



UNIT 5 Security Standards and Applications

What are Cloud Security Standards?

It was essential to establish guidelines for how work is done in the cloud due to the different security dangers facing the cloud. They offer a thorough framework for how cloud security is upheld with regard to both the user and the service provider.

- Cloud security standards provide a roadmap for businesses transitioning from a traditional approach to a cloud-based approach by providing the right tools, configurations, and policies required for security in cloud usage.
- It helps to devise an effective security strategy for the organization.
- It also supports organizational goals like privacy, portability, security, and interoperability.
- Certification with cloud security standards increases trust and gives businesses a competitive edge.

Need for Cloud Security Standards

- Ensure cloud computing is an appropriate environment: Organizations need to make sure that cloud computing is the appropriate environment for the applications as security and mitigating risk are the major concerns.
- To ensure that sensitive data is safe in the cloud: Organizations need a way to make sure that the sensitive data is safe in the cloud while remaining compliant with standards and regulations.
- No existing clear standard: Cloud security standards are essential as earlier there were no existing clear standards that can define what constitutes a secure cloud environment. Thus, making it difficult for cloud providers and cloud users to define what needs to be done to ensure a secure environment.
- Need for a framework that addresses all aspects of cloud security: There is a need for businesses to adopt a

Lack of Cloud Security Standards

- Enterprises and CSPs have been forced to fumble while relying on an endless variety of auditing needs, regulatory requirements, industry mandates, and data Centre standards to offer direction on protecting their cloud environments due to the lack of adequate cloud security standards.
- Because of this, the Cloud Security Alliance is more difficult to understand than it first appears, and its fragmented strategy does not meet the criteria for “excellent security”.

Best Practices For Cloud Security

1. Secure Access to the Cloud

Although the majority of cloud service providers have their own ways of safeguarding the infrastructure of their clients, you are still in charge of protecting the cloud user accounts and access to sensitive data for your company. Consider improving password management in your organization to lower the risk of account compromise and credential theft.



Adding password policies to your cybersecurity program is a good place to start. Describe the cybersecurity practices you demand from your staff, such as using unique, complex passwords for each account and routine password rotation.

2. Control User Access Rights

Some businesses give employees immediate access to a wide range of systems and data in order to make sure they can carry out their tasks effectively. For cybercriminals, these individuals' accounts are a veritable gold mine because compromising them can make it simpler to gain access to crucial cloud infrastructure and elevate privileges. Your company can periodically review and revoke user rights to prevent this.

3. Transparency and Employee Monitoring

You can use specialized solutions to keep an eye on the behavior of your staff in order to promote transparency in your cloud infrastructure. You can spot the earliest indications of a cloud account compromise or an insider threat by keeping an eye on what your employees are doing while they are at work. Imagine your cybersecurity experts discover a user accessing your cloud infrastructure from a strange IP address or outside of normal business hours. In that situation, they'll be able to respond to such odd activity promptly because it suggests that a breach may be imminent.

4. Data Protection

This involves data protection against unauthorized access, prevention of accidental data disclosure, and ensuring ceaseless access to crucial data in the case of failures and errors.

5. Access Management

Three capabilities that are a must in access management are the ability to identify and authenticate users, the ability to assign access rights to users, and the ability to develop and enact access control policies for all the resources.

Common Cloud Security Standards

1. NIST (National Institute of Standards and Technology)

NIST is a federal organization in the US that creates metrics and standards to boost competition in the scientific and technology industries. The National Institute of Regulations and Technology (NIST) developed the Cybersecurity Framework to comply with US regulations such as the Federal Information Security Management Act and the Health Insurance Portability and Accountability Act (HIPAA) (FISMA). NIST places a strong emphasis on classifying assets according to their commercial value and adequately protecting them.



2. ISO-27017

A development of ISO-27001 that includes provisions unique to cloud-based information security. Along with ISO-27001 compliance, ISO-27017 compliance should be taken into account. This standard has not yet been introduced to the marketplace. It attempts to offer further direction in the cloud computing information security field. Its purpose is to supplement the advice provided in ISO/IEC 27002 and various other ISO27k standards, such as ISO/IEC 27018 on the privacy implications of cloud computing, and ISO/IEC 27031 on business continuity.

3. ISO-27018

The protection of personally identifiable information (PII) in public clouds that serve as PII processors is covered by this standard. Despite the fact that this standard is especially aimed at public-cloud service providers like AWS or Azure, PII controllers (such as a SaaS provider processing client PII in AWS) nevertheless bear some accountability. If you are a SaaS provider handling PII, you should think about complying with this standard.

4. CIS controls

Organizations can secure their systems with the help of Internet Security Center (CIS) Controls, which are open-source policies based on consensus. Each check is rigorously reviewed by a number of professionals before a conclusion is reached.

To easily access a list of evaluations for cloud security, consult the CIS Benchmarks customized for particular cloud service providers. For instance, you can use the CIS-AWS controls, a set of controls created especially for workloads using Amazon Web Services (AWS).

5. FISMA

In accordance with the Federal Information Security Management Act (FISMA), all federal agencies and their contractors are required to safeguard information systems and assets. NIST, using NIST SP 800-53, was given authority under FISMA to define the framework security standards (see definition below).

6. Cloud Architecture Framework

These frameworks, which frequently cover operational effectiveness, security, and cost-value factors, can be viewed as best parties standards for cloud architects. This framework, developed by Amazon Web Services, aids architects in designing workloads and applications on the Amazon cloud. Customers have access to a reliable resource for architecture evaluation thanks to this framework, which is based on a collection of questions for the analysis of cloud environments.

7. General Data Protection Regulation (GDPR)



For the European Union, there are laws governing data protection and privacy. Even though this law only applies to the European Union, it is something you should keep in mind if you store or otherwise handle any personal information of residents of the EU.

8. SOC Reporting

A form of audit of the operational processes used by IT businesses offering any service is known as a “Service and Organization Audits 2” (SOC 2). A worldwide standard for cybersecurity risk management systems is SOC 2 reporting. Your company’s policies, practices, and controls are in place to meet the five trust principles, as shown by the SOC 2 Audit Report. The SOC 2 audit report lists security, availability, processing integrity, confidentiality, and confidentiality as security principles. If you offer software as a service, potential clients might request proof that you adhere to SOC 2 standards.

9. PCI DSS

For all merchants who use credit or debit cards, the PCI DSS (Payment Card Industry Data Security Standard) provides a set of security criteria. For businesses that handle cardholder data, there is PCI DSS. The PCI DSS specifies fundamental technological and operational criteria for safeguarding cardholder data. Cardholders are intended to be protected from identity theft and credit card fraud by the PCI DSS standard.

10. HIPAA

The Health Insurance Portability and Accountability Act (HIPAA), passed by the US Congress to safeguard individual health information, also has parts specifically dealing with information security. Businesses that handle medical data must abide by HIPAA law. The HIPAA Security Rule (HSR) is the best choice in terms of information security. The HIPAA HSR specifies rules for protecting people’s electronic personal health information that a covered entity generates, acquires, makes use of or maintains.

Organizations subject to HIPAA regulations need risk evaluations and risk management plans to reduce threats to the availability, confidentiality, and integrity of the crucial health data they manage. Assume your company sends and receives health data via cloud-based services (SaaS, IaaS, PaaS). If so, it is your responsibility to make sure the service provider complies with HIPAA regulations and that you have implemented best practices for managing your cloud setups.

11. CIS AWS Foundations v1.2

Any business that uses Amazon Web Service cloud resources can help safeguard sensitive IT systems and data by adhering to the CIS AWS Foundations Benchmark. Intelligence analysts developed a set of objective, consensus-driven configuration standards known as the CIS (Center for Internet Security) Benchmarks to help businesses improve their information security. Additionally, CIS procedures are for fortifying AWS accounts to build a solid foundation for running jobs on AWS.



12. ACSC Essential Eight

ACSC Essential 8 (also known as the ASD Top 4) is a list of eight cybersecurity mitigation strategies for small and large firms. In order to improve security controls, protect businesses' computer resources and systems, and protect data from cybersecurity attacks, the Australian Signals Directorate (ASD) and the Australian Cyber Security Centre (ACSC) developed the "Essential Eight Tactics."

Security Issues in Cloud Computing :

There is no doubt that Cloud Computing provides various Advantages but there are also some security issues in cloud computing. Below are some following Security Issues in Cloud Computing as follows.

1. **Data Loss –**

Data Loss is one of the issues faced in Cloud Computing. This is also known as Data Leakage. As we know that our sensitive data is in the hands of Somebody else, and we don't have full control over our database. So, if the security of cloud service is to break by hackers then it may be possible that hackers will get access to our sensitive data or personal files.

2. **Interference of Hackers and Insecure API's –**

As we know, if we are talking about the cloud and its services it means we are talking about the Internet. Also, we know that the easiest way to communicate with Cloud is using API. So it is important to protect the Interface's and API's which are used by an external user. But also in cloud computing, few services are available in the public domain which are the vulnerable part of Cloud Computing because it may be possible that these services are accessed by some third parties. So, it may be possible that with the help of these services hackers can easily hack or harm our data.

3. **User Account Hijacking –**

Account Hijacking is the most serious security issue in Cloud Computing. If somehow the Account of User or an Organization is hijacked by a hacker then the hacker has full authority to perform Unauthorized Activities.

4. **Changing Service Provider –**

Vendor lock-In is also an important Security issue in Cloud Computing. Many organizations will face different problems while shifting from one vendor to another. For example, An Organization wants to shift from AWS cloud to Google Cloud Services then they face various problems like shifting of all data, also both cloud services have different techniques and functions, so they also face problems regarding that. Also, it may be possible that the charges of AWS are different from Google Cloud,



etc.

5. **Lack of Skill –**

While working, shifting to another service provider, need an extra feature, how to use a feature, etc. are the main problems caused in IT Company who doesn't have skilled Employees. So it requires a skilled person to work with Cloud Computing.

6. **Denial of Service (DoS) attack –**

This type of attack occurs when the system receives too much traffic. Mostly DoS attacks occur in large organizations such as the banking sector, government sector, etc. When a DoS attack occurs, data is lost. So, in order to recover data, it requires a great amount of money as well as time to handle it.

Security in Clouds

- **Cloud Security**, also known as cloud computing security, consists of a set of policies, controls, procedures and technologies that work together to protect cloud-based systems, data, and infrastructure.
- These security measures are configured to protect cloud data, support regulatory compliance and protect customers' privacy as well as setting authentication rules for individual users and devices.
- From authenticating access to filtering traffic, cloud security can be configured to the exact needs of the business. And because these rules can be configured and managed in one place, administration overheads are reduced and IT teams empowered to focus on other areas of the business.
- The way cloud security is delivered will depend on the individual cloud provider or the cloud security solutions in place. However, implementation of cloud security processes should be a joint responsibility between the business owner and solution provider.
- For businesses making the transition to the cloud, robust cloud security is imperative. Security threats are constantly evolving and becoming more sophisticated, and cloud computing is no less at risk than an on-premise environment. For this reason, it is essential to work with a cloud provider that offers best-in-class security that has been customized for your infrastructure.



Benefits of Cloud Security

1. Centralized security: Just as cloud computing centralizes applications and data, cloud security centralizes protection. Cloud-based business networks consist of numerous devices and endpoints that can be difficult to manage when dealing with **shadow IT** or **BYOD**. Managing these entities centrally enhances traffic analysis and **web filtering**, streamlines the monitoring of network events and results in fewer software and policy updates. Disaster recovery plans can also be implemented and actioned easily when they are managed in one place.

2. Reduced costs: One of the benefits of utilizing cloud storage and security is that it eliminates the need to invest in dedicated hardware. Not only does this reduce capital expenditure, but it also reduces administrative overheads. Where once IT teams were firefighting security issues reactively, cloud security delivers proactive security features that offer protection 24/7 with little or no human intervention.

3. Reduced Administration: When you choose a reputable cloud services provider or cloud security platform, you can kiss goodbye to manual security configurations and almost constant security updates. These tasks can have a massive drain on resources, but when you move them to the cloud, all security administration happens in one place and is fully managed on your behalf.

4. Reliability: Cloud computing services offer the ultimate in dependability. With the right cloud security measures in place, users can safely access data and applications within the cloud no matter where they are or what device they are using.



Software as a Service Security

- SaaS security is cloud-based security designed to protect the data that software as service applications carry.
- It's a set of practices that companies that store data in the cloud put in place to protect sensitive information pertaining to their customers and the business itself.
- However, SaaS security is not the sole responsibility of the organization using the cloud service. In fact, the service customer and the service provider share the obligation to adhere to SaaS security guidelines published by the National Cyber Security Center (NCSC).
- SaaS security is also an important part of SaaS management that aims to reduce unused licenses, shadow IT and decrease security risks by creating as much visibility as possible.

6 SaaS Security best practices

One of the main benefits that SaaS has to offer is that the respective applications are on-demand, scalable, and very fast to implement, saving companies valuable resources and time. On top of that, the SaaS provider typically handles updates and takes care of software maintenance.

This flexibility and the fairly open access have created new security risks that SaaS security best practices are trying to address and mitigate. Below are 6 security practices and solutions that every cloud-operating business should know about.

1. Enhanced Authentication

Offering a cloud-based service to your customers means that there has to be a way for them to access the software. Usually, this access is regulated through login credentials. That's why knowing how your users access the resource and how the third-party software provider handles the authentication process is a great starting point.

Once you understand the various methods, you can make better SaaS security decisions and enable additional security features like multifactor authentication or integrate other enhanced authentication methods.

2. Data Encryption

The majority of channels that SaaS applications use to communicate employ TLS (Transport Layer Security) to protect data that is in transit. However, data that is at rest can be just as vulnerable to cyber attacks as data that is being exchanged. That's why more and more SaaS providers offer encryption capabilities that protect data in transit and at rest. It's a good idea to talk to your provider and check whether enhanced data encryption is available for all the SaaS services you use.

3. Vetting and Oversight

With a stark increase in SaaS deployment, usage and demand, new SaaS vendors emerge on a regular basis. This creates a competitive market and gives companies seeking the best SaaS solutions for their business needs the upper hand. However, too many similar products can lead to decision fatigue or rash decisions. When you choose your saas provider, apply the same review and validation process you would with other vendors and compare optional security features that might be available.



4. Discovery and Inventory

With increased digital literacy, software procurement is not only limited to IT departments but can be practiced by almost every employee. Ultimately, this leads to shadow IT and security loopholes. That's why one of the most important SaaS security practices involves maintaining a reliable inventory of what services are being used and the tracking of SaaS usage to detect unusual or unexpected activity. Automated tools within SaaS management systems can send out alerts for immediate notification.

5. Consider CASBs

It is possible that the SaaS provider that you are choosing is not able to provide the level of SaaS security that your company requires. If there are no viable alternatives when it comes to the vendor, consider cloud access security broker (CASB) tool options. This allows your company to add a layer of additional security controls that are not native to your SaaS application. When selecting a CASB —whether proxy or API-based —make sure it fits into your existing IT architecture.

6. Maintain situational awareness

Last but not least, always monitor your SaaS use. Comprehensive SaaS management tools and CASBs offer you a lot of information that can help you make the right decision when it comes to SaaS security.

Common Cloud Security Standard

Cloud Security encompasses the technologies, controls, processes, and policies which combine to protect your cloud-based systems, data, and infrastructure. It is a sub-domain of computer security and more broadly, information security.

The most well-known standard in information security and compliance is ISO 27001, developed by the International Organization for Standardization.

The ISO 27001 standard was created to assist enterprises in protecting sensitive data by best practices.

Cloud compliance is the **principle that cloud-delivered systems must be compliant with the standards their customers require**. Cloud compliance ensures that cloud computing services meet compliance requirements.

| | Infrastructure-as-a-service (IaaS) | Platform-as-a-service (PaaS) | Software-as-a-service (SaaS) |
|-------------------------|------------------------------------|------------------------------|------------------------------|
| People | You | You | You |
| Data | You | You | You |
| Applications | You | You | CSP |
| Operating system | You | CSP | CSP |
| Virtual networks | You | CSP | CSP |
| Hypervisors | CSP | CSP | CSP |
| Servers and storage | CSP | CSP | CSP |
| Physical networks | CSP | CSP | CSP |

· Naveen Singh
jh@gmail.com



Open Cloud Consortium

- The Open Cloud Consortium (OCC) is
 - A not for profit
 - Manages and operates cloud computing infrastructure to support scientific, medical, health care and environmental research.
- OCC members span the globe and include over 10 universities, over 15 companies, and over 5 government agencies and national laboratories.
- The OCC is organized into several different working groups.

The OCC Mission

- The purpose of the Open Cloud Consortium is to support the development of standards for cloud computing and to develop a framework for interoperability among various clouds.
 - The OCC supports the development of benchmarks for cloud computing.
 - Manages cloud computing testbeds, such as the Open Cloud Testbed, to improve cloud computing software and services.
 - Develops reference implementations, benchmarks and standards, such as the MalStone Benchmark, to improve the state of the art of cloud computing.
 - Sponsors workshops and other events related to cloud computing to educate the community.
-



The Open Cloud Consortium

- The **Open Commons Consortium** (*aka OCC* - formerly the **Open Cloud Consortium**) is a 501(c)(3) non-profit venture which provides cloud computing and data commons resources to support "scientific, environmental, medical and health care research."
- OCC manages and operates resources including the Open Science Data Cloud (*aka OSDC*), which is a multi-petabyte scientific data sharing resource.
- The consortium is based in [Chicago, Illinois](#), and is managed by the 501(c)3 [Center for Computational Science](#)
- **The OCC is divided into Working Groups which include:**
 - **The Open Science Data Cloud** - This is a working group that manages and operates the Open Science Data Cloud (OSDC), which is a petabyte scale science cloud for researchers to manage, analyze and share their data. Individual researchers may apply for accounts to analyze data hosted by the OSDC. Research projects with TB-scale datasets are encouraged to join the OSDC and contribute towards its infrastructure.

2. Project Matsu - Project Matsu is a collaboration between the NASA Goddard Space Flight Center and the Open Commons Consortium to develop open source technology for cloud-based processing of satellite imagery to support the earth science research community as well as human assisted disaster relief.

3. The Open Cloud Testbed - This working group manages and operates the Open Cloud Testbed. The Open Cloud Testbed (OCT) is a geographically distributed cloud testbed spanning four data centers and connected with 10G and 100G network connections. The OCT is used to develop new cloud computing software and infrastructure.

4. The Biomedical Data Commons - The Biomedical Data Commons (BDC) is cloud-based infrastructure that provides secure, compliant cloud services for managing and analyzing genomic data, electronic medical records (EMR), medical images, and other PHI data. It provides resources to researchers so that they can more easily make discoveries from large complex controlled access datasets. The BDC provides resources to those institutions in the BDC Working Group. It is an example of what is sometimes called condominium model of sharing research infrastructure in which the research infrastructure is operated by a consortium of educational and research organizations and provides resources to the consortium.

5. NOAA Data Alliance Working Group - The OCC National Oceanographic and Atmospheric Administration (NOAA) Data Alliance Working Group supports and manages the NOAA data commons and the surrounding community interested in the open redistribution of NOAA datasets.

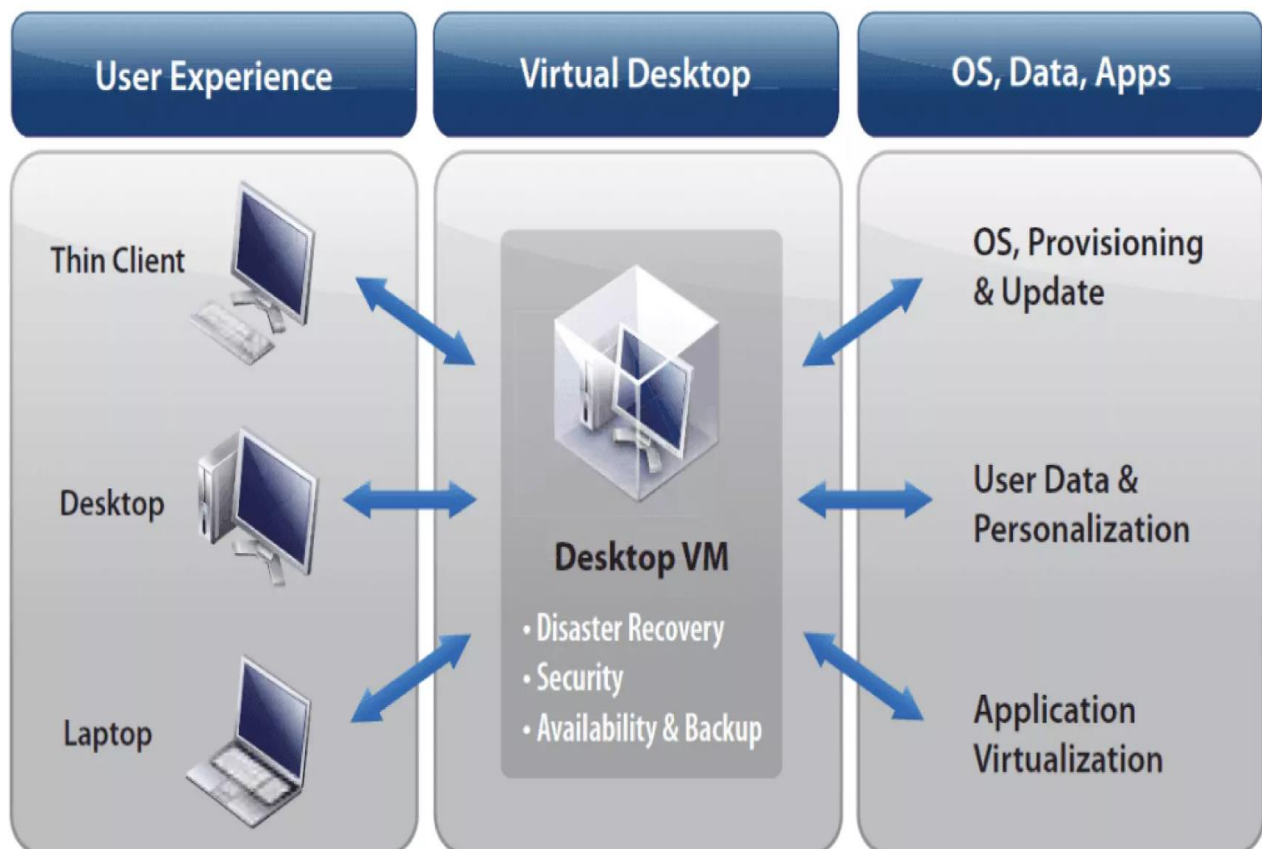
In 2015, the OCC was accepted into the Matter healthcare community at Chicago's historic Merchandise Mart. Matter is a community healthcare entrepreneurs and industry leaders working together in a shared space to individually and collectively fuel the future of healthcare innovation.

In 2015, the OCC announced a collaboration with the National Oceanic and Atmospheric Administration (NOAA) to help release their vast stores of environmental data to the general public. This effort is managed by the OCC's NOAA data alliance working group.



The Distributed management Task Force (DMTF)

- DMTF is a 501(c)(6) nonprofit industry standards organization that creates open manageability standards spanning diverse emerging and traditional IT infrastructures including cloud, virtualization, network, servers and storage. Member companies and alliance partners collaborate on standards to improve interoperable management of information technologies.
- Based in Portland, Oregon, the DMTF is led by a board of directors representing technology companies including: Broadcom Inc., Cisco, Dell Technologies, Hewlett Packard Enterprise, Intel Corporation, Lenovo, NetApp, Positive Tecnologia S.A., and Verizon.
- Founded in 1992 as the Desktop Management Task Force, the organization's first standard was the now-legacy Desktop Management Interface (DMI). As the organization evolved to address distributed management through additional standards, such as the Common Information Model (CIM), it changed its name to the Distributed Management Task Force in 1999, but is now known as, DMTF.
- The DMTF continues to address converged, hybrid IT and the Software Defined Data Center (SDDC) with its latest specifications, such as the CADF (Cloud Auditing Data Federation), CIMI (Cloud Infrastructure Management Interface), CIM (Common Information Model), DASH (Desktop and Mobile Architecture for System Hardware), MCTP (Management Component Transport Protocol), NC-SI (Network Controller Sideband Interface), OVF (Open Virtualization Format), PLDM (Platform Level Data Model), Redfish Device Enablement (RDE), Redfish (Including Protocols, Schema, Host Interface, Profiles) SMASH (Systems Management Architecture for Server Hardware) and SMBIOS (System Management BIOS).





The Distributed Management Task Force (DMTF)

- DMTF enables more effective management of millions of IT systems worldwide by bringing the IT industry together to collaborate on the development, validation and promotion of systems management standards.
- The group spans the industry with 160 member companies and organizations, and more than 4,000 active participants crossing 43 countries.
- The DMTF board of directors is led by 16 innovative, industry-leading technology companies.

The Distributed Management Task Force (DMTF)

- DMTF management standards are critical to enabling management interoperability among multi vendor systems, tools and solutions within the enterprise.
- The DMTF started the Virtualization Management Initiative (VMAN).
- The Open Virtualization Format (OVF) is a fairly new standard that has emerged within the VMAN Initiative.
- Benefits of VMAN are Lowering the IT learning curve, and Lowering complexity for vendors implementing their solutions

Standardized Approaches available to Companies due to VMAN Initiative

- Deploy virtual computer systems
- Discover and take inventory of virtual computer systems
- Manage the life cycle of virtual computer systems
- Add/change/delete virtual resources
- Monitor virtual systems for health and performance



Standards for Application Developers

- The purpose of application development standards is to ensure uniform, consistent, high-quality software solutions.
- Programming standards help to improve the readability of the software, allowing developers to understand new code more quickly and thoroughly.
- Commonly used application standards are available for the Internet in browsers, for transferring data, sending messages, and securing data.

Standards for Browsers (Ajax)

- Using Ajax, a web application can request only the content that needs to be updated in the web pages. This greatly reduces networking bandwidth usage and page load times.
- Sections of pages can be reloaded individually.
- An Ajax framework helps developers to build dynamic web pages on the client side. Data is sent to or from the server using requests, usually written in JavaScript.
- ICEfaces is an open source Ajax framework developed as Java product and maintained by <http://icefaces.org>.



ICEfaces Ajax Application Framework

- ICEfaces is an integrated Ajax application framework that enables Java EE application developers to easily create and deploy thin-client rich Internet applications in pure Java.
- To run ICEfaces applications, users need to download and install the following products:
 - Java 2 Platform, Standard Edition
 - Ant
 - Tomcat
 - ICEfaces
 - Web browser (if you don't already have one installed)

Security Features in ICEfaces Ajax Application Framework

- ICEfaces is the one of the most secure Ajax solutions available.
- It is Compatible with SSL (Secure Sockets Layer) protocol.
- It prevents cross-site scripting, malicious code injection, and unauthorized data mining.
- ICEfaces does not expose application logic or user data.
- It is effective in preventing fake form submits and SQL (Structured Query Language) injection attacks.

Data (XML, JSON)

- Extensible Markup Language (XML) allows to define markup elements.
- Its purpose is to enable sharing of structured data.
- XML is often used to describe structured data and to serialize Objects.
- XML provides a basic syntax that can be used to share information among different kinds of computers, different applications, and different organizations without needing to be converted from one to another.



Solution Stacks (LAMP and LAPP)

- LAMP is a popular open source solution commonly used to run dynamic web sites and servers.
- The acronym derives from the fact that it includes Linux, Apache, MySQL, and PHP (or Perl or Python) and is considered by many to be the platform of choice for development and deployment of high-performance web applications which require a solid and reliable foundation.
- When used in combination, they represent a solution stack of technologies that support application servers.

Linux, Apache, PostgreSQL, and PHP(or Perl or Python) (LAPP)

- The LAPP stack is an open source web platform that can be used to run dynamic web sites and servers. It is considered by many to be a more powerful alternative to the more popular LAMP stack.
- LAPP offers SSL
- Many consider the LAPP stack a more secure out-of-the-box solution than the LAMP stack.



Standards for Messaging

- A message is a unit of information that is moved from one place to another.
- Most common messaging standards used in the cloud are
 - Simple Message Transfer Protocol (SMTP)
 - Post Office Protocol (POP)
 - Internet Messaging Access Protocol (IMAP)
 - Syndication (Atom, Atom Publishing Protocol, and RSS)
 - Communications (HTTP, SIMPLE, and XMPP)

Simple Message Transfer Protocol

- Simple Message Transfer Protocol is arguably the most important protocol in use today for basic messaging. Before SMTP was created, email messages were sent using File Transfer Protocol (FTP).
- The FTP protocol was designed to transmit files, not messages, so it did not provide any means for recipients to identify the sender or for the sender to designate an intended recipient.
- SMTP was designed so that sender and recipient information could be transmitted with the message.
- SMTP is a two-way protocol that usually operates using TCP (Transmission Control Protocol) port 25.

Post Office Protocol (POP)

- SMTP can be used both to send and receive messages, but the client must have a constant connection to the host to receive SMTP messages.
- The Post Office Protocol (POP) was introduced to circumvent this situation.
- POP is a lightweight protocol whose single purpose is to download messages from a server. This allows a server to store messages until a client connects and requests them.
- Once the client connects, POP servers begin to download the messages and subsequently delete them from the server (a default setting) in order to make room for more messages.



Internet Messaging Access Protocol

- Once mail messages are downloaded with POP, they are automatically deleted from the server when the download process has finished.
- Many businesses have compulsory compliance guidelines that require saving messages. It also becomes a problem if users move from computer to computer or use mobile networking, since their messages do not automatically move where they go.
- To get around these problems, a standard called Internet Messaging Access Protocol was created. IMAP allows messages to be kept on the server but viewed and manipulated (usually via a browser) as though they were stored locally.

Standards for Security

- Security standards define the processes, procedures, and practices necessary for implementing a secure environment that provides privacy and security of confidential information in a cloud environment.
- Security protocols, used in the cloud are:
 - Security Assertion Markup Language (SAML)
 - Open Authentication (Oauth)
 - OpenID
 - SSL/TLS



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.
- SAML allows a user to log on once for affiliated but separate Web sites. SAML is designed for business-to-business (B2B) and business-to-consumer (B2C) transactions.
- 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.
- Most SAML transactions are expressed in a standardized form of XML. SAML assertions and protocols are specified using XML schema.

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.
- OAuth is a method for publishing and interacting with protected data.
- OAuth provides users access to their data while protecting account credentials.
- OAuth by itself *provides no privacy at all* and depends on other protocols such as SSL to accomplish that.

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




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.
- TLS provides endpoint authentication and data confidentiality by using cryptography.
- TLS involves three basic phases:
 - Peer negotiation for algorithm support
 - Key exchange and authentication
 - Symmetric cipher encryption and message authentication

End user Access to Cloud Computing

- In its most strict sense, end-user computing (EUC) refers **to computer systems and platforms that help non-programmers create applications**. ... What's important is that a well-designed EUC/VDI plan can allow users to access the digital platforms they need to be productive, both on-premises and working remotely in the cloud.
- An End-User Computing application or EUC is **any application that is not managed and developed in an environment that employs robust IT general controls**. ... Although the most pervasive EUCs are spreadsheets, EUCs also can include user databases, queries, scripts, or output from various reporting tools.
- Broadly, end-user computing covers a wide range of user-facing resources, such as: **desktop and notebook end user computers**; desktop operating systems and applications; wearables and smartphones; cloud, mobile, and web applications; and virtual desktops and applications.

WHAT IS END-USER COMPUTING?



EUC = computer systems & platforms meant to allow non-programmers to create working computer applications

3 Types of EUC

- 1 Traditional EUC** = the end user merely uses computer systems & applications created by developers
- 2 End-User Control** = the user's department purchases package applications & hardware for their use
- 3 End-User Development** = the user is given a set of tools that let them customize & create applications



Mobile Internet devices and the Cloud

- Mobile cloud computing uses cloud computing to deliver applications to **mobile devices**. These mobile apps can be deployed remotely using speed and flexibility and development tools.
- Mobile cloud storage is a form of cloud storage that is accessible on mobile devices such as laptops, tablets, and smartphones. Mobile cloud storage providers offer services that **allow the user to create and organize files, folders, music, and photos**, similar to other cloud computing models.
- The mobile cloud is **Internet-based data, applications and related services accessed through smartphones, laptop computers, tablets and other portable devices**. Mobile cloud computing is differentiated from mobile computing in general because the devices run cloud-based Web apps rather than native apps.
- **Locator apps and remote backup** are two types of cloud-enabled services for mobile devices
- A mobile cloud app is a software program designed to be accessible via the internet through portable devices. In terms of the real world, there are many examples of mobile cloud solutions, including: **Email**.

Hadoop (https://en.wikipedia.org/wiki/Apache_Hadoop)

- It is a collection of **open-source** software utilities that facilitates using a network of many computers to solve problems involving massive amounts of data and computation.
- It provides a **software framework** for **distributed storage** and processing of **big data** using the **MapReduce programming model**.
- Hadoop was originally designed for **computer clusters** built from **commodity hardware**, which is still the common use. It has since also found use on clusters of higher-end hardware.
- All the modules in Hadoop are designed with a fundamental assumption that hardware failures are common occurrences and should be automatically handled by the framework.
- The core of Apache Hadoop consists of a storage part, known as Hadoop Distributed File System (HDFS), and a processing part which is a MapReduce programming model.
- Hadoop splits files into large blocks and distributes them across nodes in a cluster. It then transfers **packaged code** into nodes to process the data in parallel. This approach takes advantage of **data locality**, where nodes manipulate the data they have access to.
- This allows the dataset to be **processed** faster and more efficiently than it would be in a more conventional **supercomputer architecture** that relies on a **parallel file system** where computation and data are distributed via high-speed networking.



The base Apache Hadoop framework is composed of the following modules:

- *Hadoop Common* – contains libraries and utilities needed by other Hadoop modules;
- *Hadoop Distributed File System (HDFS)* – a distributed file-system that stores data on commodity machines, providing very high aggregate bandwidth across the cluster;
- *Hadoop YARN* – (introduced in 2012) a platform responsible for managing computing resources in clusters and using them for scheduling users' applications;^{[10][11]}
- *Hadoop MapReduce* – an implementation of the MapReduce programming model for large-scale data processing.
- *Hadoop Ozone* – (introduced in 2020) An object store for Hadoop
- The term *Hadoop* is often used for both base modules and sub-modules and also the *ecosystem*, or collection of additional software packages that can be installed on top of or alongside Hadoop, such as [Apache Pig](#), [Apache Hive](#), [Apache HBase](#), [Apache Phoenix](#), [Apache Spark](#), [Apache ZooKeeper](#), [Cloudera Impala](#), [Apache Flume](#), [Apache Sqoop](#), [Apache Oozie](#), and [Apache Storm](#).
- Apache Hadoop's MapReduce and HDFS components were inspired by [Google](#) papers on [MapReduce](#) and [Google File System](#).
- The Hadoop framework itself is mostly written in the [Java programming language](#), with some native code in [C](#) and [command line](#) utilities written as [shell scripts](#). Though MapReduce Java code is common, any programming language can be used with Hadoop Streaming to implement the map and reduce parts of the user's program.^[14] Other projects in the Hadoop ecosystem expose richer user in

MapReduce

- MapReduce is a **programming model or pattern within the Hadoop framework that is used to access big data stored in the Hadoop File System (HDFS)**. ... MapReduce facilitates concurrent processing by splitting petabytes of data into smaller chunks, and processing them in parallel on Hadoop commodity servers.
- MapReduce is a programming model for processing large amounts of data in a parallel and distributed fashion. It is useful for large, long-running jobs that cannot be handled within the scope of a single request, tasks like:
 - Analyzing application logs
 - Aggregating related data from external sources
 - Transforming data from one format to another
 - Exporting data for external analysis
 - App Engine MapReduce is a community-maintained, open source library that is built on top of App Engine services, including Datastore and Task Queues. The library is available on GitHub at these locations:
 - [Java](#) source project.
 - [Python](#) source project.
- **MapReduce is a software framework** for easily writing applications which process vast amounts of data (multi-terabyte data-sets) in-parallel on large clusters (thousands of nodes) of commodity hardware in a reliable, fault-tolerant manner.
- A MapReduce *job* usually splits the input data-set into independent chunks which are processed by the *map tasks* in a completely parallel manner.
- The framework sorts the outputs of the maps, which are then input to the *reduce tasks*.
- Typically both the input and the output of the job are stored in a file-system.
- The framework takes care of scheduling tasks, monitoring them and re-executes the failed tasks.



- Typically the compute nodes and the storage nodes are the same, that is, the MapReduce framework and the Hadoop Distributed File System are running on the same set of nodes. This configuration allows the framework to effectively schedule tasks on the nodes where data is already present, resulting in very high aggregate bandwidth across the cluster.
- The MapReduce framework consists of a single master `JobTracker` and one slave `TaskTracker` per cluster-node. The master is responsible for scheduling the jobs' component tasks on the slaves, monitoring them and re-executing the failed tasks. The slaves execute the tasks as directed by the master.
- Minimally, applications specify the input/output locations and supply *map* and *reduce* functions via implementations of appropriate interfaces and/or abstract-classes. These, and other job parameters, comprise the *job configuration*.
- The Hadoop *job client* then submits the job (jar/executable etc.) and configuration to the `JobTracker` which then assumes the responsibility of distributing the software/configuration to the slaves, scheduling tasks and monitoring them, providing status and diagnostic information to the job-client.

VirtualBox

- **VirtualBox** is a general-purpose **Type-2 Hypervisor virtualization tool** for x86 and x86-64 hardware developed by Oracle Corp., targeted at server, desktop, and embedded use, that allows users and administrators to easily run multiple guest operating systems on a single host.
- VirtualBox was originally created by **Innotek GmbH**, which was acquired by Sun Microsystems in 2008, which was in turn acquired by Oracle in 2010.
- VirtualBox may be installed on Microsoft Windows, MacOS, Linux, Solaris and OpenSolaris. There are also ports to FreeBSD and Genode.
- It supports the creation and management of guest virtual machines running Windows, Linux, BSD, OS/2, Solaris, Haiku, and OSx86, as well as limited virtualization of macOS guests on Apple hardware. For some guest operating systems, a "Guest Additions" package of device drivers and system applications is available, which typically improves performance, especially that of graphics, and allows changing the resolution of the guest OS automatically when the window of the virtual machine on the host OS is resized.

Google App Engine

- Google App Engine (often referred to as GAE or simply App Engine) is a cloud computing platform as a **service for developing and hosting web applications in** Google-managed data centers. Applications are sandboxed and run across multiple servers.
- Google App Engine, which is a **platform-as-a-service (PaaS)** offering that gives software developers access to Google's scalable hosting.
- **Major Features of Google App Engine in Cloud Computing**

| | | |
|---|----------------------|-----------------------|
| Collection of Development Languages & Tools | Fully Managed | Pay-as-you-Go |
| Effective Diagnostic Services | Traffic Splitting | All Time Availability |
| Ensure Faster Time to Market | Easy to Use Platform | |
- An App Engine web application can be described as having three major parts:

| | | |
|------------------------------|------------------------------|--------------------------|
| Application instances | Scalable data storage | Scalable services |
|------------------------------|------------------------------|--------------------------|



Programming Environment for Google App Engine

- **Google App Engine** (often referred to as **GAE** or simply **App Engine**) is a cloud computing platform as a service for developing and hosting web applications in Google-managed data centers.
 - Applications are sandboxed and run across multiple servers. App Engine offers automatic scaling for web applications—as the number of requests increases for an application, App Engine automatically allocates more resources for the web application to handle the additional demand.
 - Google App Engine primarily supports Go, PHP, Java, Python, Node.js, .NET, and Ruby applications, although it can also support other languages via "custom runtimes". The service is free up to a certain level of consumed resources and only in standard environment but not in flexible environment. Fees are charged for additional storage, [bandwidth](#), or instance hours required by the application. It was first released as a preview version in April 2008 and came out of preview in September 2011.
 - The environment you choose depends on the language and related technologies you want to use for developing the application.
-

Runtimes and framework

- Google App Engine primarily supports [Go](#), [PHP](#), [Java](#), [Python](#), [Node.js](#), [.NET](#), and [Ruby](#) applications, although it can also support other languages via "custom runtimes".
 - Python web frameworks that run on Google App Engine include [Django](#), [CherryPy](#), [Pyramid](#), [Flask](#), [web2py](#) and [webapp2](#), as well as a custom Google-written webapp framework and several others designed specifically for the platform that emerged since the release.
 - Any Python framework that supports the [WSGI](#) using the CGI adapter can be used to create an application; the framework can be uploaded with the developed application. Third-party libraries written in pure Python may also be uploaded.
 - Google App Engine supports many Java standards and frameworks. Core to this is the [servlet 2.5 technology](#) using the open-source [Jetty Web Server](#), along with accompanying technologies such as [JSP](#). [JavaServer Faces](#) operates with some workarounds. A newer release of App Engine Standard Java in Beta supports Java8, Servlet 3.1 and Jetty9.
-
- Though the integrated database, [Google Cloud Datastore](#), may be unfamiliar to programmers, it is accessed and supported with [JPA](#), [JDO](#), and by the simple low-level API.
 - There are several alternative libraries and frameworks you can use to model and map the data to the database such as [Objectify](#), [Slim3](#) and [Jello framework](#).
 - The [Spring Framework](#) works with GAE. However, the Spring Security module (if used) requires workarounds. [Apache Struts 1](#) is supported, and [Struts 2](#) runs with workarounds.
 - The [Django web framework](#) and applications running on it can be used on App Engine with modification.
 - Django-nonrel aims to allow Django to work with non-relational databases and the project includes support for App Engine.