users purchase the software.

- As a result, on the customer side, there is no upfront investment in servers or software licensing.

- On the provider side, costs are rather low, compared with conventional hosting of user applications.

- The customer data is stored in the cloud that is either vendor proprietary or a publicly hosted cloud supporting PaaS and IaaS.

4.2 Resource Provisioning

- Providers supply cloud services by signing SLAs with end users.

- The SLAs must commit sufficient resources such as CPU, memory and bandwidth that the user can use for a preset period.

- Under provisioning of resources will lead to broken SLAs and penalties.

- Over provisioning of resources will lead to resource underutilization, and consequently, a decrease in revenue for the provider.

- Deploying an autonomous system to efficiently provision resources to users is a challenging problem.
- Efficient VM provisioning depends on the cloud architecture and management of cloud infrastructures.
- Resource provisioning schemes also demand fast discovery of services and data in cloud computing infrastructures.
- In a virtualized cluster of servers, this demands efficient installation of VMs, live VM migration and fast recovery from failures.
- To deploy VMs, users treat them as physical hosts with customized operating systems for specific applications.
- For example, Amazon's EC2 uses Xen as the virtual machine monitor (VMM). The same VMM is used in IBM's Blue Cloud.
- In the EC2 platform, some predefined VM templates are also provided. Users can choose different kinds of VMs from the templates.
- IBM's Blue Cloud does not provide any VM templates.

4.3 Resource Provisioning Methods

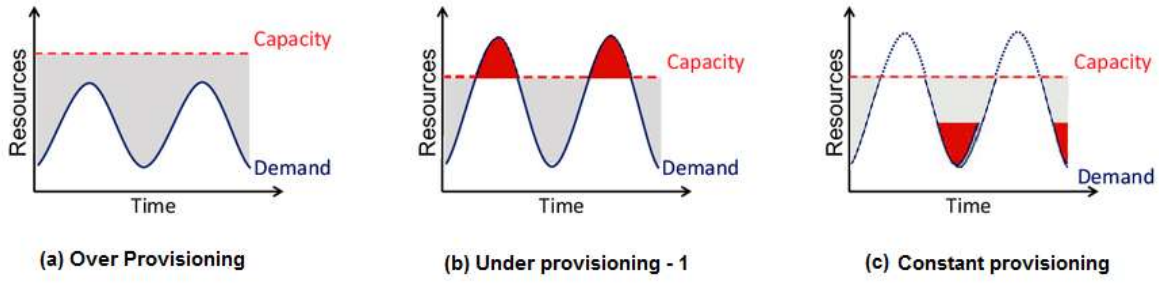

Figure 4.2 Three cases of resource provisioning

- Figure 4.2 shows three cases of static cloud resource provisioning policies.
- In case (a), over provisioning with the peak load causes heavy resource waste (shaded area).
- In case (b), under provisioning (along the capacity line) of resources results in losses by both user and provider in that paid demand by the users (the shaded area above the capacity) is not served and wasted resources still exist for those demanded areas below the provisioned capacity.
- In case (c), the constant provisioning of resources with fixed capacity to a declining user demand could result in even worse resource waste.
- The user may give up the service by canceling the demand, resulting in reduced revenue for the provider.
- Both the user and provider may be losers in resource provisioning without elasticity.
- The demand-driven method provides static resources and has been used in grid computing for many years.
- The event-driven method is based on predicted workload by time.
- The popularity-driven method is based on Internet traffic monitored.

4.3.1 Demand-Driven Resource Provisioning

- This method adds or removes computing instances based on the current utilization level of the allocated resources.
- The demand-driven method automatically allocates two Xeon processors for the user application, when the user was using one Xeon processor more than 60 percent of the time for an extended period

- In general, when a resource has surpassed a threshold for a certain amount of time, the scheme increases that resource based on demand.

- When a resource is below a threshold for a certain amount of time, that resource could be decreased accordingly.

- Amazon implements such an auto-scale feature in its EC2 platform. This method is easy to implement.

- The scheme does not work out right if the workload changes abruptly.

4.3.2 Event-Driven Resource Provisioning

- This scheme adds or removes machine instances based on a specific time event.

- The scheme works better for seasonal or predicted events such as Christmastime in the West and the Lunar New Year in the East.

- During these events, the number of users grows before the event period and then decreases during the event period.

- This scheme anticipates peak traffic before it happens.

- The method results in a minimal loss of QoS, if the event is predicted correctly.

- Otherwise, wasted resources are even greater due to events that do not follow a fixed pattern.

4.3.3 Popularity-Driven Resource Provisioning

- In this method, the Internet searches for popularity of certain applications and creates the instances by popularity demand.

- The scheme anticipates increased traffic with popularity.
- Again, the scheme has a minimal loss of QoS, if the predicted popularity is correct.
- Resources may be wasted if traffic does not occur as expected.

4.4 Global Exchange of Cloud Resources

- In order to support a large number of application service consumers from around the world, cloud infrastructure providers (i.e., IaaS providers) have established data centers in multiple geographical locations to provide redundancy and ensure reliability in case of site failures.
- For example, Amazon has data centers in the United States (e.g., one on the East Coast and another on the West Coast) and Europe.
- However, currently Amazon expects its cloud customers (i.e., SaaS providers) to express a preference regarding where they want their application services to be hosted.
- Amazon does not provide seamless/automatic mechanisms for scaling its hosted services across multiple geographically distributed data centers.
- This approach has many shortcomings.
    - First, it is difficult for cloud customers to determine in advance the best location for hosting their services as they may not know the origin of consumers of their services.
    - Second, SaaS providers may not be able to meet the QoS expectations of their service consumers originating from multiple geographical locations.
- Figure 4.3 shows the high-level components of the Melbourne group's proposed Inter Cloud architecture.

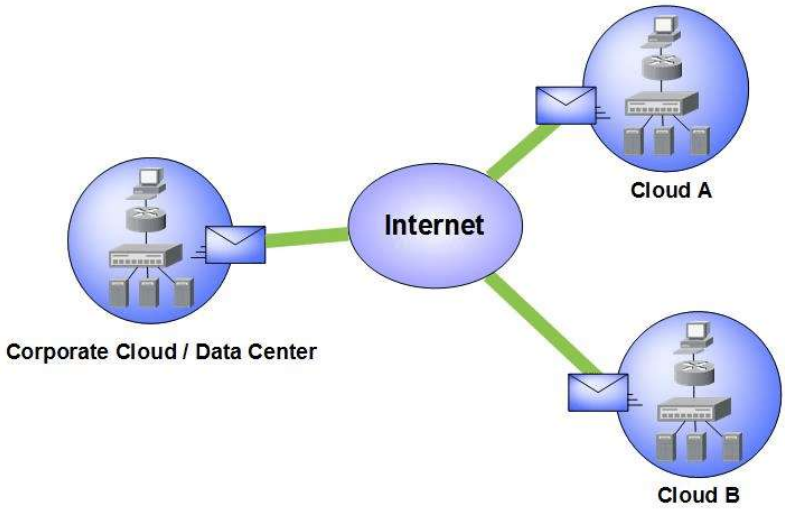

Figure 4.3 Inter cloud architecture

- In addition, no single cloud infrastructure provider will be able to establish its data centers at all possible locations throughout the world.

- As a result, cloud application service (SaaS) providers will have difficulty in meeting QoS expectations for all their consumers.

- The Cloudbus Project at the University of Melbourne has proposed InterCloud architecture supporting brokering and exchange of cloud resources for scaling applications across multiple clouds.

- By realizing InterCloud architectural principles in mechanisms in their offering,

    - Cloud providers will be able to dynamically expand or resize their provisioning capability based on sudden spikes in workload demands by leasing available computational and storage capabilities from other cloud service providers.
    - Operate as part of a market driven resource leasing federation, where application service providers such as Salesforce.com host their services based on negotiated SLA contracts driven by competitive market prices.

  - Deliver on-demand, reliable, cost-effective, and QoS-aware services based on virtualization technologies while ensuring high QoS standards and minimizing service costs.

- They need to be able to utilize market-based utility models as the basis for provisioning of virtualized software services and federated hardware infrastructure among users with heterogeneous applications.

- They consist of client brokering and coordinator services that support utility-driven federation of clouds:
  - Application scheduling
  - Resource allocation
  - Migration of workloads

- The architecture cohesively couples the administratively and topologically distributed storage and compute capabilities of clouds as part of a single resource leasing abstraction.

- The Cloud Exchange (CEx) acts as a market maker for bringing together service producers and consumers.

- It aggregates the infrastructure demands from application brokers and evaluates them against the available supply currently published by the cloud coordinators.

- It supports trading of cloud services based on competitive economic models such as commodity markets and auctions.

- An SLA specifies the details of the service to be provided in terms of metrics agreed upon by all parties, and incentives and penalties for meeting and violating the expectations, respectively.

- The availability of a banking system within the market ensures that financial transactions pertaining to SLAs between participants are carried out in a secure and dependable environment.

4.5 Security Overview

- Cloud service providers must learn from the managed service provider (MSP) model and ensure that their customer's applications and data are secure if they hope to retain their customer base and competitiveness.

- Today, enterprises are looking toward cloud computing horizons to expand their on-premises infrastructure, but most cannot afford the risk of compromising the security of their applications and data.

- For example, IDC recently conducted a survey1 (Figure 4.4) of 244 IT executives/CIOs and their line-of-business (LOB) colleagues to gauge their opinions and understand their companies' use of IT cloud services.

- Security ranked first as the greatest challenge or issue of cloud computing.

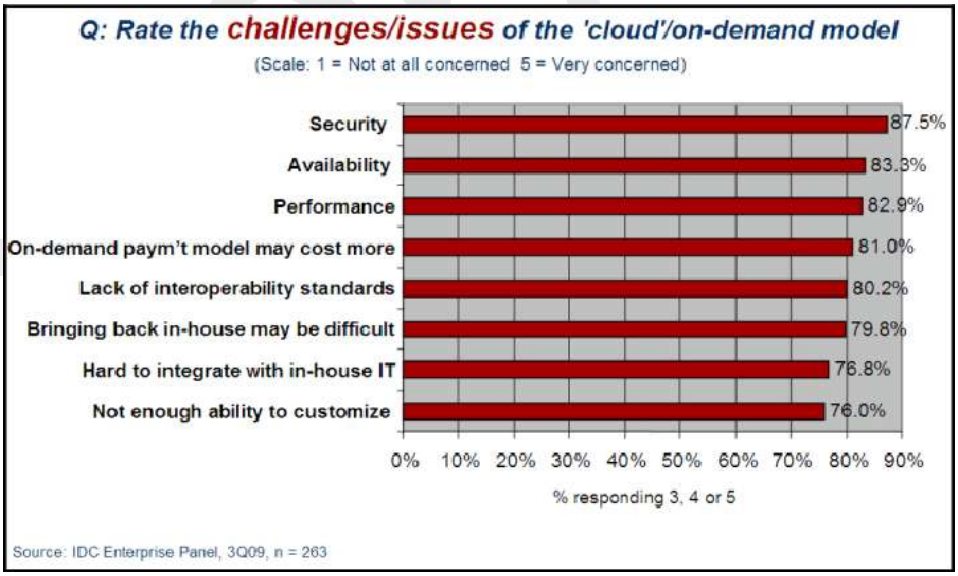

Figure 4.4 Results of IDC survey

- Moving critical applications and sensitive data to public and shared cloud environments is of great concern for those corporations that are moving beyond their data center's network perimeter defense.

- To alleviate these concerns, a cloud solution provider must ensure that customers will continue to have the same security and privacy controls over their applications and services.

- In addition, solution provider give evidence to customers that their organization and customers are secure and they can meet their service level agreements, and that they can prove compliance to auditors.

4.6 Cloud Security Challenges

- Although virtualization and cloud computing can help companies accomplish more by breaking the physical bonds between an IT infrastructure and its users, heightened security threats must be overcome in order to benefit fully from this new computing paradigm.

- Enterprise security is only as good as the least reliable partner, department and vendor.

- With the cloud model, the cloud consumer's loss control over physical security.

- In a public cloud, the consumers are sharing computing resources with other companies.

- In a shared pool outside the enterprise, users do not have any knowledge or control of where the resources run.

- Storage services provided by one cloud vendor may be incompatible with another vendor's services should you decide to move from one to the other.

- Ensuring the integrity of the data really means that it changes only in response to authorized transactions.

- The immature use of mash up technology (combinations of web services), which is fundamental to cloud applications, is inevitably going to cause unwitting security vulnerabilities in those applications.

- Since access to logs is required for Payment Card Industry Data Security Standard (PCI DSS) compliance and may be requested by auditors and regulators, security managers need to make sure to negotiate access to the provider's logs as part of any service agreement.

- Cloud applications undergo constant feature additions and users must keep up to date with application improvements to be sure they are protected.

- The speed at which applications will change in the cloud will affect both the SDLC and security.

- Security needs to move to the data level, so that enterprises can be sure their data is protected wherever it goes.

- Sensitive data is the domain of the enterprise, not the cloud computing provider.

- One of the key challenges in cloud computing is data level security.

- Most compliance standards do not envision compliance in a world of cloud computing.

- There is a huge body of standards that apply for IT security and compliance, governing most business interactions that will, over time, have to be translated to the cloud.

- SaaS makes the process of compliance more complicated, since it may be difficult for a customer to discern where its data resides on a network controlled by its SaaS provider, or a partner of that provider, which raises all sorts of compliance issues of data privacy, segregation, and security.

- Security managers will need to pay particular attention to systems that contain critical data such as corporate financial information or source code during the transition to server virtualization in production environments.

- Outsourcing means losing significant control over data, and while this is not a good idea from a security perspective, the business ease and financial savings will continue to increase the usage of these services.

- Security managers will need to work with their company's legal staff to ensure that appropriate contract terms are in place to protect corporate data and provide for acceptable service level agreements.

- Cloud based services will result in many mobile IT users accessing business data and services without traversing the corporate network.

- This will increase the need for enterprises to place security controls between mobile users and cloud based services.

- Although traditional data center security still applies in the cloud environment, physical segregation and hardware based security cannot protect against attacks between virtual machines on the same server.

- Administrative access is through the Internet rather than the controlled and restricted direct or on-premises connection that is adhered to in the traditional data center model.

- This increases risk and exposure and will require stringent monitoring for changes in system control and access control restriction.

- Proving the security state of a system and identifying the location of an insecure virtual machine will be challenging.

- The co-location of multiple virtual machines increases the attack surface and risk of virtual machine to virtual machine compromise.

- Localized virtual machines and physical servers use the same operating systems as well as enterprise and web applications in a cloud server environment, increasing the threat of an attacker or malware exploiting vulnerabilities in these systems and applications remotely.

- Virtual machines are vulnerable as they move between the private cloud and the public cloud.

- A fully or partially shared cloud environment is expected to have a greater attack surface and therefore can be considered to be at greater risk than a dedicated resources environment.

- Operating system and application files are on a shared physical infrastructure in a virtualized cloud environment and require system, file, and activity monitoring to provide confidence and auditable proof to enterprise customers that their resources have not been compromised or tampered with.

- In the cloud computing environment, the enterprise subscribes to cloud computing resources, and the responsibility for patching is the subscriber's rather than the cloud computing vendors.

- The need for patch maintenance vigilance is imperative.

- Data is fluid in cloud computing and may reside in on-premises physical servers, on-premises virtual machines, or off-premises virtual machines running on cloud computing resources and this will require some rethinking on the part of auditors and practitioners alike.

- To establish zones of trust in the cloud, the virtual machines must be self-defending, effectively moving the perimeter to the virtual machine itself.

- Enterprise perimeter security (i.e., firewalls, demilitarized zones [DMZs], network segmentation, intrusion detection and prevention systems [IDS/IPS], monitoring tools, and the associated security policies) only controls the data that resides and transits behind the perimeter.

- In the cloud computing world, the cloud computing provider is in charge of customer data security and privacy.

4.7 Software-as-a-Service Security

- Cloud computing models of the future will likely combine the use of SaaS (and other XaaS's as appropriate), utility computing and Web 2.0 collaboration technologies to leverage the Internet to satisfy their customer needs.

- New business models being developed as a result of the move to cloud computing are creating not only new technologies and business operational processes but also new security requirements and challenges as described previously.

- As the most recent evolutionary step in the cloud service model (Figure 4.5), SaaS will likely remain the dominant cloud service model for the predictable future and the area where the most critical need for security practices and oversight will reside.

- The technology analyst and consulting firm Gartner lists seven security issues which one should discuss with a cloud computing vendor.

- Privileged user access inquires about who has specialized access to data and about the hiring and management of such administrators.

- Regulatory compliance makes sure that the vendor is willing to undergo external audits and/or security certifications.

- Data location does the provider allow for any control over the location of data.

- Data segregation makes encryption is available at all stages and that these encryption schemes were designed and tested by experienced professionals.

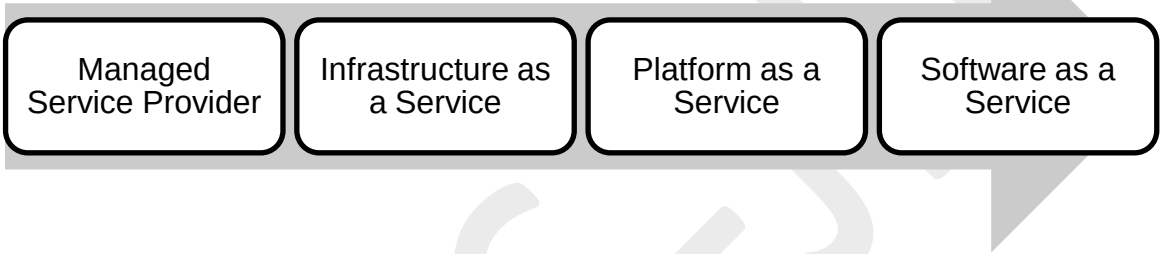

Figure 4.5 Evolution of cloud services

- Recovery is the way to find out what will happen to data in the case of a disaster. And also it covers the way to perform complete restoration.

- Investigative support does the vendor have the ability to investigate any inappropriate or illegal activity.

- Long-term viability focus on data if the company goes out of business and format and process behind the returned data.

- To address the security issues listed above, SaaS providers will need to incorporate and enhance security practices used by the managed service providers and develop new ones as the cloud computing environment evolves.

4.8 Security Governance

- A security steering committee should be developed whose objective is to focus on providing guidance about security initiatives and alignment with business and IT strategies.

- A charter for the security team is typically one of the first deliverables from the steering committee.

- This charter must clearly define the roles and responsibilities of the security team and other groups involved in performing information security functions.

- Lack of a formalized strategy can lead to an unsustainable operating model and security level as it evolves.

- In addition, lack of attention to security governance can result in key needs of the business not being met, including but not limited to, risk management, security monitoring, application security, and sales support.

- Lack of proper governance and management of duties can also result in potential security risks being left unaddressed and opportunities to improve the business being missed because the security team is not focused on the key security functions and activities that are critical to the business.

4.9 Virtual Machine Security

- In the cloud environment, physical servers are consolidated to multiple virtual machine instances on virtualized servers.

- Not only can data center security teams replicate typical security controls for the data center at large to secure the virtual machines, they can also advise their customers on how to prepare these machines for migration to a cloud environment when appropriate.

- Firewalls, intrusion detection and prevention, integrity monitoring and log inspection can all be deployed as software on virtual machines to increase protection as well as

maintain compliance integrity of servers and applications as virtual resources move from on-premises to public cloud environments.

- By deploying this traditional line of defense to the virtual machine itself, the user can enable critical applications and data to be moved to the cloud securely.

- To facilitate the centralized management of a server firewall policy, the security software loaded onto a virtual machine should include a bidirectional stateful firewall that enables virtual machine isolation and location awareness, thereby enabling a tightened policy and the flexibility to move the virtual machine from on-premises to cloud resources.

- Integrity monitoring and log inspection software must be applied at the virtual machine level.

- This approach to virtual machine security, which connects the machine back to the mother ship, has some advantages in that the security software can be put into a single software agent that provides for consistent control and management throughout the cloud while integrating seamlessly back into existing security infrastructure investments, providing economies of scale, deployment, and cost savings for both the service provider and the enterprise.

4.10 IAM

- Identity and access management is a critical function for every organization and a fundamental expectation of SaaS customers is that the principle of least privilege is granted to their data.

- The principle of least privilege states that only the minimum access necessary to perform an operation should be granted, and that access should be granted only for the minimum amount of time necessary.

- However, business and IT groups will need and expect access to systems and applications.

- The advent of cloud services and services on demand is changing the identity management landscape.

- Most of the current identity management solutions are focused on the enterprise and typically are architected to work in a very controlled, static environment.

- User-centric identity management solutions such as federated identity management make some assumptions about the parties involved and their related services.

- In the cloud environment, where services are offered on demand and they can continuously evolve, aspects of current models such as trust assumptions, privacy implications, and operational aspects of authentication and authorization, will be challenged.

- Meeting these challenges will require a balancing act for SaaS providers as they evaluate new models and management processes for IAM to provide end-to-end trust and identity throughout the cloud and the enterprise.

- Another issue will be finding the right balance between usability and security. If a good balance is not achieved, both business and IT groups may be affected by barriers to completing their support and maintenance activities efficiently.

4.11 Security Standards

- Security standards define the processes, procedures, and practices necessary for implementing a security program.

- These standards also apply to cloud related IT activities and include specific steps that should be taken to ensure a secure environment is maintained that provides privacy and security of confidential information in a cloud environment.

- Security standards are based on a set of key principles intended to protect this type of trusted environment.

- Messaging standards, especially for security in the cloud, must also include nearly all the same considerations as any other IT security endeavor.

- Security (SAML OAuth, OpenID, SSL/TLS) A basic philosophy of security is to have layers of defense, a concept known as defense in depth.

- This means having overlapping systems designed to provide security even if one system fails. An example is a firewall working in conjunction with an intrusion-detection system (IDS).

- Defense in depth provides security because there is no single point of failure and no single entry vector at which an attack can occur.

- For this reason, a choice between implementing network security in the middle part of a network (i.e., in the cloud) or at the endpoints is a false dichotomy.

- No single security system is a solution by itself, so it is far better to secure all systems.

- This type of layered security is precisely what we are seeing develop in cloud computing.

- Traditionally, security was implemented at the endpoints, where the user controlled access.

- An organization had no choice except to put firewalls, IDSs, and antivirus software inside its own network.

- Today, with the advent of managed security services offered by cloud providers, additional security can be provided inside the cloud.

4.11.1 Security Assertion Markup Language (SAML)

- SAML is an XML-based standard for communicating authentication, authorization and attribute information among online partners.
- It allows businesses to securely send assertions between partner organizations regarding the identity and entitlements of a principal.
- The Organization for the Advancement of Structured Information Standards (OASIS) Security Services Technical Committee is in charge of defining, enhancing and maintaining the SAML specifications.
- SAML is built on a number of existing standards, namely, SOAP, HTTP and XML. SAML relies on HTTP as its communications protocol and specifies the use of SOAP (currently, version 1.1).
- Most SAML transactions are expressed in a standardized form of XML.
- SAML assertions and protocols are specified using XML schema.
- Both SAML 1.1 and SAML 2.0 use digital signatures (based on the XML Signature standard) for authentication and message integrity.
- XML encryption is supported in SAML 2.0, though SAML 1.1 does not have encryption capabilities.
- SAML defines XML based assertions and protocols, bindings and profiles.

- The term SAML Core refers to the general syntax and semantics of SAML assertions as well as the protocol used to request and transmit those assertions from one system entity to another.

- SAML protocol refers to what is transmitted, not how it is transmitted.

- A SAML binding determines how SAML requests and responses map to standard messaging protocols. An important (synchronous) binding is the SAML SOAP binding.

- SAML standardizes queries for, and responses that contain, user authentication, entitlements and attribute information in an XML format.

- This format can then be used to request security information about a principal from a SAML authority.

- A SAML authority, sometimes called the asserting party. It is a platform or application that can relay security information.

- The relying party (or assertion consumer or requesting party) is a partner site that receives the security information.

- The exchanged information deals with a subject's authentication status, access authorization, and attribute information.

- A subject is an entity in a particular domain.

- A person identified by an email address is a subject, as might be a printer.

- SAML assertions are usually transferred from identity providers to service providers.

- Assertions contain statements that service providers use to make access control decisions.

- Three types of statements are provided by SAML:

    - Authentication statements
    - Attribute statements
    - Authorization decision statements

- SAML assertions contain a packet of security information in this form:

```
<saml: Asssertion A>
<Authentication>
…
</Authentication>
<Attribute>
…
</Attribute>
<Authentication>
…
</Authentication>
</saml: Asssertion A>
```

- The assertion shown above is interpreted as follows:
  Assertion A, issued at time T by issuer I, regarding subject S, provided conditions C are valid.

- Authentication statements assert to a service provider that the principal did indeed authenticate with an identity provider at a particular time using a particular method of authentication.

- Other information about the authenticated principal (called the authentication context) may be disclosed in an authentication statement.

- An attribute statement asserts that a subject is associated with certain attributes.
- An attribute is simply a name-value pair.
- An authorization decision statement asserts that a subject is permitted to perform action A on resource R given evidence E.
- A SAML protocol describes how certain SAML elements (including assertions) are packaged within SAML request and response elements
- Generally, a SAML protocol is a simple request–response protocol.
- The most important type of SAML protocol request is a query.
- A service provider makes a query directly to an identity provider over a secure back channel. For this reason, query messages are typically bound to SOAP.
- Corresponding to the three types of statements, there are three types of SAML queries:
    - Authentication query
    - Attribute query
    - Authorization decision query.
- Of these, the attribute query is perhaps most important. The result of an attribute query is a SAML response containing an assertion, which itself contains an attribute statement.

4.11.2 Open Authentication (OAuth)

- OAuth is an open protocol, initiated by Blaine Cook and Chris Messina, to allow secure API authorization in a simple, standardized method for various types of web applications.

- Cook and Messina had concluded that there were no open standards for API access delegation.

- The OAuth discussion group was created in April 2007, for the small group of implementers to write the draft proposal for an open protocol.

- DeWitt Clinton of Google learned of the OAuth project and expressed interest in supporting the effort.

- In July 2007, the team drafted an initial specification and it was released in October of the same year.

- OAuth is a method for publishing and interacting with protected data.

- For developers, OAuth provides users access to their data while protecting account credentials.

- OAuth allows users to grant access to their information, which is shared by the service provider and consumers without sharing all of their identity.

- The Core designation is used to stress that this is the baseline, and other extensions and protocols can build on it.

- By design, OAuth Core 1.0 does not provide many desired features (e.g., automated discovery of endpoints, language support, support for XML-RPC and SOAP, standard definition of resource access, OpenID integration, signing algorithms, etc.).

- This intentional lack of feature support is viewed by the authors as a significant benefit.

- The Core deals with fundamental aspects of the protocol, namely, to establish a mechanism for exchanging a user name and password for a token with defined rights and to provide tools to protect the token.

- It is important to understand that security and privacy are not guaranteed by the protocol.

- In fact, OAuth by itself provides no privacy at all and depends on other protocols such as SSL to accomplish that.

- OAuth can be implemented in a secure manner.

- In fact, the specification includes substantial security considerations that must be taken into account when working with sensitive data.

- With Oauth, sites use tokens coupled with shared secrets to access resources.

- Secrets, just like passwords, must be protected.

4.11.3 OpenID

- OpenID is an open, decentralized standard for user authentication and access control that allows users to log onto many services using the same digital identity.

- It is a single-sign-on (SSO) method of access control. As such, it replaces the common log-in process (i.e., a log-in name and a password) by allowing users to log in once and gain access to resources across participating systems.

- The original OpenID authentication protocol was developed in May 2005 by Brad Fitzpatrick, creator of the popular community web site LiveJournal.

- In late June 2005, discussions began between OpenID developers and other developers from an enterprise software company named NetMesh.

- These discussions led to further collaboration on interoperability between OpenID and NetMesh's similar Light-Weight Identity (LID) protocol.

- The direct result of the collaboration was the Yadis discovery protocol, which was announced on October 24, 2005.

- The Yadis specification provides a general-purpose identifier for a person and any other entity, which can be used with a variety of services.

- It provides syntax for a resource description document identifying services available using that identifier and an interpretation of the elements of that document.

- Yadis discovery protocol is used for obtaining a resource description document, given that identifier.

- Together these enable coexistence and interoperability of a rich variety of services using a single identifier.

- The identifier uses a standard syntax and a well established namespace and requires no additional namespace administration infrastructure.

- An OpenID is in the form of a unique URL and is authenticated by the entity hosting the OpenID URL.

- The OpenID protocol does not rely on a central authority to authenticate a user's identity.

- Neither the OpenID protocol nor any web sites requiring identification can mandate that a specific type of authentication be used; nonstandard forms of authentication such as smart cards, biometrics, or ordinary passwords are allowed.

- A typical scenario for using OpenID might be something like this:
    - A user visits a web site that displays an OpenID log in form
    - Unlike a typical log in form, which has fields for user name and password, the OpenID log in form has only one field for the OpenID identifier (which is an OpenID URL).
    - This form is connected to an implementation of an OpenID client library.
    - A user will have previously registered an OpenID identifier with an OpenID identity provider.
    - The user types this OpenID identifier into the OpenID log-in form.
    - The relying party then requests the web page located at that URL and reads an HTML link tag to discover the identity provider service URL.

- With OpenID 2.0, the client discovers the identity provider service URL by requesting the XRDS document (also called the Yadis document) with the content type application/xrds+xml, which may be available at the target URL but is always available for a target XRI.

- There are two modes by which the relying party can communicate with the identity provider: checkid_immediate and checkid_setup.

- In checkid_immediate, the relying party requests that the provider not interact with the user. All communication is relayed through the user's browser without explicitly notifying the user.

- In checkid_setup, the user communicates with the provider server directly using the same web browser as is used to access the relying party site.

- OpenID does not provide its own authentication methods, but if an identity provider uses strong authentication, OpenID can be used for secure transactions.

- SSL/TLS Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), are cryptographically secure protocols designed to provide security and data integrity for communications over TCP/IP.

- TLS and SSL encrypt the segments of network connections at the transport layer.

- Several versions of the protocols are in general use in web browsers, email, instant messaging and Voice-over-IP (VoIP).

- TLS is an IETF standard protocol which was last updated in RFC 5246.

- The TLS protocol allows client/server applications to communicate across a network in a way specifically designed to prevent eavesdropping, tampering, and message forgery.

- TLS provides endpoint authentication and data confidentiality by using cryptography.

- TLS authentication is one way in which the server is authenticated, because the client already knows the server's identity. In this case, the client remains unauthenticated.

- TLS also supports a more secure bilateral connection mode whereby both ends of the connection can be assured that they are communicating with whom they believe they are connected.

- This is known as mutual authentication.

- Mutual authentication requires the TLS client side to also maintain a certificate.

- TLS involves three basic phases:

    - Peer negotiation for algorithm support
    - Key exchange and authentication
    - Symmetric cipher encryption and message authentication

TWO MARK QUESTIONS

1. List the runtime supporting services in the cloud computing environment.

- Cluster monitoring is used to collect the runtime status of the entire cluster.
- The scheduler queues the tasks submitted to the whole cluster and assigns the tasks to the processing nodes according to node availability.
- The distributed scheduler for the cloud application has special characteristics that can support cloud applications, such as scheduling the programs written in MapReduce style.

2. Why inter cloud resource management requires runtime support system?

- The runtime support system keeps the cloud cluster working properly with high efficiency.
- Runtime support is software needed in browser-initiated applications applied by thousands of cloud customers.

3. Differentiate between over provisioning and under provisioning.

- Under provisioning of resources will lead to broken SLAs and penalties.
- Over provisioning of resources will lead to resource underutilization, and consequently, a decrease in revenue for the provider.
- over provisioning with the peak load causes heavy resource waste (shaded area).
- under provisioning (along the capacity line) of resources results in losses by both user and provider in that paid demand by the users (the shaded area above the capacity) is not served and wasted resources still exist for those demanded areas below the provisioned capacity.

4. List the various resource provisioning methods.

- demand-driven resource provisioning
- Event-Driven Resource Provisioning

- Popularity-Driven Resource Provisioning
- Global Exchange of Cloud Resources

5. What is demand-driven resource provisioning?

- This method adds or removes computing instances based on the current utilization level of the allocated resources.
- The demand-driven method automatically allocates two Xeon processors for the user application, when the user was using one Xeon processor more than 60 percent of the time for an extended period

6. Write short notes on cloud security.

- Cloud service providers must learn from the managed service provider (MSP) model and ensure that their customer's applications and data are secure if they hope to retain their customer base and competitiveness.
- Security ranked first as the greatest challenge or issue of cloud computing.

7. List the challenges in cloud security.

- Enterprise security is only as good as the least reliable partner, department, or vendor.
- With the cloud model, users lose control over physical security.
- In a public cloud, the users are sharing computing resources with other companies.
- In a shared pool outside the enterprise, users don't have any knowledge or control of where the resources run.
- Storage services provided by one cloud vendor may be incompatible with another vendor's services should you decide to move from one to the other.
- Ensuring the integrity of the data really means that it changes only in response to authorized transactions.

8. List the seven security issues with respect to cloud computing vendor.

- Privileged user access
- Regulatory compliance
- Data location
- Data segregation
- Recovery
- Investigative support
- Long-term viability

9. What is the purpose of security governance?

- A security steering committee should be developed whose objective is to focus on providing guidance about security initiatives and alignment with business and IT strategies.
- A charter for the security team is typically one of the first deliverables from the steering committee.

10. How to perform virtual machine security?

- Firewalls, intrusion detection and prevention, integrity monitoring, and log inspection can all be deployed as software on virtual machines to increase protection and maintain compliance integrity of servers and applications as virtual resources move from on-premises to public cloud environments.
- Integrity monitoring and log inspection software must be applied at the virtual machine level.

11. Define IAM.

- Identity and access management is a critical function for every organization, and a fundamental expectation of SaaS customers is that the principle of least privilege is granted to their data.

12. Why cloud requires security standards?

- Security standards define the processes, procedures, and practices necessary for implementing a security program.
- These standards also apply to cloud related IT activities and include specific steps that should be taken to ensure a secure environment is maintained that provides privacy and security of confidential information in a cloud environment.

13. What is SAML?

- Security Assertion Markup Language (SAML) is an XML-based standard for communicating authentication, authorization, and attribute information among online partners.
- It allows businesses to securely send assertions between partner organizations regarding the identity and entitlements of a principal.

14. List the types of statements are provided by SAML.

- Authentication statements
- Attribute statements
- Authorization decision statements

15. Describe about SAML protocol.

- A SAML protocol describes how certain SAML elements (including assertions) are packaged within SAML request and response elements
- SAML protocol is a simple request–response protocol.
- The most important type of SAML protocol request is a query.

16. List the types of SAML queries.

- Authentication query
- Attribute query
- Authorization decision query.

17. What is OAuth?

- OAuth (Open authentication) is an open protocol, initiated by Blaine Cook and Chris Messina, to allow secure API authorization in a simple, standardized method for various types of web applications.
- OAuth is a method for publishing and interacting with protected data.
- OAuth allows users to grant access to their information, which is shared by the service provider and consumers without sharing all of their identity.

18. What is the purpose of OpenID?

- OpenID is an open, decentralized standard for user authentication and access control that allows users to log onto many services using the same digital identity.
- It is a single-sign-on (SSO) method of access control.
- An OpenID is in the form of a unique URL and is authenticated by the entity hosting the OpenID URL.

19. Why cloud environment need SSL/TLS?

- SSL/TLS Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), are cryptographically secure protocols designed to provide security and data integrity for communications over TCP/IP.
- TLS and SSL encrypt the segments of network connections at the transport layer.

20. What is mutual authentication?

- TLS also supports a more secure bilateral connection mode whereby both ends of the connection can be assured that they are communicating with whom they believe they are connected. This is known as mutual authentication.
- Mutual authentication requires the TLS client side to also maintain a certificate.