

Cloud_Computing_Unit_1.pdf

Cloud_Computing_Unit_2.pdf

Cloud_Computing_Unit_3.pdf

Cloud_Computing_Unit4.pdf

cloud_cube_model_v1.0.pdf

CloudComputingnew.ppt

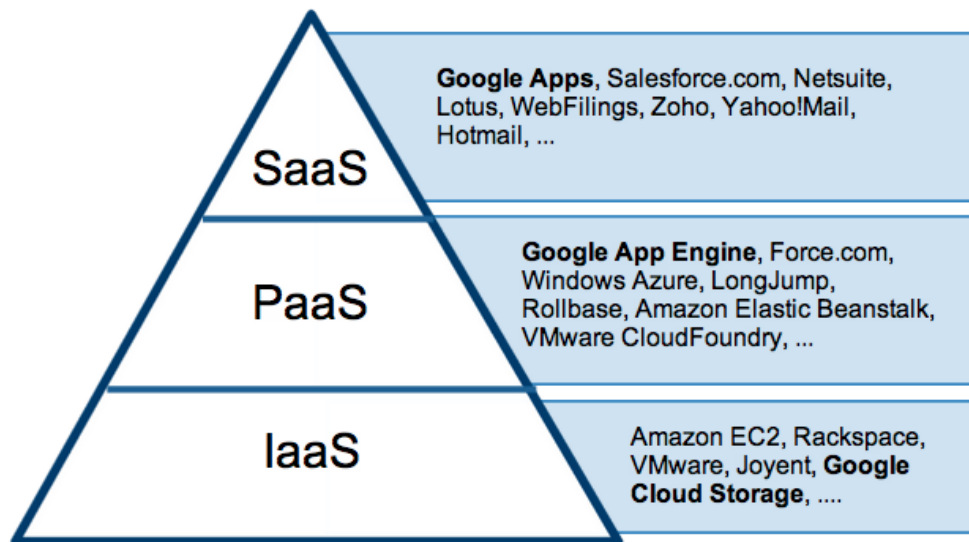
02a_CloudPrimer1.pdf

Cloud Computing

What is Cloud Computing?

Cloud computing is the use of computing resources (hardware and software) that are delivered as a service over a network (typically the Internet). The “cloud” in cloud computing can be defined as the set of hardware, networks, storage, services, and interfaces that combine to deliver aspects of computing as a service. Cloud services include the delivery of software, infrastructure, and storage over the Internet. Most of the cloud computing has following features:

- **It's virtual**
- **It can be secure**
- **It's flexible and scalable**
- **It can be affordable**
- **It can be secure and affordable**

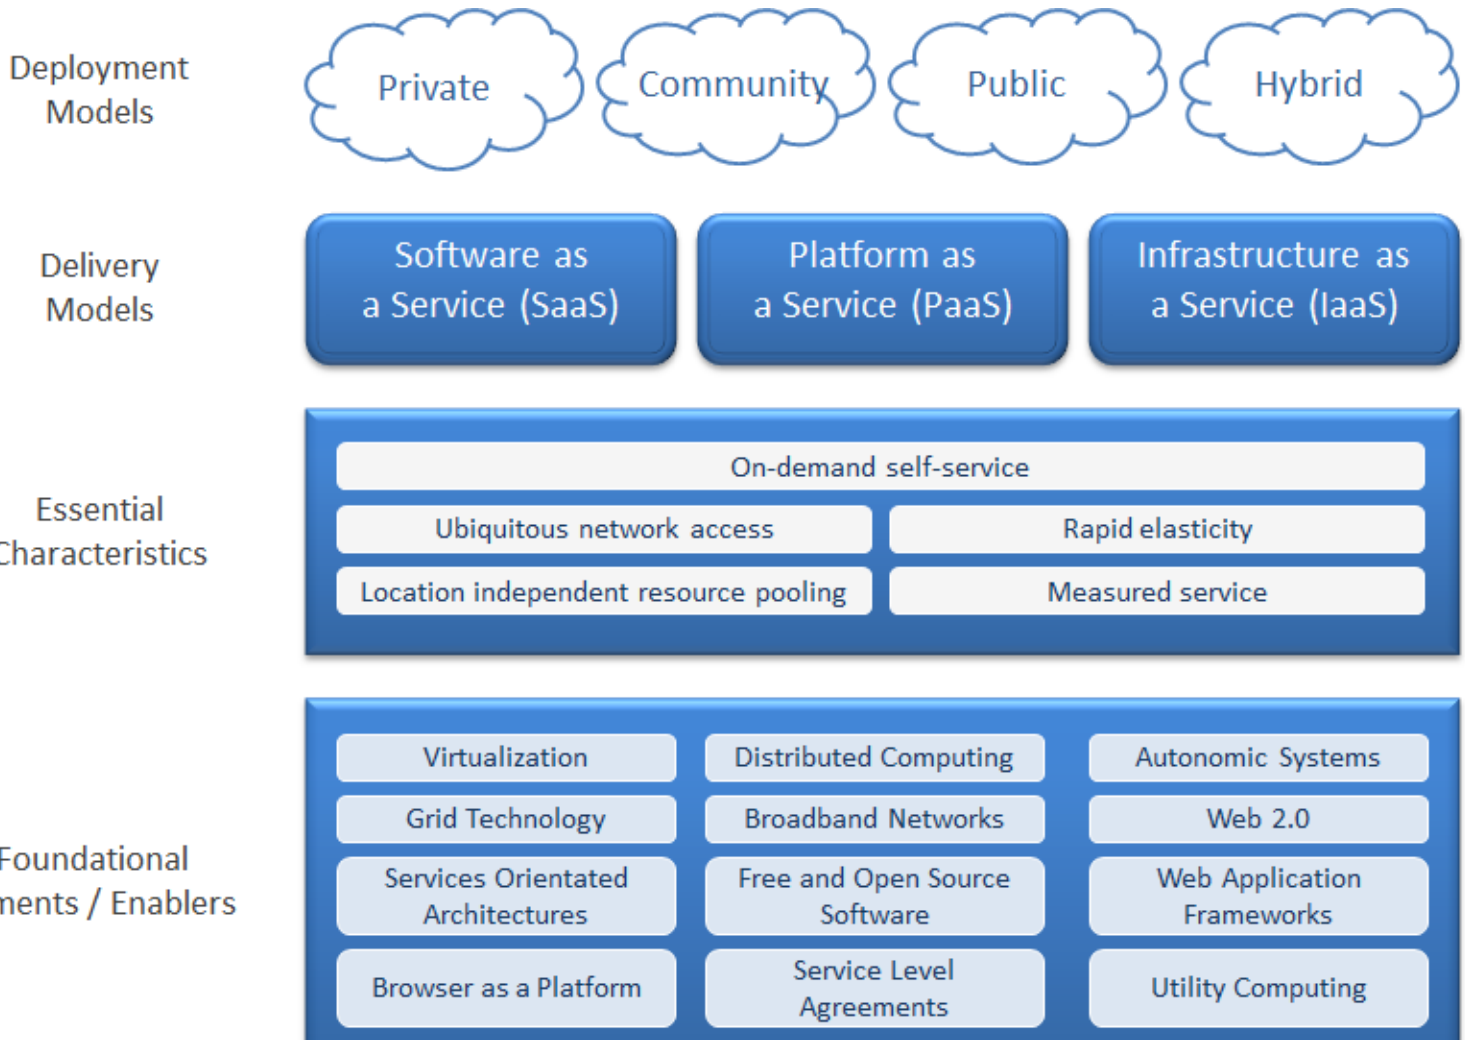


Most IT departments are forced to spend a significant portion of their time on frustrating implementation, maintenance, and upgrade projects that too often don't add significant value to the company's bottom line. Increasingly, IT teams are turning to cloud computing technology to minimize the time spent on lower-value activities and allow IT to focus on strategic activities with greater impact on the business

In a cloud computing system, there's a significant workload shift. Local computers no longer have to do all the heavy lifting when it comes to running applications. The network of computers that make up the cloud handles them instead. Hardware and software demands on the user's side decrease. The only thing the user's computer needs to be able to run is the cloud computing system's interface software, which can be as simple as a Web browser, and the cloud's network takes care of the rest. Some of the advantages of cloud based services are:

- **Proven Web-services integration.** By their very nature, cloud computing technology is much easier and quicker to integrate with your other enterprise applications (both traditional software and cloud computing infrastructure-based), whether third-party or homegrown.
- **World-class service delivery.** Cloud computing infrastructures offer much greater scalability, complete disaster recovery, and impressive uptime numbers.
- **No hardware or software to install:** a 100% cloud computing infrastructure. The beauty of cloud computing technology is its simplicity... and in the fact that it requires significantly fewer capital expenditures to get up and running.
- **Faster and lower-risk deployment.** You can get up and running in a fraction of the time with a cloud computing infrastructure. No more waiting months or years and spending millions of dollars before anyone gets to log into your new solution. Your cloud computing technology applications are live in a matter of weeks or months, even with extensive customization or integration.
- **Support for deep customizations.** Some IT professionals mistakenly think that cloud computing technology is difficult or impossible to customize extensively, and therefore is not a good choice for complex enterprises. The cloud computing infrastructure not only allows deep customization and application configuration, it preserves all those customizations even during upgrades. And even better, cloud computing technology is ideal for application development to support your organization's evolving needs.
- **Empowered business users.** Cloud computing technology allows on-the-fly, point-and-click customization and report generation for business users, so IT doesn't spend half its time making minor changes and running reports.
- **Automatic upgrades that don't impact IT resources.** Cloud computing infrastructures put an end to a huge IT dilemma: If we upgrade to the latest-and-greatest version of the application, we'll be forced to spend time and resources (that we don't have) to rebuild our customizations and integrations. Cloud computing technology doesn't force you to decide between upgrading and preserving all your hard work, because those customizations and integrations are automatically preserved during an upgrade.

CLOUD AT A GLANCE



Emergence/ History of Cloud Computing

One of the first milestones for cloud computing was the arrival of Salesforce.com in 1999, which pioneered the concept of delivering enterprise applications via a simple website. The services firm paved the way for both specialist and mainstream software firms to deliver applications over the internet.

The next development was Amazon Web Services in 2002, which provided a suite of cloud-based services including storage, computation and even human intelligence through the Amazon Mechanical Turk.

Then in 2006, Amazon launched its Elastic Compute cloud (EC2) as a commercial web service that allows small companies and individuals to rent computers on which to run their own computer applications.

Another big milestone came in 2009, as Web 2.0 hit its stride, and Google and others started to offer browser-based enterprise applications, though services such as Google Apps.

The most important contribution to cloud computing has been the emergence of "killer apps" from leading technology giants such as Microsoft and Google. When these companies deliver services in a way that is reliable and easy to consume, the knock-on effect to the industry as a whole is a wider general acceptance of online services.

Then came mature virtualization technologies in 2009-tilldate that changed landscape of cloud computing. Private, Public and hybrid cloud were dominant cloud types in Enterprise Level. Server and Storage Consolidation were major works in the cloud industry.

Cloud Based Services

- **Hosted Desktops:** The hosted virtual desktop is just like normal PC desktop, except that applications, data and user profile are stored in a secure data center. It is designed to replace traditional desktop PC environment, and provides the same level of functionality and performance as normal PC.
- **Hosted Websites/ Emails:** Websites and emails are hosted in cloud based servers. The main advantages of hosted websites and emails is the saving of cost used for procuring new servers, managing servers, power system, experts, maximum uptime, scalability and security.
- **Hosted Telephony:** Hosted telephony is the way of taking all the telephone infrastructure from on premises to some cloud based telephony system. The advantages of this technique are that it is easy for the service provider to upgrade the system, there's very little local infrastructure needed at your office location and it provides good

flexibility especially across multiple location. Hosted PBX companies handle call routing, or switching, at their own location and are responsible for managing all of the PBX equipment and software involved in the virtual PBX service. Hosted PBX services can function over the Public Switched Telephone Network (PSTN) over the Internet (hosted IP PBX via Internet telephony, or VoIP), or over a combination of the two. Because it is a hosted system, there's no high initial cost for buying the expensive hardware. There's also never an upgrade cost for changing hardware. Unlike traditional PBX services, hosted PBX makes advanced business technology available for even small businesses.

- **Cloud Storage:** Cloud storage is a service model in which data is maintained, managed and backed up remotely and made available to users over a network (typically the Internet). Availability and Scalability are the major advantages of using cloud based storage. By availability means, any user can access their data residing in cloud from any location via Internet. By Scalability means, user can increase their storage space as per need and load by paying they use the storage quantity.

Grid Computing:

Grid computing is a computer network in which each computer's resources are shared with every other computer in the system. Processing power, memory and data storage are all community resources that authorized users can tap into and leverage for specific tasks. A grid computing system can be as simple as a collection of similar computers running on the same operating system or as complex as inter-networked systems comprised of every computer platform you can think of. A grid computer is connected through a super-fast network and share the devices like disk drives, mass storage, printers and RAM.

In general, a grid computing system requires:

- **At least one computer, usually a server, which handles all the administrative duties for the system:** This type of computer are sometimes referred as a control node. All the administrative tasks required for computing are handled by this control node. The control node must prioritize and schedule tasks across the network. It's the control node's job to determine what resources each task will be able to access. The control node must also monitor the system to make sure that it doesn't become overloaded.
- **A network of computers running special grid computing network software:** These computers act both as a point of interface for the user and as the resources the system will look into for different applications.
- **A collection of computer software called middleware:** The purpose of middleware is to allow different computers to run a process or application across the entire network of machines. Middleware is the workhorse of the grid computing system. Without it, communication across the system would be impossible.

Grid Vs Cloud Computing

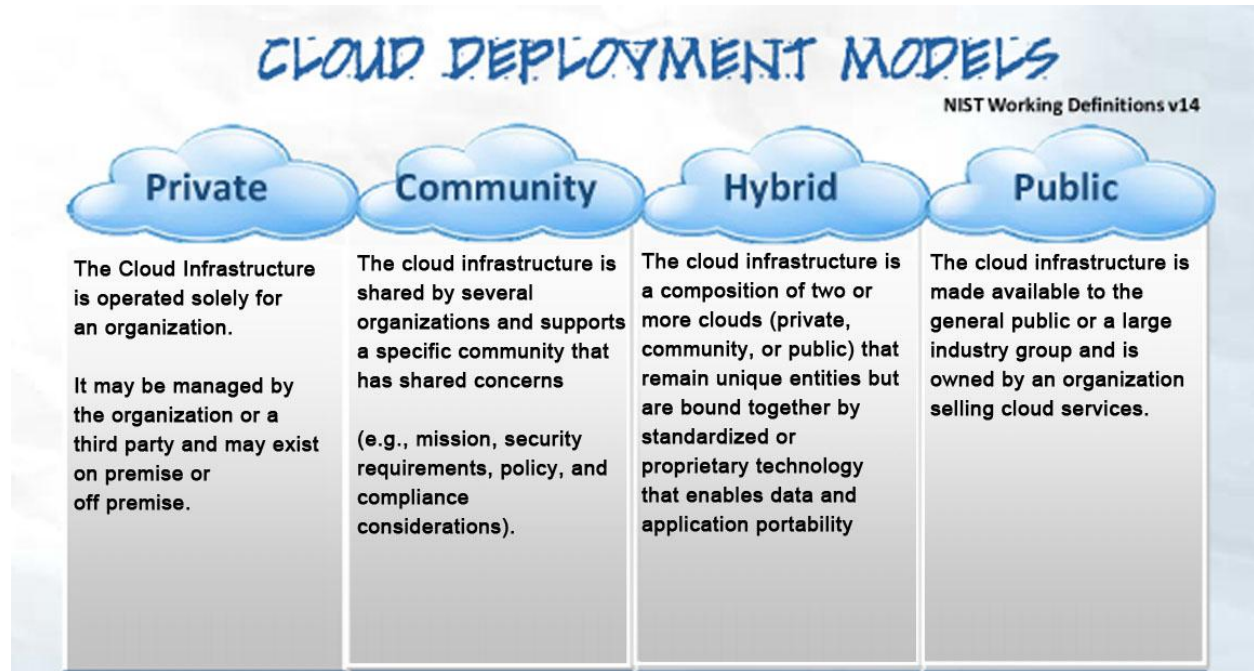
Parameter	Grid computing	Cloud computing
Goal	Collaborative sharing of resources	Use of everything as a service
Level of abstraction	Low (more details)	High (eliminate details)
Degree of scalability	Normal	High
Multitask	Yes	Yes
Transparency	Low	High
Time to	Not real-time	Real-time services
Requests type	Few but large allocation	Lots of small allocation
Virtualization	Not a necessity	Vital
Portal accessible	Via a DNS system	Only using IP (no DNS registered)
Transmission	Suffered from internet delays	Was significantly fast
Security	Low (grid certificate service)	High (Virtualization)
Infrastructure	Low level command	High level services (SaaS)
Operating System	Any standard OS	A hypervisor (VM) on which multiple OSs run
Ownership	Multiple	Single
Interconnection network	Mostly internet with latency and low bandwidth	Dedicated, high-end with low latency and high bandwidth
Service negotiation	SLA based	SLA based
User management	Decentralized and also Virtual Organization (VO)-based	Centralized or can be delegated to third party
Resource management	Distributed	Centralized/Distributed
Allocation/Scheduling	Decentralized	Both centralized/decentralized
Failure management	Limited (often failed tasks/applications are restarted)	Strong (VMs can be easily migrated from one node to other)
Pricing of services	Dominated by public good or privately assigned	Utility pricing, discounted for larger Customers
Type of service	CPU, network, memory, bandwidth, device, storage,...	IaaS, PaaS, SaaS, Everything as a service
Example of real world	SETI, BOINC, Folding@home, GIMPS	Amazon Web Service (AWS), Google apps
Response Time	Can't be serviced at a time and need to be scheduled	Real-time

Critical object	Computer resource	Service
Number of users	Few	More
Resource	Limited (because hardware are limited)	Unlimited
Configuration	Difficult as users haven't administrator privilege	Very easy to configure
Future	Cloud computing	Next generation of internet

Components of Cloud Computing

- a. **The Client- The End User:** Everything ends with the client. The hardware components, the application and everything else developed for cloud computing will be used in the client. Client systems has some application installed which enables them to connect to cloud software or some infrastructure.
- b. **The Service (Functions in Cloud Computing):** Cloud computing always has a purpose. One of the main reasons cloud computing become popular is due to the adoption of businesses as the easier way to implement business processes. It has some standard service or procedure of interfacing/ connecting client computers with cloud infrastructure. Cloud computing is all about processes and the services launched through cloud computing always has to deal with processes with an expected output.
- c. **The Application:** Application is the core of what users are going to use. It is the mainstay of what users are wanting for their daily operations. Application are normally a program that users use to connect cloud infrastructure either with web interface or any application interface. In simple way, we can define application as a software that end user uses to do their operations in which their main data resides in cloud.
- d. **The Platform:** The platform is where all the applications and services are based upon. The platform usually comes as the programming language such as Ajax (Asynchronous JavaScript and XML) or Ruby on Rails. In simple way, the platform is the cloud infrastructure where it provides application and service the base to operate. It is the environment provided by cloud vendors which enables all the application to operate and services to operate.
- e. **The Storage:** The last and most critical components in cloud computing is the storage. Everything that the application knows and the functions that could be provided by service are possible through storage. Modern day cloud storage is based on highly virtualized infrastructure and has the same characteristics as cloud computing in terms of agility, scalability, elasticity and multi-tenancy. Some cloud storage systems are small operations, while others are so large that the physical equipment can fill up an entire warehouse.

Cloud Computing Deployment Model



Public Cloud: This is the deployment model that most commonly described as cloud computing. In this model, all of the physical resources are owned and operated by a third party cloud computing provider. The provider services multiple clients that may consist of individuals or corporations utilizing these resources through the public Internet. Services can be dynamically provisioned and are billed based on usage alone. This model provides the highest degree of cost savings while requiring the least amount of overhead.

This model is best suited for business requirements wherein it is required to manage load spikes, host SaaS applications, utilize interim infrastructure for developing and testing applications, and manage applications which are consumed by many users that would otherwise require large investment in infrastructure from businesses.

Private Cloud: Private cloud computing systems emulate public cloud service offerings within an organization's boundaries to make services accessible for one designated organization. Private cloud computing systems make use of virtualization solutions and focus on consolidating distributed IT services often within data centers belonging to the company. The chief advantage of these systems is that the enterprise retains full control

over corporate data, security guidelines, and system performance. This model doesn't bring much in terms of cost efficiency: it is comparable to buying, building and managing your own infrastructure. Still, it brings in tremendous value from a security point of view. In addition to security reasons, this model is adopted by organizations in cases where data or applications are required to conform to various regulatory standards, which may require data to be managed for privacy and audits that govern the corporation.

Hybrid Cloud: This can be a combination of private and public clouds that support the requirement to retain some data in an organization, and also the need to offer services in the cloud. A company may use internal resources in a private cloud maintain total control over its proprietary data. It can then use a public cloud storage provider for backing up less sensitive information. At the same time, it might share computing resources with other organizations that have similar needs. By combining the advantages of the other models, the hybrid model offers organizations the most flexibility.

This model is also used for handling cloud bursting, which refers to a scenario where the existing private cloud infrastructure is not able to handle load spikes and requires a fallback option to support the load. Hence, the cloud migrates workloads between public and private hosting without any inconvenience to the users.

Community Cloud: In the community deployment model, the cloud infrastructure is shared by several organizations with the same policy and compliance considerations. This helps to further reduce costs as compared to a private cloud, as it is shared by larger group.

A community cloud contains features of the public and private cloud models. Like a public cloud, the community cloud may contain software, data storage, and computing resources that are utilized by multiple organizations. Where this model differs from the public model is that the infrastructure is only utilized by a group of organizations that are known to each other. Similarly to a private cloud, these organizations are responsible for the operation of their own infrastructure. The community cloud model can provide greater cost savings than the private cloud while offering some of its security features. This model is best suited for organizations that share common requirements such as security or legal compliance policies. It can be managed by the member organizations or by a third party provider.

Benefits of Using Cloud Model/ Why Switch from Traditional IT to cloud/ Goal of Cloud Computing

- **Reduced Spending on Technology Infrastructure:** Moving to cloud computing may reduce the cost of managing and maintaining your IT systems. Rather than purchasing expensive systems and equipment for your business, you can reduce your costs by using the resources of your cloud computing service provider. You may be able to reduce your operating costs because
 - the cost of system upgrades, new hardware and software may be included in your contract
 - you no longer need to pay wages for expert staff
 - your energy consumption costs may be reduced
 - There are fewer time delays.
- **Globalizing Workspace/ Easy Accessibility:** Globalizing your workspace or system may add additional agility and effectiveness to your system. For example, you have the ability to access data from home, on holiday, or via the commute to and from work (providing you have an internet connection). If you need access to your data while you are off-site, you can connect to your virtual office, quickly and easily.
- **Improve Flexibility and Scalability:** Your business can scale up or scale down your operation and storage needs quickly to suit your situation, allowing flexibility as your needs change. Rather than purchasing and installing expensive upgrades yourself, your cloud computer service provider can handle this for you. Using the cloud frees up your time so you can get on with running your business.
- **Better Resource Utilization:** Using technologies such as virtualization and distributed computing, computing resources can be used fully up to their potentials. In cloud model, you can pay whatever you use for your operational purpose.
- **Backup and Disaster Recovery:** Since all your data is stored in the cloud, backing it up and restoring the same is relatively much easier than storing the same on a physical device. Furthermore, most cloud service providers are usually competent enough to handle recovery of information. Hence, this makes the entire process of backup and recovery much simpler than other traditional methods of data storage.

Also adding to backup management, cloud infrastructure can be used as a very strong disaster recovery site as they do have enough infrastructure and technology to be a potent DR Site.

- **Guaranteed Uptime and Strong Service Level Agreements (SLA):** Cloud infrastructure by nature is built for robustness and high availability. Uptime is a major attribute for rapid popularity of cloud computing as all the services provided by cloud vendors should be guaranteed to be working all the time. Hence, the service level agreements between the cloud provider and customer are supposed to be very strong in case of availability and security. By achieving strong SLA's, customer can operate very effectively with zero downtime.
- **Helps smaller business compete:** Cloud infrastructure is a global platform where smaller and large scale business both can use the resources. Hence, smaller business industries can be equally competent in regard to IT Infrastructure and can focus more on their business side. By decreasing their capital expenses (CAPEX) they can move forward by using financials to business development.

Legal Issues in Cloud Computing:

- **Confidentiality:** Data in enterprise world is as important as anything. Placing your data in cloud infrastructure is supposed to be vulnerable and insecure. Hence, before and after moving to cloud infrastructure, organizations should carefully judge whether their data is managed confidentially or not.
- **Liability and responsibility:** Liability and responsibility is another legal issues that has to be addressed by cloud vendor as well as customers. It should be regularly monitored to investigate that whether cloud vendors has performed their duties in accordance to Service Level Agreements (SLA) or not. As is cloud definition, customer has to ensure that cloud vendors has sufficient infrastructure, proper backup policy, business continuity plan and all the prerequisite to host valuable data.
- **Compliance:** Before customers will entrust their IT needs to managed or cloud services, they need two things: first, assurance that cloud infrastructure is secure and compliant, and second, visibility into their own security and compliance stance in cloud or managed infrastructure. Cloud vendors should ensure the security and compliance of their customer with powerful incident management capabilities, immediate alerts about suspicious activities, and access to detailed

forensic data. It should give its customer all the components required to deliver the compliance and security reports and dashboards they demand.

- **Data protection, safety and recovery:** Data in cloud as said should be safe enough to be trusted and protected from various attacks. Safety and protection is not only enough for operating in cloud but it should have some standard recovery mechanism to recover data in case of failure of system.
- **Copyright and Ownership:** Even though data may be residing in cloud infrastructure in any part of the world, data should be owned by customer and it should have legal obligation of being owned by customer themselves. Data once migrated to the cloud data centers should be completely owned and should be protected by some copyright. Customers should be aware of intentional duplication of data, data being copied or any leakage of data.
- **Data portability:** What if customers want to shift data/ app to other cloud vendors? Data portability is a major hurdle for any customer to migrate from on cloud vendor to another. Is there any legal obligation of cloud vendor regarding the move or not? What if existing cloud vendor do not allow customer to migrate data to other provider? These questions should be clear enough for both customer and vendor and there should be defining answer for these questions.
- **Right to Audit:** IT audit in cloud infrastructure is a necessity for maintaining compliance of cloud vendor as well as customer. Before moving to the cloud, and ideally during the procurement process, you should know your risk appetite and how it feeds the control environment — and then determine the potential cloud provider's risk appetite. Security, Risk, Compliance are some of the factors that customers need to check periodically for the risk free operations. And while moving to cloud data centers, customers should be legally enforced to have the right to audit their hardware, software, systems and applications.
- **Termination or Suspension Contract:** Cloud computing agreement can be terminated on the account of various reasons. The contract may expire at the end of its stipulated term or it may be terminated for default or material breach of terms of contract. User may also want to terminate the contract to migrate to a better or more cost effective cloud computing service. The user's data is most vulnerable after the termination of contract and in most cases service provider has no legal duty or liability to handle the user's data properly unless stipulated otherwise in the cloud computing contract. Hence a careful steps should be taken whenever for various reason a customer terminates the services from cloud service provider.

Characteristics of Cloud Computing

- **Service Oriented:** The defining characteristics of cloud computing is the service oriented feature. All the IT related services are hosted in cloud infrastructure. Companies should not have to buy expensive servers, network equipment's and invest on expensive manpower. All they need is to subscribe to any cloud service provider and get what they want. In this way, we can decrease our capital expenditure and move to operate via Operating expenditure.
- **Broad Network Access:** Cloud Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms such as mobile phones, laptops and PDAs.
- **On Demand:** A consumer can provision computing capabilities, such as server processing and network storage, as needed automatically without requiring human interaction with each service's provider. computer services such as email, applications, network or server service can be provided without requiring human interaction with each service provider. Cloud service providers providing on demand self-services include Amazon Web Services (AWS), Microsoft, Google, IBM and Salesforce.com.
- **Reliability, Elasticity and scalability:** The cloud is reliable in the sense that the infrastructure setup for cloud is robust and backed up for high availability. It is some resilient replication and backup strategy that is targeted for huge customer base. The cloud is **elastic**, meaning that resource allocation can get bigger or smaller depending on demand. Elasticity enables **scalability**, which means that the cloud can scale upward for peak demand and downward for lighter demand. Scalability also means that an application can scale when adding users and when application requirements change.
- **Resource Pooling (Processor, Memory, and Storage):** Cloud infrastructure should have features of resource pooling i.e. resources (CPU, Memory, Disk) should be categorized in a hierarchy as per the need of computing. Resource pooling is mainly used for utilizing servers up to its potential. Since most of the times server

resources are unused, we can use the concept of virtualization to pool its resources.

- **Measured Service (Pay per Use):** Cloud computing resource usage can be measured, controlled, and reported providing transparency for both the provider and consumer of the utilized service. Cloud computing services use a metering capability which enables to control and optimize resource use. This implies that just like air time, electricity or municipality water IT services are charged per usage metrics – pay per use. The more you utilize the higher the bill. Just as utility companies sell power to subscribers, and telephone companies sell voice and data services, IT services such as network security management, data center hosting or even departmental billing can now be easily delivered as a contractual service.
- **Multi Tenancy:** Multi tenancy refers to a principle in IT infrastructure where a single instance of the software runs on a server, serving multiple client organizations (tenants). With a multitenant architecture, a software application is designed to virtually partition its data and configuration, and each client organization works with a customized virtual application instance. Each customer does its own work without interfering other customer even though they are hosted at the same platform.

Challenges in Cloud Computing/ Security Issues

- Privacy
- Data Security
- Ownership
- RASP (Reliability, Availability, Scalability and Performance)
- Data Recovery and Backup
- Cross Country Data Migration and Portability
- Multiplatform Support
- Intellectual Property
- Misuse
- Real Time Processing
- Compliance

Distributed Computing in Grid and Cloud Computing

A **distributed computing** consists of multiple software components that are on multiple computers, but run as a single system. The computers that are in a distributed system can be physically close together and connected by a local network, or they can be geographically distant and connected by a wide area network. A distributed system can consist of any number of possible configurations, such as mainframes, personal computers, workstations, minicomputers, and so on. The goal of distributed computing is to make such a network work as a single computer.

Distributed computing, as one can imagine, is where the computing elements of a network are spread over a large geographical area. Both cloud and grid computing are prime examples of distributed computing architectures.

Another type of distributed computing is known as **grid computing**. Grid computing consists of many computers operating together remotely and often simply using the idle processor power of normal computers. Grid provides the sharing of:

- Computational resources
- Storage elements
- Specific applications
- Equipment
- Other

While there are many similarities between grid and cloud computing, it is the differences that matter most. Grid computing is better suited for organizations with large amounts of data being

requested by a small number of users (or few but large allocation requests), whereas cloud computing is better suited to environments where there are a large number of users requesting small amounts of data (or many but small allocation requests).

Cloud computing is basically a sales and distribution model for various types of resources over the internet, while distributed computing can be identified as a type of computing, which uses a group of machines to work as a single unit to solve a large scale problem. Distributed computing achieves this by breaking the problem up to simpler tasks, and assigning these tasks to individual nodes.

Compared to other distributed systems such as grids or clusters, cloud computing solutions give enterprises significantly more flexibility. They can dispense with IT infrastructures of their own and only have to pay for the resources and services they actually use (“pay-per-use”/ “pay as you go”). These can be dynamically adapted to changed business requirements and processes with the help of virtualization technologies and service oriented, distributed software systems.

CHAPTER 2

Communication as a Service (CaaS)

Communications as a Service (CaaS) is an outsourced enterprise communications solution that can be leased from a single vendor. Such communications can include voice over IP (VoIP or Internet telephony), instant messaging (IM), collaboration and videoconference applications using fixed and mobile devices. CaaS has evolved along the same lines as Software as a Service (SaaS).

CaaS brings social networking, cloud computing, and smartphones together, providing cloud-technologies that let users communicate via voice, text, and rich media on whatever device they prefer to use. To compete in this marketplace, software vendors, enterprises, and service providers must introduce communications-enhanced services that meet a surging need for value, efficiency, cost reduction, and convenience.

Through the hosted model, the CaaS provider manages everything, from the telecommunications infrastructure to the software integration platform for delivering communications offered at a guaranteed Quality of Service (QoS). As a result, businesses can selectively deploy communications features, applications, and devices on a pay-as-you-go, as-needed basis, reducing risks while eliminating capital costs associated with new services.

CaaS offers flexibility and expandability that small and medium-sized business might not otherwise afford, allowing for the addition of devices, modes or coverage on demand. The network capacity and feature set can be changed from day to day if necessary so that functionality keeps pace with demand and resources are not wasted. There is no risk of the system becoming obsolete and requiring periodic major upgrades or replacement

Advantages of Communication as a Service (CaaS)

- **Fully Integrated Enterprise Class Unified Communication:** By managing the LAN/WAN, the vendor can guarantee consistent Quality of Service (QoS) from the desktop across the VoIP backbone and back again. Advanced Unified Communications features such as Outlook integration, soft phones, real-time presence, chat, multimedia conferencing, video calling, unified messaging and mobility are also part of a standard CaaS deployment. And with CaaS, the feature set can continue to evolve. Development and introduction of new features and applications are faster, easier and more economical because the service provider is doing the work for multiple end users across a scalable platform
- **No Upfront Capital Expenses:** Since cloud services are supposed to lower capital expenditure and focus more on operating expenditure, by implementing CaaS, consumers can build up their communication infrastructure without any upfront cost. They just need to pay it as a service.
- **Flexibility in Features:** Since cloud is a multi-tenant architecture, cloud vendors have to manage multiple customers and look after the features that they want. What this allows cloud vendor is to add more advanced features and flexibility in their service model. Economies of scale also

mean that the service provider is not tied to a single vendor investment and can leverage best-of-breed providers like Cisco, Microsoft and Nortel much more economically than an independent enterprise.

- **No Risk of Obsolescence:** Technology changes rapidly and are obsolete within few years of introduction. With CaaS, companies are always privileged with new technologies as cloud vendors keep on updating their equipment's and technologies to sustain in market.
- **No Data Center Cost:** As a prime advantages of cloud computing, while using CaaS infrastructure, organization need not invest on expensive servers, cooling system and electric equipment. With monthly/ yearly recurring cost, organization can dramatically cut down the management cost of data center as well.
- **Guaranteed Business Continuity:** With CaaS, organization can be hugely benefitted with guaranteed business continuity as cloud service providers proactively plans for Business Continuity Planning for their customers. Service uptime is guaranteed even if any catastrophic disaster strikes.

Unified Communication

Unified Communications (UC) might be defined as communications integrated to optimize business performance, and can describe a seamless set of voice, video and Web collaboration applications designed to enable advanced connectivity between employees, customers, partners and other stakeholders available on any device.

In the large enterprise setting, Unified Communications allow employees to send and receive messages on various medium – guided by presence information or employee preference – switching calls between home lines and a mobile device, or accessing a voicemail via text or email. When well implemented, UC can save money, optimize business processes and improve communications by better connecting devices, media and applications, reducing latency, and opening new forms of collaboration.

Some of the benefits of Unified Communication are as below:

- a. **Efficient Communication:** Users can communicate more efficiently by having access to all communications at one time and being free to share, forward, or manage them in the way that's most convenient or effective for the given communication.
- b. **Anywhere Access:** Unified messaging provides alternative methods of accessing communications. By merging e-mail, voice, and other communications, users can get voice messages in e-mail, have e-mail dictated over the phone, or access communications via the Web.
- c. **Collaboration:** Collaborative approach of communication can be helpful in achieving immediacy, simplicity and interoperability of communication. In this way, business can sense and respond more quickly to changing environment.
- d. **Business Process Integration:** Business processes can be integrated to each other for more efficiency. Many business entities and functions can be incorporated in one platform to make processes more agile and responsive.

Infrastructure as a Service

Infrastructure as a Service (IaaS) is one of the three fundamental service models of cloud computing alongside Platform as a Service (PaaS) and Software as a Service (SaaS).

Infrastructure as a Service is a provision model in which an organization outsources the equipment used to support operations, including storage, hardware, servers and networking components. The service provider owns the equipment and is responsible for housing, running and maintaining it. The client typically pays on a per-use basis. Clients are able to self-provision this infrastructure, using a Web-based graphical user interface that serves as an IT operations management console for the overall environment. API access to the infrastructure may also be offered as an option.

A typical Infrastructure as a Service offering can deliver the following features and benefits:

- **Scalability;** resource is available as and when the client needs it and, therefore, there are no delays in expanding capacity or the wastage of unused capacity
- **No investment in hardware;** the underlying physical hardware that supports an IaaS service is set up and maintained by the cloud provider, saving the time and cost of doing so on the client side
- **Utility style costing;** the service can be accessed on demand and the client only pays for the resource that they actually use
- **Location independence;** the service can usually be accessed from any location as long as there is an internet connection and the security protocol of the cloud allows it
- **Physical security of data center locations;** services available through a public cloud, or private clouds hosted externally with the cloud provider, benefit from the physical security afforded to the servers which are hosted within a data center
- **No single point of failure;** if one server or network switch, for example, were to fail, the broader service would be unaffected due to the remaining multitude of hardware resources and redundancy configurations. For many services if one entire data center were to go offline, never mind one server, the IaaS service could still run successfully.

On Demand Computing: On-demand (OD) computing is an increasingly popular enterprise model in which computing resources are made available to the user as needed. The resources may be maintained within the user's enterprise, or made available by a service provider. The on-demand model evolved to overcome the challenge of being able to meet fluctuating resource demands efficiently. Because demand for computing resources can vary drastically from one time to another, maintaining sufficient resources to meet peak requirements can be costly.

On-demand computing products are rapidly becoming prevalent in the marketplace. Computer Associates, HP, IBM, Microsoft, and Sun Microsystems are among the more prominent on-demand


vendors. These companies refer to their on-demand products and services by a variety of names. Concepts such as grid computing, utility computing, autonomic computing, and adaptive management seem very similar to the concept of on-demand computing.

The major advantage of On Demand Computing (ODC) is low initial cost, as computational resources are essentially rented when they are required. This provides cost savings over purchasing them outright.

(Amazon EC2): Amazon Elastic Compute Cloud (EC2) is a central part of Amazon cloud computing platform Amazon Web Services (AWS). EC2 allows users to rent virtual computers on which to run their own computer applications. EC2 allows scalable deployment of applications by providing a Web service through which a user can boot an Amazon Machine Image to create a virtual machine, which Amazon calls an "instance", containing any software desired. A user can create, launch, and terminate server instances as needed, paying by the hour for active servers, hence the term "elastic".

It reduces the time required to obtain and boot new server instances to minutes, allowing customers to quickly scale capacity as their computing demands dictate. It changes the economics of computing by allowing clients to pay only for capacity they actually use. It provides developers the tools needed to build failure-resilient applications and isolate themselves from common failure scenarios.

To use the EC2, a subscriber creates an Amazon Machine Image (AMI) containing the operating system, application programs and configuration settings. Then the AMI is uploaded to the Amazon Simple Storage Service (Amazon S3) and registered with Amazon EC2, creating a so-called AMI identifier (AMI ID). Once this has been done, the subscriber can requisition virtual machines on an as-needed basis. Capacity can be increased or decreased in real time from as few as one to more than 1000 virtual machines simultaneously. Billing takes place according to the computing and network resources consumed.



The idea of EC2 is to lighten the cost of buying servers to host a system, but more importantly to eliminate the wasted time systems engineers devote to managing hard assets. Instead of buying servers to increase capacity or add new features, you simply buy more gigabytes on EC2. Amazon sells it in subscription form, with subscriptions based on how much capacity you use.

Characteristics of Amazon EC2

http://en.wikipedia.org/wiki/Amazon_ec2

- a. Persistent Storage
- b. Elastic IP Addresses
- c. Amazon Cloud Watch
- d. Automated Scaling
- e. Reliability

Monitoring as a Service

Monitoring-as-a-Service is an outsourced service to provide security mainly to platforms that are run on the Internet for conducting business. MaaS became highly popular in the last decade. Since the advent of Cloud Computing, its popularity has increased even more. Safe monitoring involves protecting a company or other institution / organization from cyber threats, in which a team prepared is crucial to maintain the confidentiality, integrity and access to IT assets.

Many industry regulations require organizations to monitor their security environment, server logs, and other information assets to ensure the integrity of these systems. However, conducting effective security monitoring can be a daunting task because it requires advanced technology, skilled security experts, and scalable processes—none of which come cheap. MaaS security monitoring services offer real-time, 24/7 monitoring and nearly immediate incident response across a security infrastructure—they help to protect critical information assets of their customers

Protection against Internal and External Threats

Security monitoring services can improve the effectiveness of a customer security infrastructure by actively analyzing logs and alerts from infrastructure devices around the clock and in real time. Typical services provided by many MaaS vendors are described below.

Early Detection

An early detection service detects and reports new security vulnerabilities shortly after they appear. Generally, the threats are correlated with third-party sources, and an alert or report is issued to customers. This report is usually sent by email to the person designated by the company. Security vulnerability reports, aside from containing a detailed description of the vulnerability and the platforms affected, also include information on the impact the exploitation of this vulnerability would have on the systems or applications previously selected by the company receiving the report. Most often, the report also indicates specific actions to be taken to minimize the effect of the vulnerability, if that is known.

Platform, Control, and Services Monitoring

Platform, control, and services monitoring is often implemented as a dashboard interface and makes it possible to know the operational status of the platform being monitored at any time. It is accessible from a web interface, making remote access possible. Each operational element that is monitored usually provides an operational status indicator, always taking into account the critical impact of each element. This service aids in determining which elements may be operating at or near capacity or beyond the limits of established parameters. By detecting and identifying such problems, preventive measures can be taken to prevent loss of service.

Intelligent Log Centralization and Analysis

Intelligent log centralization and analysis is a monitoring solution based mainly on the correlation and matching of log entries. Such analysis helps to establish a baseline of operational performance and provides an index of security threat. Alarms can be raised in the event an incident moves the established baseline parameters beyond a stipulated threshold. These types of sophisticated tools are used by a team of security experts who are responsible for incident response once such a threshold has been crossed and the threat has generated an alarm or warning picked up by security analysts monitoring the systems.

Vulnerabilities Detection and Management

Vulnerabilities detection and management enables automated verification and management of the security level of information systems. The service periodically performs a series of automated tests for the purpose of identifying system weaknesses that may be exposed over the Internet, including the possibility of unauthorized access to administrative services, the existence of services that have not been updated, the detection of vulnerabilities such as phishing, etc. The service performs periodic follow-up of tasks performed by security professionals managing information systems security and provides reports that can be used to implement a plan for continuous improvement of the systems security level.

Continuous System Patching/Upgrade

Security posture is enhanced with continuous system patching and upgrading of systems and application software. New patches, updates, and service packs for the equipment operating system are necessary to maintain adequate security levels and support new versions of installed products. Keeping abreast of all the changes to all the software and hardware requires a committed effort to stay informed and to communicate gaps in security that can appear in installed systems and applications.

Intervention, Forensics, and Help Desk Services

Quick intervention when a threat is detected is crucial to mitigating the effects of a threat. This requires security engineers with ample knowledge in the various technologies and with the ability to support applications as well as infrastructures on a 24/7 basis. MaaS platforms routinely provide this service to their customers. When a detected threat is analyzed, it often requires forensic analysis to determine what it is, how much effort it will take to fix the problem, and what effects are likely to be seen. When problems are encountered, the first thing customers tend to do is pick up the phone. Help desk services provide assistance on questions or issues about the operation of running systems. This service includes assistance in writing failure reports, managing operating problems, etc.

Real-Time Log Monitoring Enables Compliance

Security monitoring services can also help customers comply with industry regulations by automating the collection and reporting of specific events of interest, such as log-in failures. Regulations and industry guidelines often require log monitoring of critical servers to ensure the integrity of confidential data. MaaS provider's security monitoring services automate this time consuming process.

Platform as a Service

Platform as a service (PaaS) is a category of cloud computing services that provide a computing platform and a solution stack as a service. Along with software as a service (SaaS) and infrastructure as a service (IaaS), it is a service model of cloud computing. PaaS offerings facilitate the deployment of applications without the cost and complexity of buying and managing the underlying hardware and software and provisioning hosting capabilities.

Technically, a PaaS is an Application Platform comprised of an operating system, middleware and other software that allows applications to run on the cloud with much of the management, security, scaling and other stack related headaches abstracted away. This allows you to focus on two things: customers and developing your application. Let the PaaS deal with system administration details like setting up servers or VMs, installing libraries or frameworks, configuring testing tools, etc.

Platform as a Service allows users to create software applications using tools supplied by the provider. PaaS services can consist of preconfigured features that customers can subscribe to; they can choose to include the features that meet their requirements while discarding those that do not. PaaS works on top of IaaS and will do all of that work automatically.

The Traditional On-Premises Model

The traditional approach to developing and running applications on-premises has always been complex, expensive and risky. Producing your own solution never brought any guarantee of success. Each application has been designed to meet their specific requirements within each business. Each solution requires a specific programming hardware, an operating system, a database, often a middle-ware package, mail and web servers, etc. Once environment was created in hardware and software, a team of developers had to navigate a complex programming platform to build their own applications. Additionally, a team of network, database and system management was needed to keep everything in perfect driving conditions. Inevitably, developers were forced to change the application on behalf of a detail of the business, generating new cycles of testing before being distributed.

PaaS model offers a choice of faster and more cost-effective application development and delivery. Furthermore, PaaS provides all the infrastructure needed to run applications over the Internet. Just like Google, iTunes and Youtube, this cloud computing model allows new functionality to be delivered in emerging markets through web browsers. PaaS is based on a model of mediation or subscription, and users only pay for what they use. The PaaS model in its range also includes other facilities such as, application design and development, testing, deployment and hosting as well as integration, security, scalability, storage, status management, control panel, etc.

Characteristics of PaaS

1. Multi-tenant architecture

A PaaS offering must be multi-tenanted. A multi-tenant platform is one that uses common computing resources including hardware, operating system, software (i.e. application code), and a single underlying database with a shared schema to support multiple customers simultaneously. This is in direct contrast to the traditional client/server architecture, which requires an entire stack of hardware and software to be dedicated to each tenant (customer).

2. Customizable /Programmable User Interface

PaaS offering should provide the ability to construct highly flexible user interfaces via a simple “drag & drop” methodology that permits the creation and configuration of UI components on the fly. Furthermore, given the growing set of Web devices, additional flexibility to use other technologies such as CSS, AJAX and Adobe Flex to specify the appearance of the application’s interface should be available to the UI designer.

3. Unlimited Database Customizations

Database used by application should have option of customization for more flexibility in application development. Specifying relationships between objects, a key requirement of any sophisticated business application, must be possible through the declarative Web-based interface. Other mandatory functions include the ability to incorporate validation rules and permissions at the object/field level and the ability to specify auditing behavior.

4. Automation

PaaS environments automate the process of deploying applications to infrastructure, configuring application components, provisioning and configuring supporting technology like load balancers and databases, and managing system change based on policies set by the user. While IaaS is known for its ability to shift capital costs to operational costs through outsourcing, only PaaS is able to cut down costs across the development, deployment and management aspects of the application lifecycle.

5. Security

The PaaS offering should provide a flexible access control system that allows detailed control over what users of the SaaS application can see and the data each user can access. Definition of access from the application level (including tabs, menus, objects, views, charts, reports and workflow actions) to the individual field level should be possible. Defining an access control model should be possible through the creation of groups and roles and the assignment of users to either groups or roles.

6. Runtime Framework:

This is the “software stack” aspect of PaaS, and perhaps the aspect that comes first to mind for most people. The PaaS runtime framework executes end-user code according to policies set by the application owner and cloud provider. PaaS runtime frameworks come in many flavors, some based on traditional application runtimes, others based on 4GL and visual programming concepts, and some with pluggable support for multiple application runtimes.

Software as a Service

Software as a service , sometimes referred to as "on-demand software", is a software delivery model in which software and associated data are centrally hosted on the cloud. SaaS is typically accessed by users using a thin client via a web browser.

In this model, the software is not hosted on the customers' individual computers. Under the SaaS model, a vendor is responsible for the creation, updating, and maintenance of software. Customers buy a subscription to access it, which includes a separate license, or seat, for each person that will use the software.

SaaS has become a common delivery model for many business applications, including accounting, collaboration, customer relationship management (CRM), management information systems(MIS), enterprise resource planning (ERP), invoicing, human resource management (HRM), content management (CM) and service desk management.

The emergence of SaaS as an effective software-delivery mechanism creates an opportunity for IT departments to change their focus from deploying and supporting applications to managing the services that those applications provide.

Unlike traditional software which is conventionally sold as a perpetual license with an up-front cost (and an optional ongoing support fee), SaaS providers generally price applications using a subscription fee, most commonly a monthly fee or an annual fee. Consequently, the initial setup cost for SaaS is typically lower than the equivalent enterprise software. SaaS vendors typically price their applications based on some usage parameters, such as the number of users ("seats") using the application. However, because in a SaaS environment customers' data reside with the SaaS vendor, opportunities also exist to charge per transaction, event, or other unit of value.

Benefits of SaaS (Hope now you can explain all):

- Save money by not having to purchase servers or other software to support use.
- Focus Budgets on competitive advantage rather than infrastructure.
- Monthly obligation rather than up front capital cost.
- Reduced need to predict scale of demand and infrastructure investment up front as available capacity matches demand.
- Multi-Tenant efficiency
- Flexibility and scalability
- Security

Characteristics of SaaS

- **Simple and quick system implementation:** As SaaS sits on top of PaaS and IaaS, deploying any enterprise level software becomes easy and quick as SaaS inherits all the features of underlying infrastructure. Also since SaaS is scalable, any user request can be addressed with required elasticity.
- **Transparent and low pricing:** With common infrastructure, cloud vendor can leverage their infrastructure with lower cost and transparency. End result of this directly impacts customer cost of software implementation as they get it cheaper with enhanced security and high availability.
- **Multitenant Architecture:** A multitenant architecture, in which all users and applications share a single, common infrastructure and code base that is centrally maintained. Because SaaS vendor clients are all on the same infrastructure and code base, vendors can innovate more quickly and save the valuable development time previously spent on maintaining numerous versions of outdated code.
- **Easy software maintenance and Customization:** The ability for each user to easily customize applications to fit their business processes without affecting the common infrastructure. Because of the way SaaS is architected, these customizations are unique to each company or user and are always preserved through upgrades. That means SaaS providers can make upgrades more often, with less customer risk and much lower adoption cost.
- **Better Access:** Improved access to data from any networked device while making it easier to manage privileges, monitor data use, and ensure everyone sees the same information at the same time.

CHAPTER 3

Managed Service

A managed service provider (MSP) is a third-party contractor that delivers network-based services, applications and equipment to enterprises, residences or other service providers.

Managed service providers can be hosting companies or access providers that offer IT services such as fully outsourced network management arrangements, including IP telephony, messaging and call center management, virtual private networks (VPNs), managed firewalls and monitoring/reporting of network servers. Most of these services can be performed from outside a company's internal network with a special emphasis placed on integration and certification of Internet security for applications and content. MSPs serve as outsourcing agents for companies, especially other service providers like ISPs, that don't have the resources to constantly upgrade or maintain faster and faster computer networks.

Managed services providers can offer services such as alerts, security, patch management, data backup and recovery for different client devices: desktops, notebooks, servers, storage systems, networks and applications. Offloading routine infrastructure management to an experienced managed services professional lets you concentrate on running your business, with fewer interruptions due to IT issues.

MSPs act as an extension of your IT department, taking care of routine IT infrastructure monitoring and management around the clock—freeing up your IT staff to focus on higher-value projects. By proactively monitoring and maintaining your systems, an MSP can help you avoid many technology problems in the first place. Should an issue occur, an experienced MSP can troubleshoot and resolve it more efficiently.

Evolution from Managed Service to Cloud Computing

Managed service providers are feeling a bit of pressure from cloud computing. After years of investing in infrastructure and business model changes to deliver remote managed services, MSPs are seeing the cloud as a threat to their livelihood. In many cases, the cloud requires no remote infrastructure or on-premise equipment. The fear is the cloud could render MSPs obsolete.

Cloud computing could also be a managed service. Cloud computing consists of security monitoring, storage management, network administration and the all managed services need not necessarily be cloud computing. Cloud computing puts its efforts in creating a technical solution. It is a technical model that delivers technical access to the computing resources. Thus, cloud computing can be defined as a technical solution.

Managed services, on the other hand, is a contract based relationship. The definition and delivery of the service is done on a repeating revenue basis. It means that managed services are recurring revenues from well-defined services that are predictable.

Some types of services that managed service provider's offer are help desk assistance and network administration services. Managed services allow the Information Technology staff of the company to

focus their efforts and energy in to the core activities of the company rather than dealing with the IT challenges. A managed service is the management of technology like telephony, IT, applications and others. However, the definition of a managed service is changing.

There are, however a lot of similarities that exist between a managed service and cloud computing. While cloud computing is generally a technical implementation that decides how the infrastructure and applications are to be delivered over the private and public networks, it can also be a business model. The same contract agreement holds ground in a cloud computing relationship. Managed services can also be a technical implementation.

Cloud Optimized Infrastructure is based around 5 key belief that offer capabilities that differentiate as an MSP from cloud computing:

- Flexible licensing which supports the elastic expansion/contraction of cloud-based services and accommodates the billing implications
- Low overhead in deployment and use across server, network, and storage resources, making it a great fit for virtual machine environments that are a key supporting technology in cloud-based computing
- Non-disruptive scalability which accommodates the need for server, storage and other infrastructure growth on both the end user and cloud provider sides without impacting client-side production servers
- Enterprise multi tenancy that provides for the secure delivery of reliable services to multiple customers with a scalable management model
- Broad heterogeneous support that maximizes cloud provider market opportunities by covering a wide range of server, storage, and application environments found in customer settings

Single Purpose Architecture to Multipurpose Architecture

In the early days of MSPs, the providers would actually go onto customer sites and perform their services on customer-owned premises. Over time, these MSPs specialized in implementation of infrastructure and quickly figured out ways to build out data centers and sell those capabilities off in small chunks commonly known as monthly recurring services, in addition to the basic fees charged.

With virtualization and other supporting technologies, MSP's quickly convinced their customers to shift their data centers to multipurpose and multitenant architecture. Not only multipurpose architecture was effective for their customers but also was huge cost saving initiative.

Data Center Virtualization

Data center virtualization is a method of moving information storage from physical servers to virtual ones, often in a different location. In the past, large companies would keep physical servers on site that held huge amounts of corporate information. These servers were expensive, both to purchase and maintain. With data center virtualization, it became possible to separate both the hardware and location from the data. This cuts costs and increases the data's availability.

Data center virtualization actually comes from a combination of two different technologies; high-speed data transfer and server virtualization. Without both of these components, data center virtualization becomes highly impractical.

There are three areas of IT where virtualization is making inroads, network virtualization, storage virtualization and server virtualization:

- **Network virtualization** is a method of combining the available resources in a network by splitting up the available bandwidth into channels, each of which is independent from the others, and each of which can be assigned (or reassigned) to a particular server or device in real time. The idea is that virtualization disguises the true complexity of the network by separating it into manageable parts, much like your partitioned hard drive makes it easier to manage your files.
- **Storage virtualization** is the pooling of physical storage from multiple network storage devices into what appears to be a single storage device that is managed from a central console. Storage virtualization is commonly used in storage area networks (SANs).
- **Server virtualization** is the masking of server resources, including the number and identity of individual physical servers, processors, and operating systems, from server users. The server administrator uses a software application to divide one physical server into multiple isolated virtual environments. The virtual environments are sometimes called virtual private servers, but they are also known as guests, instances, containers or emulations.

Benefits of Data Center Virtualization

In addition to the cost savings that result from reducing the number of servers, the following benefits are typically realized as well:

- Enables the consolidation of physical servers, slashing the costs of operating a data center. This includes reducing the costs of server upgrades, management, power, space, and storage.

- Reduction in data center space and in data center equipment such as PDUs, air conditioning units, etc.
- Reduction in the number of network, HBAs and SAN switches.
- Provides true high-availability for all servers without requiring duplicate hardware and clustering software.
- Integrates the test/development and production environments while significantly enhancing the test/development process.
- Facilitates true disaster recovery for all servers.
- Eliminates the need for maintenance windows for physical server troubleshooting or upgrades and enables faster server provisioning.
- Enhances security and provides regulatory compliance benefits.

Cloud Data Center

A cloud data center has three distinct characteristics that differentiate it from traditional DC. It is sold on demand, typically by the minute or the hour; it is elastic -- a user can have as much or as little of a service as they want at any given time; and the service is fully managed by the provider (the consumer needs nothing but a personal computer and Internet access). Significant innovations in virtualization and distributed computing, as well as improved access to high-speed Internet and a weak economy, have accelerated interest in cloud computing.

The multiplier effect of Internet data, trends of enterprises information transformation and tremendous load bring forth traditional DC huge challenges: how to reduce operation and maintenance cost, how to meet demand of high capacity, high security and high efficiency. Attributing to resource on-demand, flexible and dynamic structure, cloud computing is the right technology to resolve all these issues. By continuously improving core technologies including virtualization, elastic computing and high-density computing etc, cloud vendors have creatively developed Modular cloud computing Data Center of trusty, efficient, smart, ultra-bandwidth and green end to end solution.

Cloud Data Center is now evolving beyond being merely a model of technology delivery to becoming a new operating model where business decision makers are empowered to procure infrastructure on demand, and where IT becomes an internal service provider delivering increased business agility without compromising security or control.

Ultimately, cloud services are attractive because the cost is likely to be far lower than providing the same service from your traditional data center.

Traditional Corporate Data Center	Cloud Data Center
Thousands of different applications	Few applications (maybe even just one)
Mixed hardware environment	Homogeneous hardware environment
Multiple management tools	Standardized management tools
Frequent application patching and updating	Minimal application patching and updating
Complex workloads	Simple workloads
Multiple software architectures	Single standard software architecture

Service Oriented Architecture


If a service-oriented architecture is to be effective, we need a clear understanding of the term **service**. A service is a function that is well-defined, self-contained, and does not depend on the context or state of other services.

A service-oriented architecture is essentially a collection of services. These services communicate with each other. The communication can involve either simple data passing or it could involve two or more services coordinating some activity. SOA defines how to integrate widely disparate applications for a Web-based environment and uses multiple implementation platforms. **Rather than defining an API**, SOA defines the interface in terms of protocols and functionality.

Why SOA?

The reality in IT enterprises is that infrastructure is heterogeneous across operating systems, applications, system software, and application infrastructure. Some existing applications are used to run current business processes, so starting from scratch to build new infrastructure isn't an option.

Enterprises should quickly respond to business changes with agility; leverage existing investments in applications and application infrastructure to address newer business requirements; support new channels of interactions with customers, partners, and suppliers; and feature an architecture that supports organic business. SOA with its loosely coupled nature allows enterprises to plug in new services or upgrade existing services in a granular fashion to address the new business requirements, provides the option to make the services consumable across different channels, and exposes the existing



enterprise and legacy applications as services, thereby safeguarding existing IT infrastructure investments.

For example, a core banking application provides a Fund Transfer service, then the other banking applications such as Treasury, Payment Gateway, ATM Switching, and so on can call or invoke Fund Transfer service without need to worry about where the Fund Transfer is located in the network. This contrasts with the Tight Coupling approach. Each application defining their own Fund Transfer, the problem come when the Transfer Fund logic is change. It will be difficult and require high cost (and time, of course) to set the new logic into each application.

Service-oriented architectures are not new. The first service-oriented architectures are usually considered to be the Distributed Component Object Model (DCOM) or Object Request Brokers (ORBs), which were based on the Common Object Requesting Broker Architecture (CORBA) specification. The introduction of SOA provides a platform for technology and business units to meet business requirements of the modern enterprise. With SOA, your organization can use existing application systems to a greater extent and may respond faster to change requests. These benefits are attributed to several critical elements of SOA:

1. Free-standing, independent components
2. Combined by loose coupling
3. Message (XML)-based instead of API-based
4. Physical location, etc., not important

Characteristics of SOA

- In SOA, Services should be **independent** of other services. Altering a service should not affect calling service.
- Services should be **self-contained**. When we talk about a Register Customer service it means, service will do all the necessary work for us, we are not required to care about anything.
- Services should be able to **define themselves**. Services should be able to answer a question what it does? It should be able to tell client what all operations it does, what all data types it uses and what kind of responses it will return.
- Services should be **published** into a location (directory) where anyone can search for it.
- As said, SOA comprises of collection services which communicate via **standard Messages**. Standard messages make them platform independent. (Here standard doesn't mean standard across Microsoft it means across all programming languages and technologies.)
- Services should be able to communicate with each other **asynchronously**.
- Services should support **reliable messaging**. Means there should be a guarantee that request will be reached to correct destination and correct response will be obtained.
- Services should support **secure communication**.

About Open Source Software

Open-source software (OSS) is computer software with its source code made available and licensed with an open-source license in which the copyright holder provides the rights to study, change and distribute the software for free to anyone and for any purpose. Open-source software is very often developed in a public, collaborative manner.

The basics behind the Open Source Initiative is that when programmers can read, redistribute and modify the source code for a piece of software, the software evolves. Open source sprouted in the technological community as a response to proprietary software owned by corporations.

Proprietary software is privately owned and controlled. In the computer industry, proprietary is considered the opposite of open. A proprietary design or technique is one that is owned by a company. It also implies that the company has not divulged specifications that would allow other companies to duplicate the product.

Open Source is a certification standard issued by the Open Source Initiative (OSI) that indicates that the source code of a computer program is made available free of charge to the general public. OSI dictates that in order to be considered "OSI Certified" a product must meet the following criteria:

Some Points to Know about Open Source Software

- The author or holder of the license of the source code cannot collect royalties on the distribution of the program.
- The distributed program must make the source code accessible to the user.
- The author must allow modifications and derivations of the work under the program's original name.
- No person, group or field of endeavor can be denied access to the program.
- The rights attached to the program must not depend on the program's being part of a particular software distribution.
- The licensed software cannot place restrictions on other software that is distributed with it

Linux, Apache and other open-source applications have long been used to power Web and file servers. But when it comes to managing the data center, many companies have held back. Now, though, some users have turned into big believers that open source works here, too.

The following open source packages take a more holistic approach by integrating all of the necessary functionality into a single package (including virtualization, management, interfaces, and security). When added to a network of servers and storage, these packages produce flexible cloud computing and storage infrastructures (IaaS).

Eucalyptus

One of the most popular open source packages for building cloud computing infrastructures is Eucalyptus (for Elastic Utility Computing Architecture for Linking Your Programs to Useful Systems). What makes it unique is that its interface is compatible with Amazon Elastic Compute Cloud (Amazon EC2—Amazon's cloud computing interface). Additionally, Eucalyptus includes Walrus, which is a cloud storage application compatible with Amazon Simple Storage Service (Amazon S3—Amazon's cloud storage interface).

Eucalyptus supports KVM/Linux and Xen for hypervisors and includes the Rocks cluster distribution for cluster management.

OpenNebula

OpenNebula is another interesting open source application (under the Apache license) developed at the Universidad Complutense de Madrid. In addition to supporting private cloud construction, OpenNebula supports the idea of hybrid clouds. A hybrid cloud permits combining a private cloud infrastructure with a public cloud infrastructure (such as Amazon) to enable even higher degrees of scaling.

OpenNebula supports Xen, KVM/Linux, and VMware and relies on elements like libvirt for management and introspection.

Nimbus

Nimbus is another IaaS solution focused on scientific computing. With Nimbus, you can lease remote resources (such as those provided by Amazon EC2) and manage them locally (configure, deploy VMs, monitor, etc.). Nimbus morphed from the Workspace Service project (part of Globus.org). Being dependent on Amazon EC2, Nimbus supports Xen and KVM/Linux.

Xen Cloud Platform

Citrix has integrated Xen into an IaaS platform, using Xen as the hypervisor while incorporating other open source capabilities such as the Open vSwitch. An interesting advantage to the Xen solution is the

focus on standards-based management (including OVF, Distributed Management Task Force [DTMF], the Common Information Model [CIM], and Virtualization Management Initiative [VMAN]) from the project Kensho. The Xen management stack supports SLA guarantees, along with detailed metrics for charge-back.

OpenQRM

Our penultimate solution is OpenQRM, which is categorized as a data center management platform. OpenQRM provides a single console to manage an entire virtualized data center that is architecturally pluggable to permit integration of third-party tools. OpenQRM integrates support for high availability (through redundancy) and supports a variety of hypervisors, including KVM/Linux, Xen, VMware, and Linux VServer.

OpenStack

Today, the leading IaaS solution is called OpenStack. OpenStack was released in July 2010, and has quickly become the standard open-source IaaS solution. OpenStack is a combination of two cloud initiatives from Rackspace Hosting (Cloud Files) and NASA's Nebula platform. OpenStack is being developed in the Python language, and is under active development under the Apache license.

Apache

The Apache HTTP Server, commonly referred to as Apache, is a web server software program notable for playing a key role in the initial growth of the World Wide Web. In 2009 it became the first web server software to surpass the 100 million website milestone. Apache was the first viable alternative to the Netscape Communications Corporation web server (currently named Oracle iPlanet Web Server). Typically Apache is run on a Unix-like operating system and was developed for use on Linux. The Apache HTTP Server Project is a collaborative software development effort aimed at creating a robust, commercial-grade, feature-rich and freely-available source code implementation of an HTTP (Web) server.

Advantages of OSS

Open-source software is free to use, distribute, and modify. It has lower costs, and in most cases this is only a fraction of the cost of their proprietary counterparts.

Open-source software is more secured as the code is accessible to everyone. Anyone can fix bugs as they are found, and users do not have to wait for the next release. The fact that it is continuously analyzed by a large community produces secure and stable code.

Open source is not dependent on the company or author that originally created it. Even if the company fails, the code continues to exist and be developed by its users. Also, it uses open standards accessible to everyone; thus, it does not have the problem of incompatible formats that exist in proprietary software.

Lastly, the companies using open-source software do not have to think about complex licensing models and do not need anti-piracy measures like product activation or serial number.

Disadvantages of OSS

The main disadvantage of open-source software is not being straightforward to use. Open-source operating systems like Linux cannot be learned in a day. They require effort and possibly training from your side before you are able to master them. You may need to hire a trained person to make things easier, but this will incur additional costs.

There is a shortage of applications that run both on open source and proprietary software; therefore, switching to an open-source platform involves a compatibility analysis of all the other software used that run on proprietary platforms. In addition, there are many ongoing parallel developments on open source software. This creates confusion on what functionalities are present in which versions.

Lastly, many of the latest hardware are incompatible to the open-source platform; so you have to rely on third-party drivers.

Unit 4

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. This is particularly true for the SaaS provider.

With the cloud model, you lose control over physical security. In a public cloud, you are sharing computing resources with other companies. In a shared pool outside the enterprise, you don't have any knowledge or control of where the resources run. Simply because you share the environment in the cloud, may put your data at risk.

Storage services and its security provided by one cloud vendor may be incompatible with another vendor's services should you decide to move from one to the other. Vendors are known for creating what the hosting world calls "sticky services"—services that an end user may have difficulty transporting from one cloud vendor to another (e.g., Amazon's "Simple Storage Service" [S3] is incompatible with IBM's Blue Cloud, or Google, or Dell). If information is encrypted while passing through the cloud, who controls the encryption/decryption keys? Is it the customer or the cloud vendor? Most customers probably want their data encrypted both ways across the Internet using SSL (Secure Sockets Layer protocol). They also most likely want their data encrypted while it is at rest in the cloud vendor's storage pool. Be sure that you, the customer, control the encryption/decryption keys, just as if the data were still resident on your own servers.

Data integrity means ensuring that data is identically maintained during any operation (such as transfer, storage, or retrieval). Put simply, data integrity is assurance that the data is consistent and correct. Ensuring the integrity of the data really means that it changes only in response to authorized transactions.

Having proper **fail-over technology** is a component of securing the cloud that is often overlooked. The company can survive if a non-mission critical application goes offline, but this may not be true for mission-critical applications. Core business practices provide competitive differentiation. **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. Many compliance regulations require that data not be intermixed with other data, such as on shared servers or databases. Some countries have strict limits on what data about its citizens can be stored and for how long, and some banking regulators require that customers' financial data remain in their home country.

Government policy will need to change in response to both the opportunity and the threats that cloud computing brings. This will likely focus on the off-shoring of personal data and protection of privacy, whether it is data being controlled by a third party or off-shored to another country. There will be a corresponding drop in security as the traditional controls such as VLANs (virtual local-area networks) and firewalls prove less effective during the transition to a virtualized environment. 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 isn't 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. Placing large amounts of sensitive data in a globally accessible cloud leaves organizations open to large distributed threats—attackers no longer have to come onto the premises to steal data, and they can find it all in the one “virtual” location.

Virtualization efficiencies in the cloud require virtual machines from multiple organizations to be co-located on the same physical resources. 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.

The dynamic and fluid nature of virtual machines will make it difficult to maintain the consistency of security and ensure the auditability of records. The ease of cloning and distribution between physical servers could result in the propagation of configuration errors and other vulnerabilities. Proving the security state of a system and identifying the location of an insecure virtual machine will be challenging. Regardless of the location of the virtual machine within the virtual environment, the intrusion detection and prevention systems will need to be able to detect malicious activity at virtual machine level. 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 vendor's. The need for patch maintenance vigilance is imperative. Lack of due diligence in this regard could rapidly make the task unmanageable or impossible, leaving you with "virtual patching" as the only alternative.

Software-as-a-service (SaaS) security issues

Cloud computing models of the future will likely combine the use of SaaS (and other as a service as appropriate), utility computing, and Web 2.0 collaboration technologies to leverage the Internet to satisfy their customers' 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. SaaS will likely remain the dominant cloud service model for the foreseeable future and the area where the most critical need for security practices and oversight will reside. Just as with an managed service provider, corporations or end users will need to research vendors' policies on data security before using vendor services to avoid losing or not being able to access their data. Seven security issues which one should discuss with a cloud-computing vendor:

1. Privileged user access—Inquire about who has specialized access to data, and about the hiring and management of such administrators.

2. Regulatory compliance—Make sure that the vendor is willing to undergo external audits and/or security certifications.

3. Data location—Does the provider allow for any control over the location of data?

4. Data segregation—Make sure that encryption is available at all stages, and that these encryption schemes were designed and tested by experienced professionals.

5. Recovery—Find out what will happen to data in the case of a disaster. Do they offer complete restoration? If so, how long would that take?

6. Investigative support—Does the vendor have the ability to investigate any inappropriate or illegal activity?

7. Long-term viability—What will happen to data if the company goes out of business? How will data be returned, and in what format?

To address the security issues listed above along with others mentioned earlier in the topic, 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.

The baseline security practices for the SaaS environment as currently formulated are discussed in the following sections

Security Management

Lack of clearly defined roles and responsibilities, and agreement on expectations, can result in a general feeling of loss and confusion among the security team about what is expected of them, how their skills and experienced can be leveraged, and meeting their performance goals. Morale among the team and pride in the team is lowered, and security suffers as a result.

Risk Management

Effective risk management entails identification of technology assets; identification of data and its links to business processes, applications, and data stores; and assignment of ownership and custodial responsibilities. Actions should also include maintaining a repository of information assets. Owners have authority and accountability for information assets including protection requirements, and custodians implement confidentiality, integrity, availability, and privacy controls. A formal risk assessment process should be created that allocates security resources linked to business continuity.

Risk/ Vulnerability Assessment

Security risk assessment is critical to helping the information security organization make informed decisions when balancing the dueling priorities of business utility and protection of assets. Lack of attention to completing formalized risk assessments can contribute to an increase in information security audit findings, can jeopardize certification goals, and can lead to inefficient and ineffective selection of security controls that may not adequately mitigate information security risks to an acceptable level. A formal information security risk management process should proactively assess information security risks as well as plan and manage them on a periodic or as-needed basis.

Security Monitoring and Incident Response

Centralized security information management systems should be used to provide notification of security vulnerabilities and to monitor systems continuously through automated technologies to identify potential issues. They should be integrated with network and other systems monitoring processes (e.g., security information management, security event management, security information and event management, and security operations centers that use these systems for dedicated 24/7/365 monitoring).

Management of periodic, independent third-party security testing should also be included. Many of the security threats and issues in SaaS center around application and data layers, so the types and sophistication of threats and attacks for a SaaS organization require a different approach to security monitoring than traditional infrastructure and perimeter monitoring. The organization may thus need to expand its security monitoring capabilities to include application- and data-level activities. This may also require subject-matter experts in applications security and the unique aspects of maintaining privacy in the cloud. Without this capability and expertise, a company may be unable to detect and prevent security threat and attacks to its customer data and service stability.

Incident response is an organized approach to addressing and managing the aftermath of a security breach or attack (also known as an incident). The goal is to handle the situation in a way that limits damage and reduces recovery time and costs. An incident response plan includes a policy that defines, in specific terms, what constitutes an incident and provides a step-by-step process that should be followed when an incident occurs. An organization's incident response is conducted by the computer incident response team, a carefully selected group that, in addition to security and general IT staff, may include representatives from legal, human resources, and public relations departments.

Security Architecture Design

Security Architecture is one component of a products/systems overall architecture and is developed to provide guidance during the design of the product/system.

A security architecture framework should be established with consideration of processes (enterprise authentication and authorization, access control, confidentiality, integrity, non-repudiation, security management, etc.), operational procedures, technology specifications, people and organizational management, and security program compliance and reporting.

A security architecture document should be developed that defines security and privacy principles to meet business objectives. Documentation is required for management controls and metrics specific to asset classification and control, physical security, system access controls, network and computer management, application development and maintenance, business continuity, and compliance.

The creation of a secure architecture provides the engineers, data center operations personnel, and network operations personnel a common blueprint to design, build, and test the security of the applications and systems. Design reviews of new changes can be better assessed against this architecture to assure that they conform to the principles described in the architecture, allowing for more consistent and effective design reviews.

Vulnerability Assessment

Vulnerability assessment classifies network assets to more efficiently prioritize vulnerability-mitigation programs, such as patching and system upgrading. It measures the effectiveness of risk mitigation by setting goals of reduced vulnerability exposure and faster mitigation. Vulnerability management should be integrated with discovery, patch management, and upgrade management processes to close vulnerabilities before they can be exploited.

A vulnerability assessment attempts to identify the exposed vulnerabilities of a specific host, or possibly an entire network. The vulnerabilities may be due to configuration problems or missing software patches.

Vulnerability Assessment in cloud should be done in periodic basis with predefined service level agreement. Customers should be allowed to test cloud infrastructure before and after they outsource their infrastructure to cloud.

Data Privacy and Security

Cloud computing has transformed the way organizations approach IT, enabling them to become more agile, introduce new business models, provide more services, and reduce IT costs. Cloud computing technologies can be implemented in a wide variety of architectures, under different service and deployment models, and can coexist with other technologies and software design approaches.

Maintaining control over the data is paramount to cloud success. A decade ago, enterprise data typically resided in the organization's physical infrastructure, on its own servers in the enterprise's data center, where one could segregate sensitive data in individual physical servers. Today, with virtualization and the cloud, data may be under the organization's logical control, but physically reside in infrastructure owned and managed by another entity.

This shift in control is the number one reason new approaches and techniques are required to ensure organizations can maintain data security. When an outside party owns, controls, and manages infrastructure and computational resources, how can you be assured that business or regulatory data remains private and secure, and that your organization is protected from damaging data breaches—and feel you can still completely satisfy the full range of reporting, compliance, and regulatory requirements?

Some of the points to keep data private and secure in cloud infrastructure are as below:

1. Avoid storing sensitive information in the cloud.
2. Read the user agreement to find out how your cloud service storage works.
3. Password sensitivity
4. Encrypt your data
5. Use Encrypted cloud services

Application Security

Application security is one of the critical success factors for SaaS company. This is where the security features and requirements are defined and application security test results are reviewed. Application security processes, secure coding guidelines, training, and testing scripts and tools are typically a collaborative effort between the security and the development team.

Although product engineers will likely focus on the application layer, the security design of the application itself, and the infrastructure layers interacting with the application, the security team should provide the security requirements for the product development engineers to implement. This should be a collaborative effort between the security and product development team. External penetration testers are used for application source code reviews, and attack and penetration tests provide an objective review of the security of the application as well as assurance to customers that attack and penetration tests are performed regularly. Fragmented and undefined collaboration on application security can result in lower-quality design, coding efforts, and testing results.

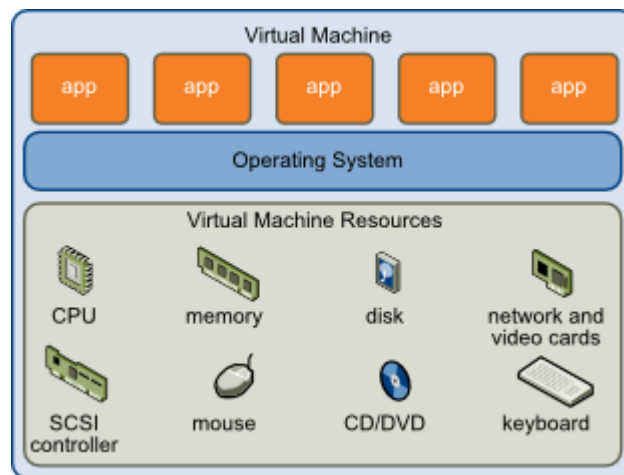
Some of the things that we should consider while moving to cloud application are:

- a. Risks associated with cloud application
- b. The fact that someone is managing and controlling your critical application
- c. The perimeter of cloud is different and multitenant
- d. Application should be protected with industry standard firewall and security products
- e. Insecure Interfaces and Application Program Interface (API's)
- f. Denial of Service (DOS) attack

Virtual Machine Security

Virtual machines are the containers in which applications and guest operating systems run. By design, all VMware virtual machines are isolated from one another. This isolation enables multiple virtual machines to run securely while sharing hardware and ensures both their ability to access hardware and their uninterrupted performance.

Each virtual machine is isolated from other virtual machines running on the same hardware. Although virtual machines share physical resources such as CPU, memory, and I/O devices, a guest operating system on an individual virtual machine cannot detect any device other than the virtual devices made available to it,



Virtual Machine Isolation

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 and 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, you 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.

A further area of concern with virtualization has to do with the potential for undetected network attacks between VMs collocated on a physical server. Unless you can monitor the traffic from each VM, you can't verify that traffic isn't possible between those VMs. In essence, network virtualization must deliver an appropriate network interface to the VM. That interface might be a multiplexed channel with all the switching and routing handled in the network interconnect hardware.

Disaster Recovery

A Disaster Recovery Plan (DRP) is a business plan that describes how work can be resumed quickly and effectively after a disaster. Disaster recovery planning is just part of business continuity planning and applied to aspects of an organization that rely on an IT infrastructure to function.

The overall idea is to develop a plan that will allow the IT department to recover enough data and system functionality to allow a business or organization to operate - even possibly at a minimal level.

A **disaster recovery plan (DRP)** documents policies, procedures and actions to limit the disruption to an organization in the wake of a disaster. Just as a disaster is an event that makes the continuation of normal functions impossible, a disaster recovery plan consists of actions intended to minimize the negative effects of a disaster and allow the organization to maintain or quickly resume mission-critical functions.

To better understand and evaluate disaster recovery strategies, it is important to define two terms: recovery time objective (RTO) and recovery point objective (RPO).

RTO

The recovery time objective (RTO) is the maximum amount of time allocated for restoring application functionality. This is based on business requirements and is related to the importance of the application. Critical business applications require a low RTO.

RPO

The recovery point objective (RPO) is the acceptable time window of lost data due to the recovery process. For example, if the RPO is one hour, you must completely back up or replicate the data at least every hour. Once you bring up the application in an alternate datacenter, the backup data may be missing up to an hour of data. Like RTO, critical applications target a much smaller RPO.

Some of the points why Disaster Recovery is needed?

- a. Machines, hardware and even data centers fail.
- b. Much like machines, humans are not perfect. They make mistakes. In case of mistakes, DR may help resume business from back date.
- c. Customers expect perfection as they don't want disruption in services
- d. DR enabled organizations will attract more customers.

Disaster Recovery Management/ Planning Steps

- **Count the costs.** Although data center downtime is harmful to any company that relies on its IT services, it costs some companies more than others. Your disaster recovery plan should enable a fast return to service, but it shouldn't cost you more than you are losing in downtime costs.
- **Evaluate the types of threats you face and how extensively they can affect your facility.** Malicious attacks can occur anywhere, but you may also face threats peculiar to your location, such as weather events (tornadoes, hurricanes, floods and so on), earthquakes or other dangers. Part of preparing for a disaster is to know what is likely to occur and how those threats could affect your systems. Evaluating these situations beforehand allows you to better take appropriate action should one of these events occur.
- **Know what you have and how critical it is to operations.** Responding to a disaster in your data center is similar to doing so in medicine: you need to treat the more serious problems first, then the more minor ones. By determining which systems are most critical to your data center, you enable your IT staff to prioritize and make the best use of the precious minutes and hours immediately following an outage. Not every system need be functional immediately following a disaster.
- **Identify critical personnel and gather their contact information.** Who do you most want to be present in the data center following an outage? Who has the most expertise in a given area and the greatest ability to oversee some part of the recovery effort? Being able to get in touch with these people is crucial to a fast recovery. Collect their contact information and, just as importantly, keep it up to date. If it's been a year or more since you last checked, some of that contact information is likely out of date. Every minute you spend trying to find important personnel is time not spent on recovery.
- **Train your employees.** Knowledge of how to implement disaster recovery procedures is obviously important when an outage occurs. To this end, prepare by training personnel—

and not just in their respective areas of expertise. Everyone should have some broad-based knowledge of the recovery process so that it can be at least started even if not everyone is present.

- **Ensure that everyone knows the disaster recovery plan and understands his or her role.** Announcing the plan and assigning roles is not something you should do after a disaster strikes; it should be done well in advance, leaving time for personnel to learn their roles and to practice them. Almost nothing about a disaster event should be new (aside from some contingencies of the moment, perhaps): the IT staff should implement disaster recovery as a periodic task (almost) like any other.
- **Practice.** Needless to say, this is perhaps the most critical part of preparation for a downtime event. The difference between knowing your role and being able to execute it well is simply practice. You may not be able to shut down your data center to simulate precisely all of the conditions you will face in an outage, but you can go through many of the procedures nevertheless. Some recommendations prescribe semiannual drills, at a minimum, to practice implementing the disaster recovery plan. If there's one thing you take from this article, it's that you should practice your disaster recovery plan—don't expect it to unfold smoothly when you need it (regardless of how well laid-out a plan it is) if you haven't given it a trial run or two.
- **Automate where possible.** Your staff is limited, so it can only do so much. The more that your systems can do on their own in a recovery situation, the faster the recovery will generally be. This also leaves less room for human error—particularly in the kind of stressful atmosphere that exists following a disaster.
- **Follow up after a disaster.** When a downtime event does occur, evaluate the performance of the personnel and the plan to determine if any improvements can be made. Update your plan accordingly to enable a better response in the future. Furthermore, investigate the cause of the outage. If it's an internal problem, take necessary measures to correct equipment issues to avoid the same problem occurring again.



Cloud Cube Model: Selecting Cloud Formations for Secure Collaboration

Problem

Cloud computing offers massive scalability - in virtual computing power, storage, and applications resources - all at almost immediate availability and low cost, and business managers are demanding their IT operations assess the benefits this new computing model can represent. As with all new technologies, there are new risks to be discovered and old risks to be re-evaluated. Many articles have already been published, and several new groups have formed to address cloud computing. In line with our continuing strategy for enabling secure business collaboration, the Jericho Forum members are addressing how to collaborate securely in the clouds.

There are several “cloud formations” - or forms of cloud computing. Each offers different characteristics, varying degrees of flexibility, different collaborative opportunities, and different risks. Thus one of the key challenges that businesses face when considering cloud computing as an option is to determine how to choose the cloud formation best suited to their various types of business operations.

The Jericho Forum’s objectives related to cloud computing are distinctive – enabling secure collaboration in the appropriate cloud formations best suited to the business needs.

With this in mind, the aim of this paper is to:

- point out that not everything is best implemented in clouds; it may be best to operate some business functions using a traditional non-cloud approach
- explain the different cloud formations that the Jericho Forum has identified
- describe key characteristics, benefits and risks of each cloud formation
- provide a framework for exploring in more detail the nature of different cloud formations and the issues that need answering to make them safe and secure places to work in.

This is very much “work-in-progress”, which we hope will enable all stakeholders, but particularly business decision-makers, to appreciate the key considerations that need to be taken into account when deciding which parts of their business could be operated in which of the available cloud formations.

Why should I care?

Cloud computing suppliers claim they are responding to customer demands for assurances on the security of the services they provide. Some even claim that, because they know that their customers place high priority on the security of the data they own, the security of the cloud services they offer is often significantly better than that of the customer’s own IT systems.

While this may well be true, it is critical that cloud customers select the right cloud formations for their needs, to ensure they remain secure, able to collaborate safely with their selected parties as their evolving business needs require, and compliant to applicable regulatory requirements - including on the use and location of their data.

The joy of the cloud model is that it can deliver great advantages, but only if you know where in the different formations of cloud you need to be in order to achieve the right flexibility for your business needs. For example, if a cloud vendor were to cease providing a service, how effortlessly could you move to another provider or use your cloud-based capability to provide you with seamless disaster recovery and business continuity?

The Jericho Forum is actively encouraging solution providers and vendors to develop the missing capabilities and services to ensure customers are protected from the stormier implications of clouds. In Feb 2009, we delivered a practical framework geared to showing how to create the right Collaboration Oriented Architecture¹ (COA) to assure secure business collaboration in de-perimeterised environments. For the Jericho Forum, the natural evolution from this is to address how to follow a well-structured path towards enabling secure business collaboration without becoming vulnerable to issues which may put at risk your data, or your ability to work with your chosen business parties, or your regulatory compliance.

Recommendation / response

Protecting your Data

First, it is necessary to classify your data so as to know what rules must apply to protecting it:

- it's sensitivity - must it only exist at specific trust levels? If so, which?
- What regulatory/compliance restrictions apply – e.g. Must it stay within your national boundary? Does it have to stay in Safe Harbours? etc.

We can only meet this requirement if we have universally adopted standards for:

- a data classification model that is sufficiently easy for all originators of data to use – for example the G8 Traffic Light Protocol².
- an associated standard for managing trust levels
- standardised metadata that signals to “cloud security” what security needs be applied to each item of data.

With an understanding on what security you need to apply to your data, you're in a position to decide:

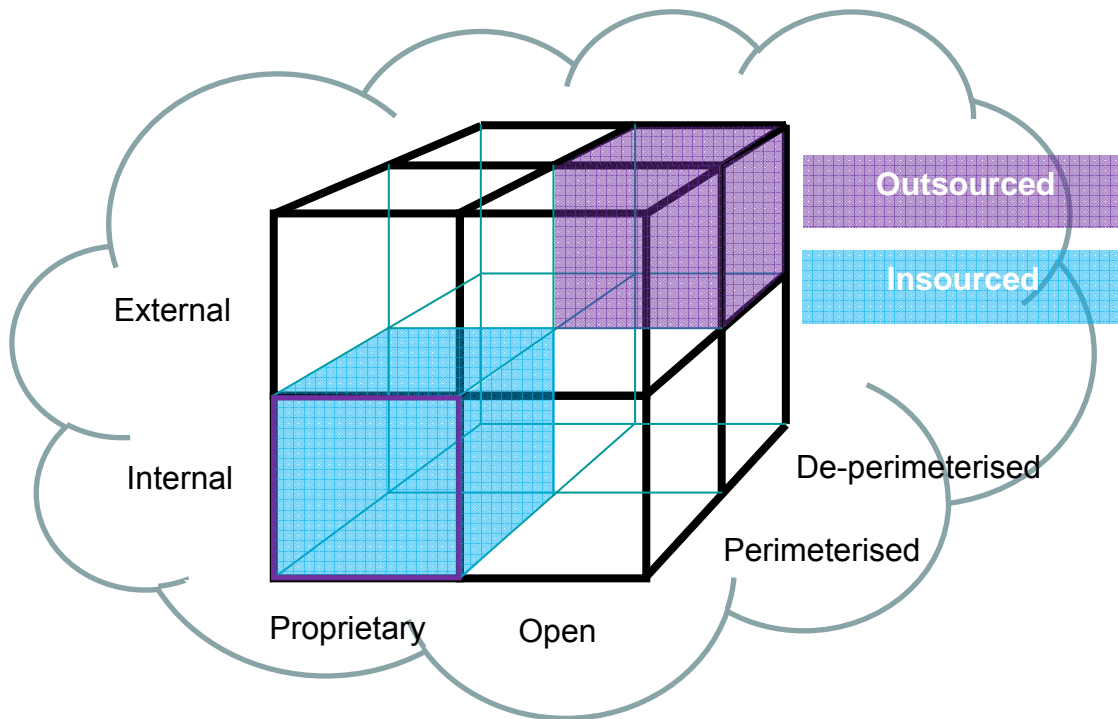
- what data and processes to move to the Clouds
- at what level you want to operate in the Clouds? Cloud models separate layers of business service from each other, for example, Infrastructure / Platform / Software / Process.
- which Cloud Formations are best suited to your needs.

¹ Collaboration Oriented Architectures (COA) Framework – see COA papers freely available from the Jericho Forum Web site at <http://www.opengroup.org/jericho/publications.htm>

² See COA paper on Trust Management – available as free download from <http://www.opengroup.org/jericho/publications.htm>

Cloud Formations – the Cloud Cube Model

The Jericho Forum has identified 4 criteria to differentiate cloud formations from each other and the manner of their provision. The Cloud Cube Model summarises these 4 dimensions, which are explained in turn in the rest of this paper.




The Cloud Cube Model

Cloud Cube Model Dimensions

Dimension: Internal (I) / External (E)

This is the dimension that defines the physical location of the data: where does the cloud form you want to use exist - inside or outside your organization's boundaries.

- If it is within your own physical boundary then it is Internal.
- If it is not within your own physical boundary then it is External.




For example, virtualised hard disks in an organisation's data centre would be internal, while Amazon SC3³ would be external at some location "off-site".

Note: Be wary of making a false assumption that Internal is more secure than External. The effective use of both is likely to provide the most secure usage model.

³ Amazon's SC3 - on-demand storage solution.

Dimension: Proprietary (P) / Open (O)

This is the dimension that defines the state of ownership of the cloud technology, services, interfaces⁴, etc. It indicates the degree of interoperability, as well as enabling “data/application transportability” between your own systems and other cloud forms, and the ability to withdraw your data from a cloud form or to move it to another without constraint. It also indicates any constraints on being able to share applications.

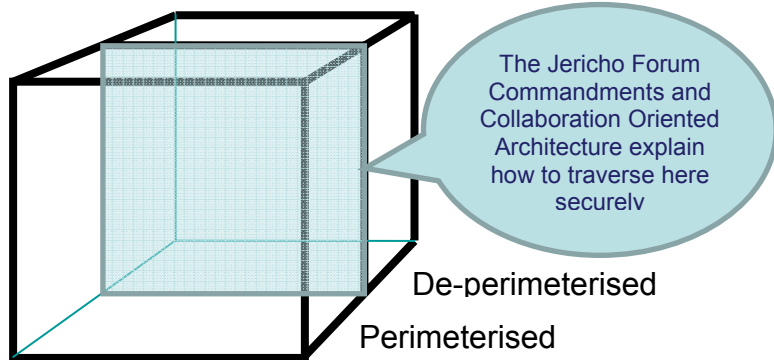


- Proprietary means that the organisation providing the service is keeping the means of provision under their ownership. As a result, when operating in clouds that are proprietary, you may not be able to move to another cloud supplier without significant effort or investment. Often the more innovative technology advances occur in the proprietary domain. As such the proprietor may choose to enforce restrictions through patents and by keeping the technology involved a trade secret.
- Clouds that are Open are using technology that is not proprietary, meaning that there are likely to be more suppliers, and you are not as constrained in being able to share your data and collaborate with selected parties using the same open technology. Open services tend to be those that are widespread and consumerised, and most likely a published open standard, for example email (SMTP).

An as yet unproven premise is that the clouds that most effectively enhance collaboration between multiple organisations will be Open.

Dimension: Perimeterised (Per) / De-perimeterised (D-p) Architectures

The third dimension represents the “architectural mindset” - are you operating inside your traditional IT perimeter or outside it? De-perimeterisation has always related to the gradual failure / removal / shrinking / collapse of the traditional silo-based IT perimeter.



- Perimeterised implies continuing to operate within the traditional IT perimeter, often signalled by “network firewalls”. As has been discussed in previous published Jericho Forum papers, this approach inhibits collaboration. In effect, when operating in the perimeterised areas, you may simply extend your own organisation’s perimeter into the external cloud computing domain using a VPN and operating the virtual server in your own IP domain, making use of your own directory services to control access. Then, when the computing task is completed you can withdraw your perimeter back to its original traditional position. We consider this type of system perimeter to be a traditional, though virtual, perimeter.

⁴ On the assumption that it is how you get your data in and out that is key to being able to move to another service, assuming the service does the same, it does not matter if internal code is the same.

- De-perimeterised, assumes that the system perimeter is architected following the principles outlined in the Jericho Forum’s Commandments and Collaboration Oriented Architectures Framework. The terms Micro-Perimeterisation and Macro-Perimeterisation will likely be in active use here - for example in a de-perimeterised frame the data would be encapsulated with meta-data and mechanisms that would protect the data from inappropriate usage. COA-enabled systems allow secure collaboration. In a de-perimeterised environment an organisation can collaborate securely with selected parties (business partner, customer, supplier, outworker) globally over any COA capable network


The de-perimeterised areas in our Cloud Cube Model use both internal and external domains but the collaboration or sharing of data should not be seen as internal or external – rather it is controlled by and limited to the parties that the using organisations select. For example, in the future frame, one organisation will not feel uncomfortable about allowing data into the internal COA-compliant domain of a collaborating organisation; rather, they will be confident that the data will be appropriately protected.

This means:

- You can operate in any of the four cloud formations so far described (I/P,I/O,E/P,E/O) with either of two architectural mindsets - Perimeterised or De-perimeterised.
- The top-right E/O/D-p cloud formation is likely to be the “sweet spot” where optimum flexibility and collaboration can be achieved.
- A Proprietary cloud provider will likely want to keep you in the left side of the cube, achieved by either continuous innovation that adds value, or by limiting the means of migrating from the proprietary domain. The ability to move from that top-left cloud form to the “sweet-spot” top-right cloud form will require a rare interface because facilitating you making this move is going to be rarely in the cloud supplier’s best business interests.

While the underlying intent remains the same, an added distinction in describing De-perimeterised cloud usage arises in that the detailed description changes based on the level of abstraction at which you choose to operate.

At the heart of all cloud forms is the concept of abstraction. Cloud models separate one layer of business from another, e.g. process from software, platform from infrastructure, etc. We show an example model here with four levels of abstraction; we can expect other models identifying different layers and abstraction levels to emerge to suit different business needs. Most cloud computing activities today are occurring at the lower layers of the stack, so today we have more maturity at the lower level.



Example from one customer:

An early experience of using the external proprietary cloud form represented by Amazon SC3 has involved a combination of Perimeterised Amazon virtual servers and De-perimeterised public data sets to create private results that are then repatriated to the Internal non-cloud environment.

Dimension: Insourced / Outsourced

We define a 4th dimension that has 2 states in each of the 8 cloud forms: Per(IP,IO,EP,EO) and D-p(IP,IO,EP,EO), that responds to the question “Who do you want running your Clouds?”:

- **Outsourced:** the service is provided by a 3rd party
- **Insourced:** the service is provided by your own staff under your control

These 2 states describe who is managing delivery of the cloud service(s) that you use. This is primarily a policy issue (i.e. a business decision, not a technical or architectural decision) which must be embodied in a contract with the cloud provider. In the Cloud Cube Model diagram we show this 4th dimension by 2 colors; any of the 8 cloud forms can take either color.

Note: Few organisations that are traditionally bandwidth, software or hardware providers will be able to easily make the move to becoming Cloud Service providers. It takes a new mindset - a new culture - to be a Cloud Service Provider, just as it takes a new mindset to be a Cloud Service User. We can expect early growing pains caused by both of these transitions. This leaves a market for traditional service providers to help organisations to make this cloud transition.

Given the ease with which a user within your business can procure cloud services – just by tendering a valid credit card - it is absolutely essential that your business develops the agility to rapidly set up legally binding collaboration agreements, and to close them equally rapidly as soon as they are no longer needed⁵. Will it be possible in the future to design a cloud data capsulation approach that means if the cloud provider accepts the data capsule then they automatically accept the terms that the data came with – for example “do not process outside the data owner’s national boundary”?

Note: When closing down an agreement with a provider, care should be taken to ensure that the data is appropriately deleted from the cloud service provider’s infrastructure (including backups), otherwise a data leak risk will remain. “Data Repatriation” is a key new capability.

Other attributes like Offshore and Onshore are also relevant to cloud computing, but in this paper we have focused on the 4 dimensions identified in our Cloud Cube Model.

Background / rationale

Key questions customers need to ask their Cloud Computing suppliers so as to be confident that they are securely collaboratively enabled and compliant with applicable regulations:

1. Where in our cloud cube model is my cloud supplier operating when providing each of their services?⁶
2. How will my cloud supplier assure that when using their services I am operating in a cloud form that has and will maintain the features I expect?
3. How can I ensure that my data and the cloud services will continue to be available, in the event of the provider’s bankruptcy or change in business direction.

⁵ See COA paper: Trust Levels - Business Impact Level, available from the Jericho Forum publications page at <http://www.opengroup.org/jericho/publications.htm> . We expect this agility will require automated set-up and close-down of contractual agreements.

⁶ This matters not just in terms of the characteristics of each cloud form, but also is driven by regulatory requirements that will need to be maintained up-to-date, not least to align with and account for the changes that cloud computing is bringing.

It's important for business managers in their decision-making:

- To understand how and why using any cloud form will return the value-add they want to achieve
- To set out their Cloud Computing requirements clearly, and know what to expect as a result, so they can achieve the great benefits that cloud computing can offer. Entering into any cloud form without establishing the actual business objectives – especially what collaborative flexibility and security they want - may well result in significant problems.
- Moving data, both sensitive and confidential, into the cloud also has legal and compliance issues. These too should be fully understood by all parties before the decision to move to a cloud service is made. It may be that while the cost associated with the cloud service is significantly lower, the business risk is too high.

Challenges to the industry

The major cloud services providers (including Google, Amazon, Salesforce.com, Yahoo) should work with the infrastructure suppliers (including Symantec, IBM, HP, Cisco, Juniper, Microsoft, SAP, TATA), and the Jericho Forum and other relevant consumer interest groups, to develop the services, solutions, and open standards-based interfaces, that customers need for secure open cloud computing (i.e. the right-hand side of the Cloud Cube Model).

These solutions should be based on the Jericho Forum's commandments (design principles) and COA Framework. Cloud computing has undoubted business capabilities and advantages for each cloud formation, and the Jericho Forum's COA Framework explains how to provide the business view. Cloud computing also has the potential to provide the right technical enabling and control capabilities for safe and secure business collaboration, and again the Jericho Forum's COA Framework describes the key technology components involved.

Working in common cause on these areas, the industry can build management of trust into cloud computing, so that everyone prospers – safely and securely - in all cloud formations.

The way forward

We need to provide examples of practical business uses of each cloud form, in use-cases which illustrate the exploitation of the distinguishing features of each of the cloud forms described in this paper. Use-case examples are the subject of another Jericho Forum paper.

Note that the dimensions used in this paper to define cloud forms are primarily business decisions, which may or may not be supported by technologies. Architectural and technology questions to support secure Cloud Computing lie at the next level down from defining these Cloud Cube dimensions. For example, we anticipate that you can't get to that sweet spot in the top right cloud formation – and probably all cloud formations - without appropriate cloud based identity, reputation, authentication, access & authorisation, and governance & compliance. These issues are among the many related topics requiring further study.



Cloud Computing

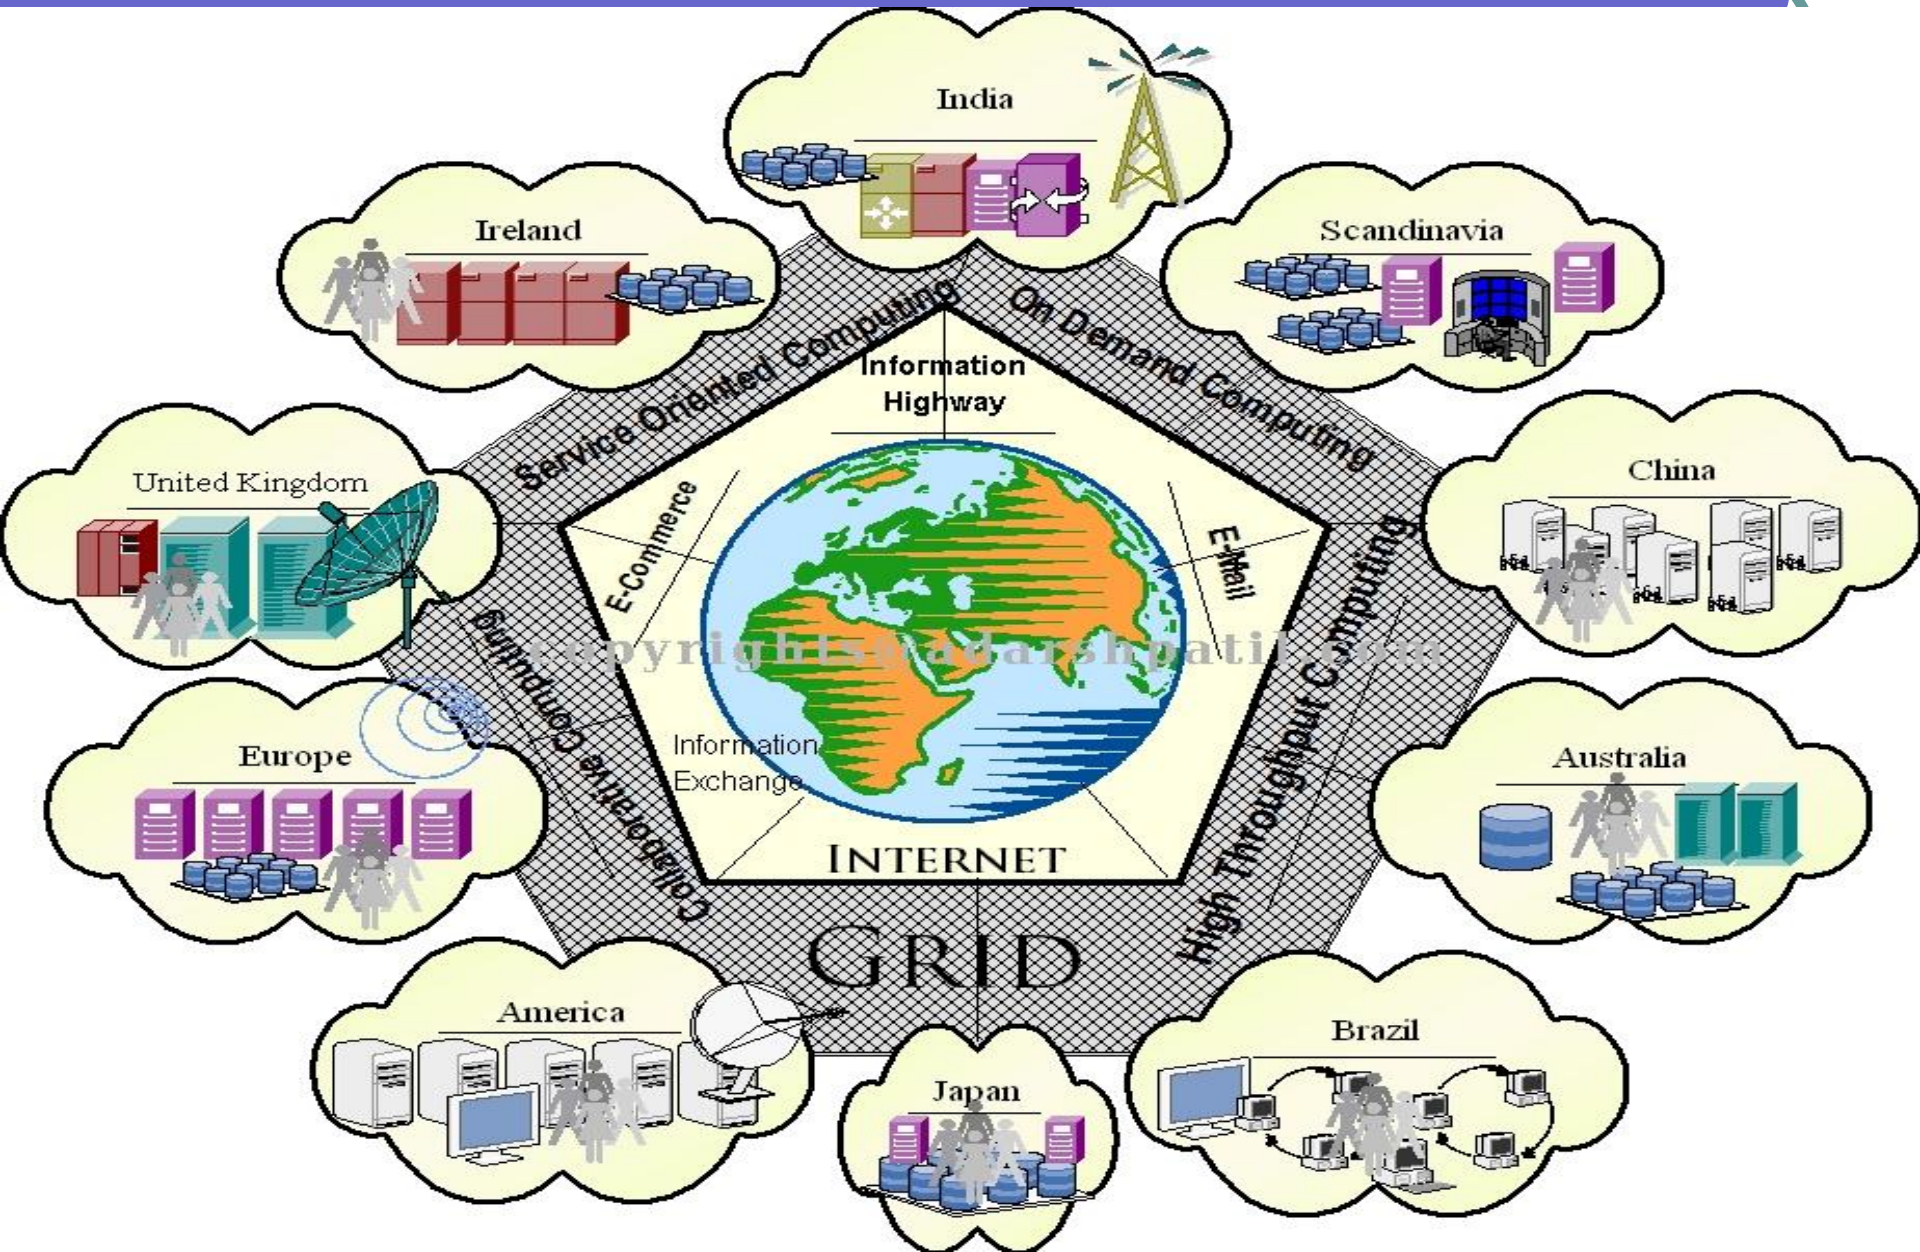
Outline

- Definitions of Cloud computing
- Architecture of Cloud computing
- Benefits of Cloud computing
- Opportunities of Cloud Computing
- Cloud computing – Google Apps
- Grid computing vs Cloud computing

Grid Computing

- Grid is a highly heterogeneous distributed system which is formed by integrating shared resources across many industrial and academic domains by accepting a usage policy.
- Grid resources are
 - End users to Research scientists (**Collaborative Computing**)
 - Software – including open source, custom built and proprietary (Middleware / Network Enable Servers/ Problem solving Environments / Databases) (**Grid Computing System**).
 - Applications running on the resources (**Service Oriented Computing**)
 - Hardware (**Utility Computing**)
 - Silicon/Metals used from PCs/Macs/Mainframes to Workstations to Clusters to Cluster of Workstations/Cluster of Clusters
 - All the hardware components related to Medical Science (X-ray), Astronomy (Telescope), Physics (LHC Project)

Understanding Grid



Definitions

Source	Definition
Gartner	“a style of computing in which massively scalable IT-related capabilities are provided “as a service” using Internet technologies to multiple external customers” (Gartner 2008b)
IDC	“an emerging IT development, deployment and delivery model, enabling real-time delivery of products, services and solutions over the Internet (i.e., enabling cloud services)” (Gens 2008)
The 451 Group	“a service model that combines a general organizing principle for IT delivery, infrastructure components, an architectural approach and an economic model – basically, a confluence of grid computing, virtualization, utility computing, hosting and software as a service (SaaS)” (Fellows 2008)
Merrill Lynch	“the idea of delivering personal (e.g., email, word processing, presentations.) and business productivity applications (e.g., sales force automation, customer service, accounting) from centralized servers” (Merrill Lynch 2008)

Definitions

- Cloud computing is using the internet to access someone else's software running on someone else's hardware in someone else's data center.

Lewis Cunningham^[2]

Definitions

- A large-scale distributed computing paradigm that is driven by economies of scale, in which a pool of abstracted, virtualized, dynamically scalable, managed computing power, storage, platforms, and services are delivered on demand to external customers over the Internet.

Ian Foster^[9]

Definitions

- A Cloud is a type of parallel and distributed system consisting of a collection of interconnected and virtualised computers that are dynamically provisioned and presented as one or more unified computing resources based on service-level agreements established through negotiation between the service provider and consumers.

Rajkumar Buyya^[10]

Outline

- Definitions of Cloud computing
- **Architecture of Cloud computing**
- Benefits of Cloud computing
- Opportunities of Cloud Computing
- Cloud computing – Google Apps
- Grid computing vs Cloud computing

Architecture

- Cloud Service Models
- Cloud Deployment Models
- Essential Characteristics of Cloud Computing

Architecture

Broad
Network Access

Rapid Elasticity

Measured Service

On-Demand
Self-Service

Resource Pooling

**Essential
Characteristics**

Software as a
Service (SaaS)

Platform as a
Service (PaaS)

Infrastructure as a
Service (IaaS)

**Service
Models**

Public

Private

Hybrid

Community

**Deployment
Models**

NIST Visual Model of Cloud Computing Definition

Essential Characteristics^[7]

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

Essential Characteristics^[7]

- Broad network access.
 - Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs) as well as other traditional or cloudbased software services.

Essential Characteristics^[7]

- Resource pooling.
 - The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.

Essential Characteristics^[7]

- **Rapid elasticity.**

- Capabilities can be rapidly and elastically provisioned - in some cases automatically - to quickly scale out; and rapidly released to quickly scale in.
- To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Essential Characteristics^[7]

- **Measured service.**
 - Cloud systems automatically control and optimize resource usage by leveraging a metering capability at some level of abstraction appropriate to the type of service.
 - Resource usage can be monitored, controlled, and reported - providing transparency for both the provider and consumer of the service.

Cloud Service Models

SPI Model

- Cloud **S**oftware as a Service (**SaaS**)
- Cloud **P**latform as a Service (**PaaS**)
- Cloud **I**nfrastructure as a Service (**IaaS**)

Infrastructure as a Service (IaaS)

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

Platform as a Service (PaaS)

- The capability provided to the consumer is to deploy onto the cloud infrastructure consumer created or acquired applications created using programming languages and tools supported by the provider.
- The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Software as a Service (SaaS)

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

Cloud Deployment Models

- Public Cloud.
- Private Cloud.
- Community Cloud.
- Hybrid Cloud.

Public Cloud

- The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Private Cloud

- The cloud infrastructure is operated solely for a single organization. It may be managed by the organization or a third party, and may exist on-premises or off-premises.

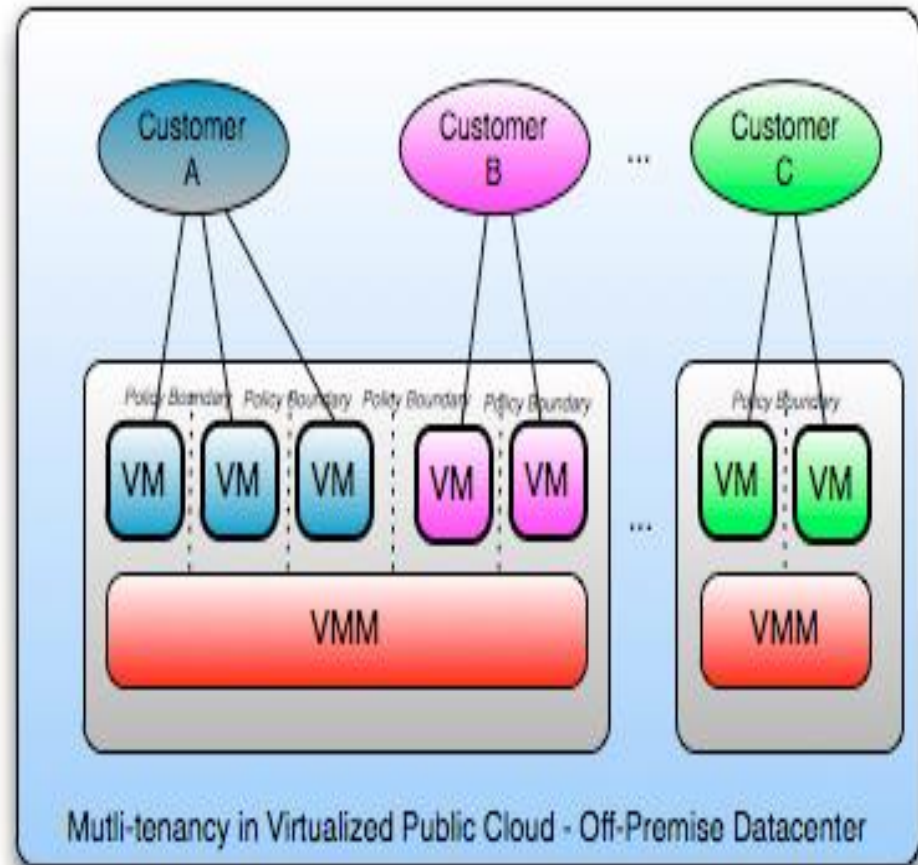

Community Cloud

- The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, or compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

Hybrid Cloud

- The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

Private VS Public Cloud



Private Cloud of Company XYZ with 3 business units, each with different security, SLA, governance and chargeback policies on shared infrastructure

Public Cloud Provider with 3 business customers, each with different security, SLA, governance and billing policies on shared infrastructure

Outline

- Definitions of Cloud computing
- Architecture of Cloud computing
- **Benefits of Cloud computing**
- Opportunities of Cloud Computing
- Cloud computing – Google Apps
- Grid computing vs Cloud computing

Benefits of Cloud Computing

- Business Benefits of Cloud Computing
- Technical Benefits of Cloud Computing

Business Benefits

- Almost zero upfront infrastructure investment
- Just-in-time Infrastructure
- More efficient resource utilization
- Usage-based costing
- Reduced time to market


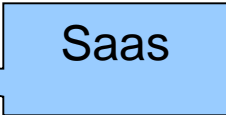
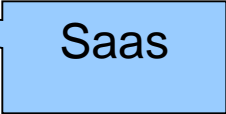
Technical Benefits

- Automation – “Scriptable infrastructure”
- Auto-scaling
- Proactive Scaling
- More Efficient Development lifecycle
- Improved Testability
- Disaster Recovery and Business Continuity

Outline

- Definitions of Cloud computing
- Architecture of Cloud computing
- Benefits of Cloud computing
- **Opportunities of Cloud Computing**
- Cloud computing – Google Apps
- Grid computing vs Cloud computing

Opportunities of Cloud Computing

- End consumers. 
- Business customers. 
- Developers and Independent Software Vendors (ISVs). 

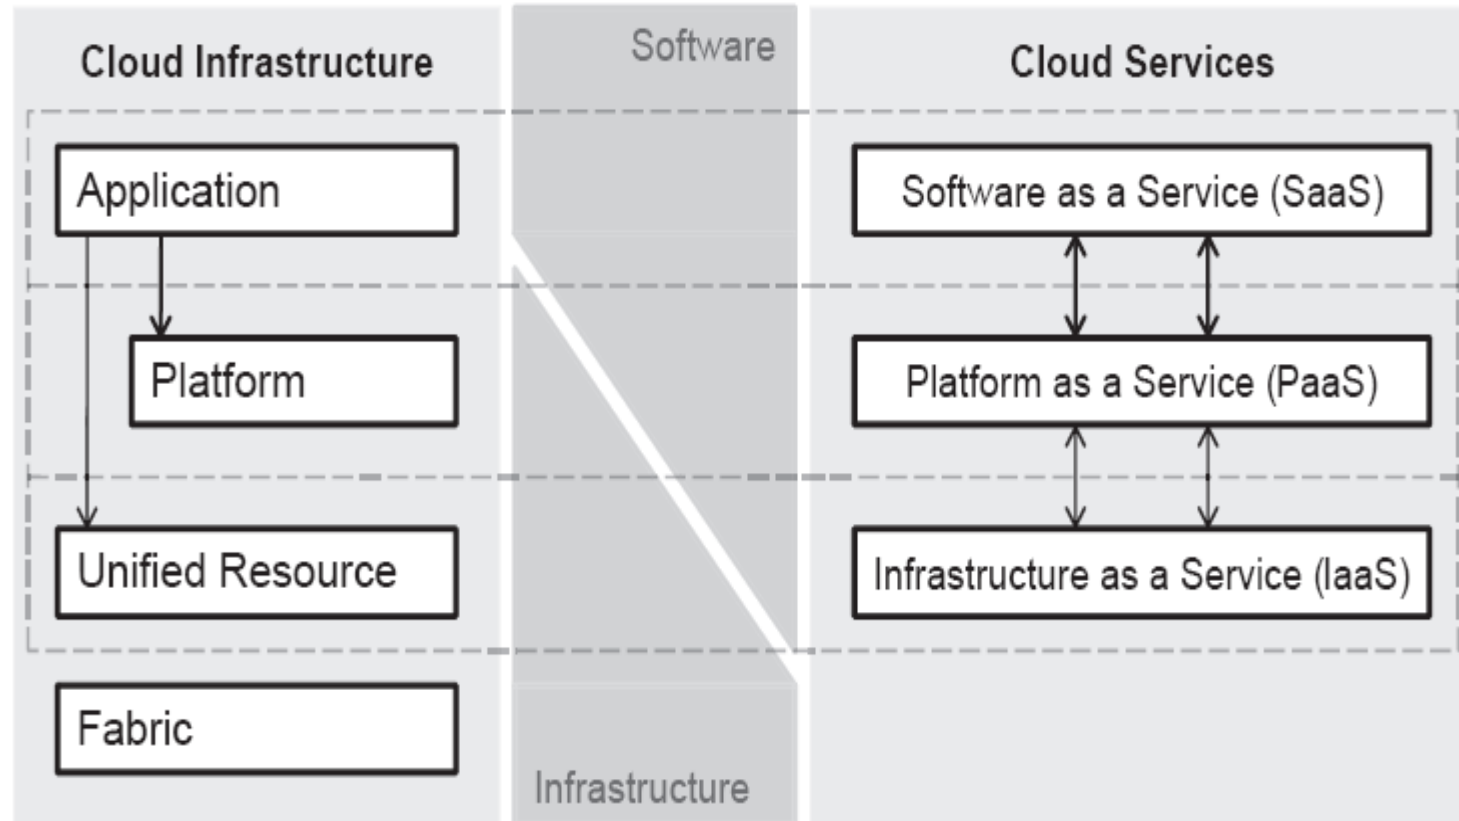
Outline

- Definitions of Cloud computing
- Architecture of Cloud computing
- Benefits of Cloud computing
- Cloud computing – Google Apps
- Grid computing vs Cloud computing

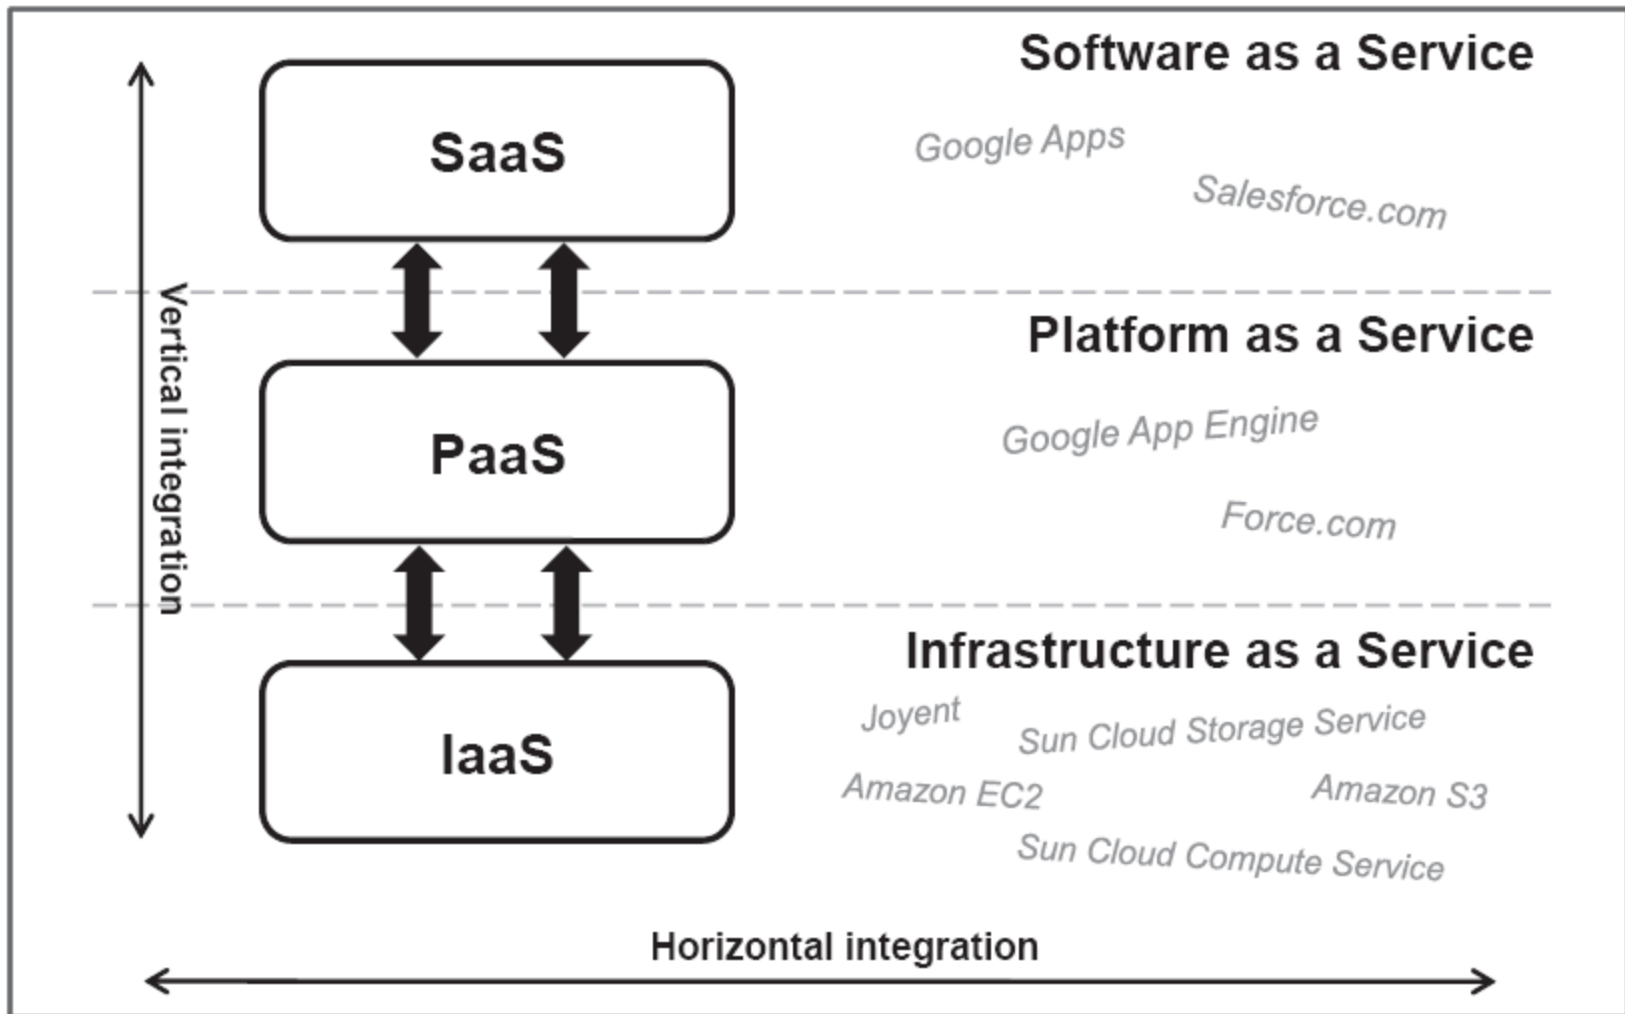
Cloud computing – Google Apps

- Email, chat.
- **Google App Engine**

Google App Engine



Google App Engine



Google App Engine

- Google App Engine?
- Create application.

Google App Engine?

- Google App Engine enables you to build web applications on the same scalable systems that power Google applications. App Engine applications are **easy to build**, **easy to maintain**, and **easy to scale** as your traffic and data storage needs grow.

Easy to build →
Write local,
upload server

Easy to scale →
how many user,
how much data

easy to maintain → 10
year (data & application)

Google App Engine?

- Cost → ?
 - Pay only for what you actually use.
 - Exceed the free quota of **500 MB of storage** and around **5M pageviews** per month.
 - Trial? → 1GB store & 5M pageviews

Create application

- build an App Engine application using standard Java web technologies, such as servlets and JSP.
- create an App Engine Java project with Eclipse → use the Google Plugin for Eclipse for App Engine development. (Use SDK)
- use the App Engine datastore with the Java Data Objects (JDO) standard interface.
- upload your app to App Engine.

Outline

- Definitions of Cloud computing
- Architecture of Cloud computing
- Benefits of Cloud computing
- Opportunities and Challenges of Cloud Computing
- Cloud computing – Google Apps
- Grid computing vs Cloud computing

Grid computing vs Cloud computing

- Same
- Difference

same

- Increase computing.
- Increase store.

difference

- Business model
- Architecture
- Application.

Business model

- Cloud → consumption basis.
- Grid → project-oriented

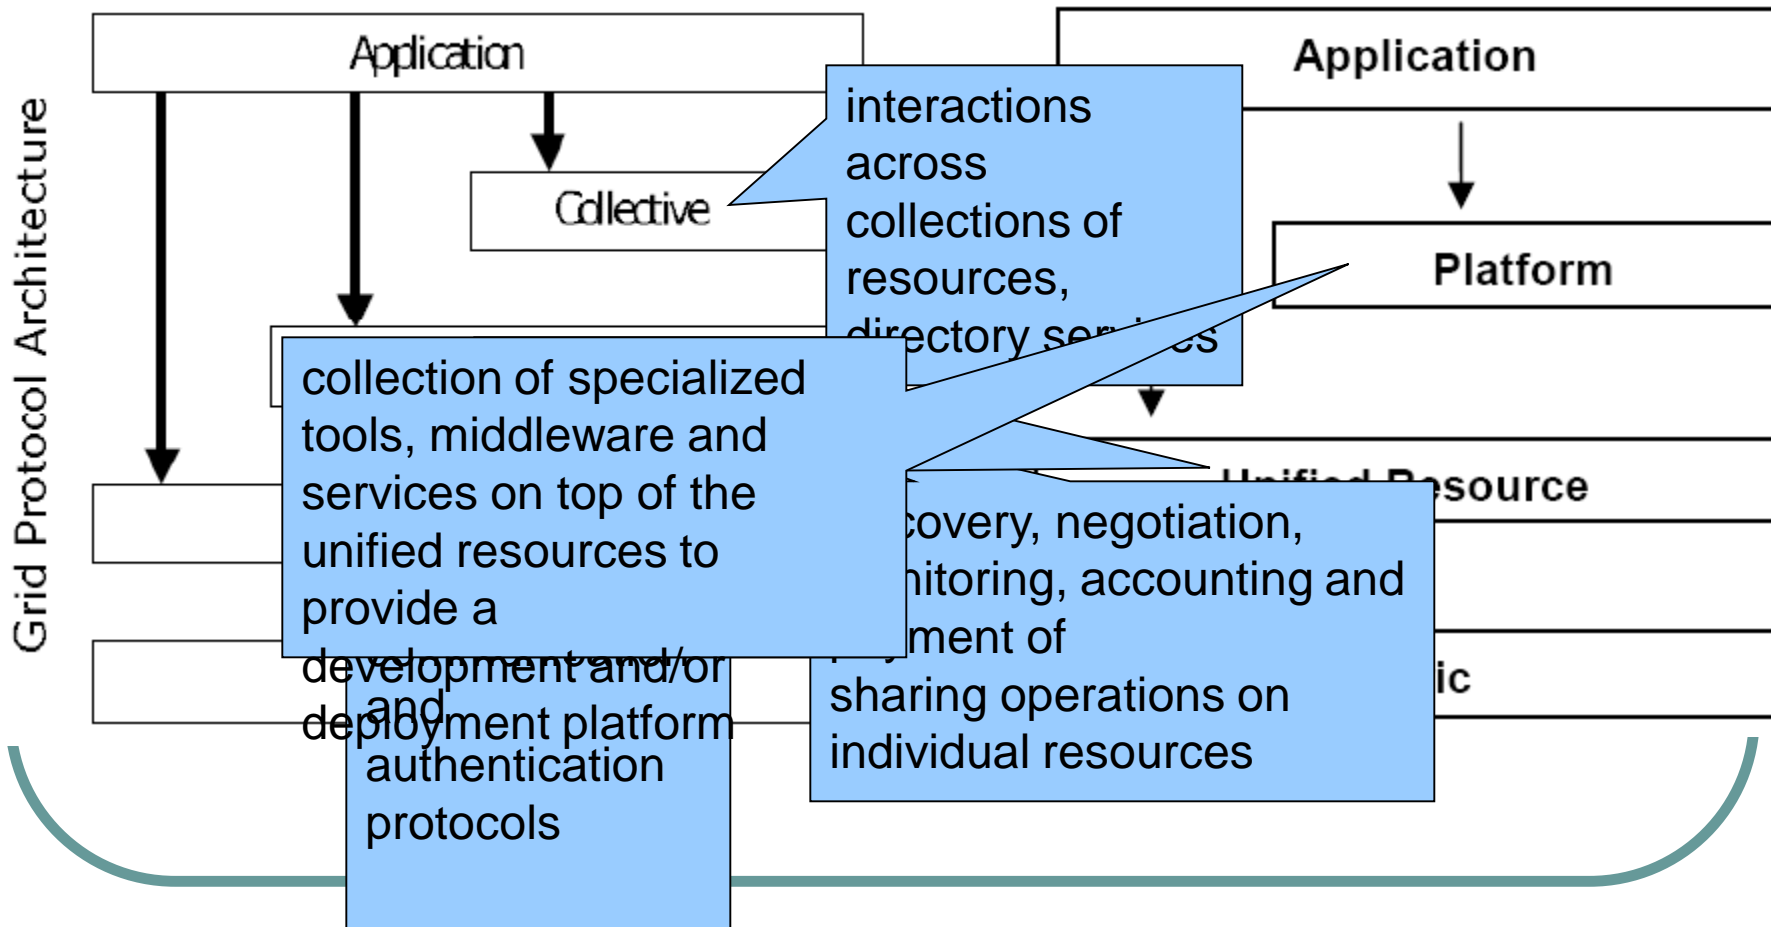
Grid → academia
or government labs

TeraGrid : number
of service units

Cloud → IBM,
Google, Microsoft ...

Hour, storage,
view...

Architecture

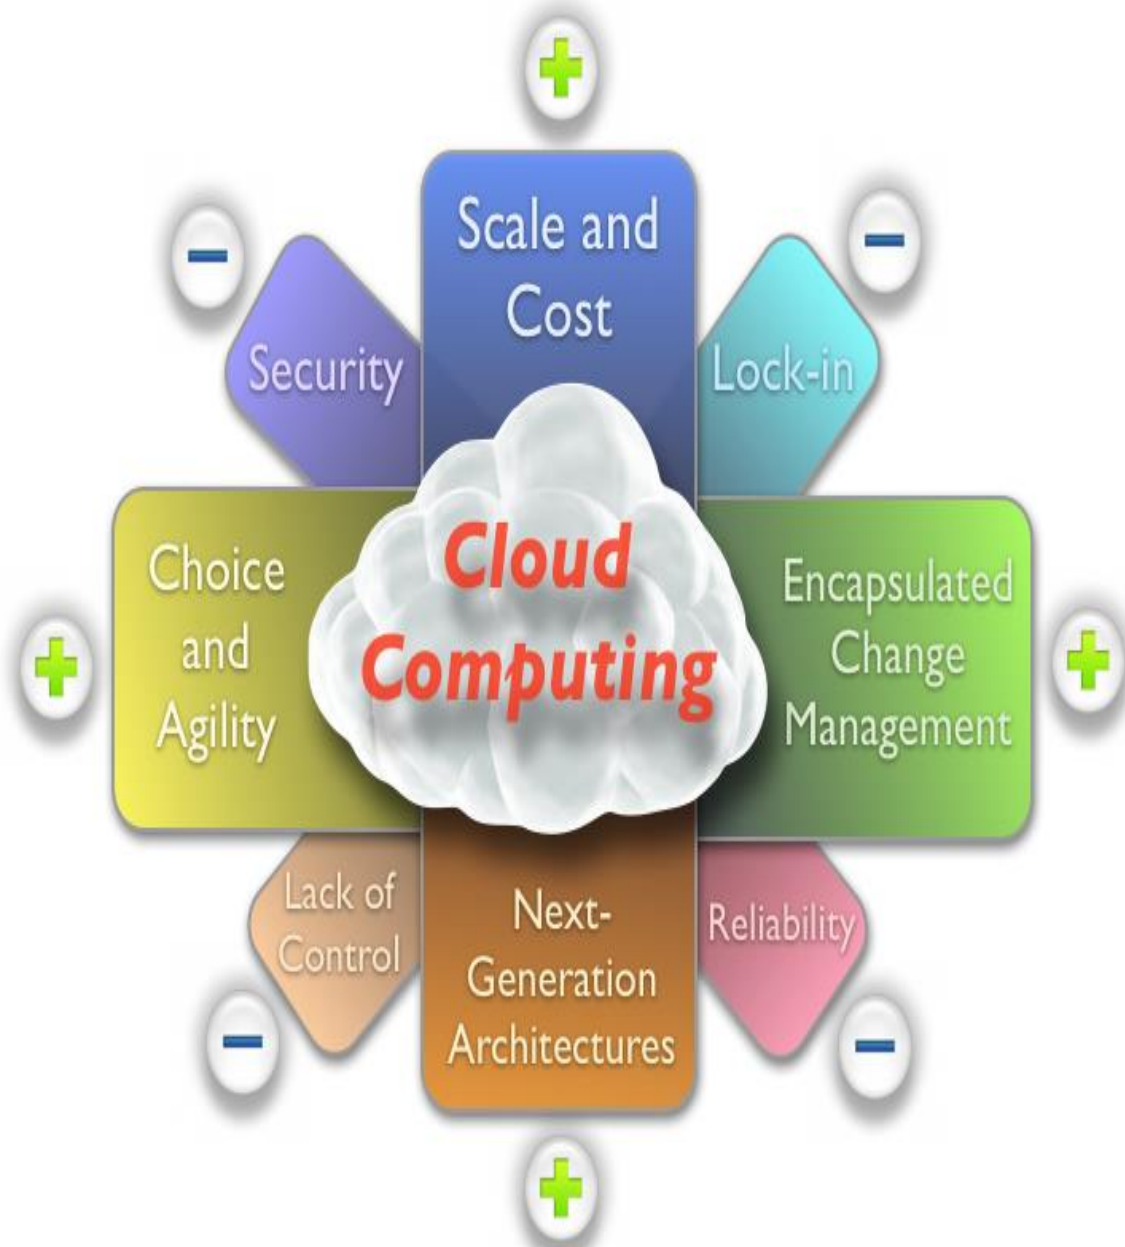


Application

- Grid Computing emerged in eScience to solve scientific problems requiring HPC.
- Cloud Computing is rather oriented towards applications that run permanently and have varying demand for physical resources while running.
 - the well-known CRM SaaS Salesforce.com.



Pros and Cons



Mobile Cloud Computing


Lecture 02a

Cloud Computing I

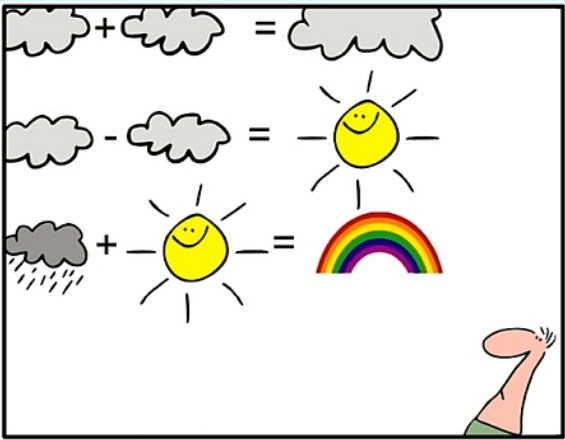


吳 秀 陽
Shiow-yang Wu

What is Cloud Computing?



- Computing with cloud?



geek and poke

**SIMPLY EXPLAINED - PART 17:
CLOUD COMPUTING**

Mobile Cloud Computing Cloud Computing I 2

What is Cloud Computing?



- Walking on the Clouds ? Computing under the clouds?



Mobile Cloud Computing

Cloud Computing I 3

What is Cloud Computing?



Mobile Cloud Computing

Cloud Computing I 4



**T. Sridhar. Cloud Computing—A Primer,
Part 1: Models and Technologies. *The
Internet Protocol Journal*, Volume 12,
No.3, Sep 2009.**

Cloud Computing Definition



- “Cloud computing is a model for enabling **convenient**, **on-demand** network **access** to a **shared pool** of configurable computing **resources** (for example, networks, servers, storage, applications, and services) that can be **rapidly provisioned** and **released** with **minimal management effort** or service provider interaction.” -- NIST
- “Cloud computing is the delivery of computing as a **service** rather than a **product**, whereby shared resources, software, and information are provided to computers and other devices as a **metered service** over a network.
- Cloud computing provides **computation**, **software**, **data access**, and **storage** resources without requiring cloud users to know the location and other details of the computing infrastructure.” -- Wiki

Characteristics of Cloud Comp



- **Elasticity** and **scalability**: Expand and reduce resources according to your specific service requirement.
- **Pay-per-use**: Pay for cloud services only when you use them.
- **On-demand**: Cloud services are invoked only when you need them. They are not permanent parts of your IT infrastructure—a significant advantage. No need to have dedicated resources waiting to be used.

Characteristics of Cloud Comp




- **Resiliency**: Completely isolate the failure of server and storage resources from cloud users. Work is migrated to a different physical resource in the cloud with or without user awareness and intervention.
- **Multitenancy**(多租戶): Can host the cloud services for multiple users with different requirements within the same infrastructure.
- **Workload movement**: Cloud-computing providers can migrate workloads across servers—both inside the data center and across data centers.

Cloud Computing Context



- Shifting from **capital expenditures** (CapEx, ie buying resources for internal IT) to an **operating expense** (OpEx) model, where you **pay for usage** of resources.

Figure 1: Cloud Computing Context



Mobile Cloud Computing

Cloud Computing I 9

What about Hosted Services?



- **Hosted services** — a model in which servers, storage, and networking infrastructure are **shared** across multiple tenants and over a **remote connection** with the ability to **scale** (done manually by provider).
- **Cloud computing** is different in that it offers a **pay-per-use** model and **rapid** (and **automatic**) **scaling up** or **down** of resources along with **workload migration**.
- Cloud computing is much more elastic and scalable with much more flexible business model.

Mobile Cloud Computing

Cloud Computing I 10

Virtualization and Cloud Comp



- **Virtualization** is the key to the success of CC.
- Virtualization software is used to run multiple **Virtual Machines (VMs)** on a single physical server to provide the same functions as multiple physical machines.
- The software is known as a **hypervisor**, which performs the abstraction of the hardware to the individual VMs.
- It was first invented and popularized by IBM in the 1960s for running multiple software contexts on its mainframe computers.

Hypervisor

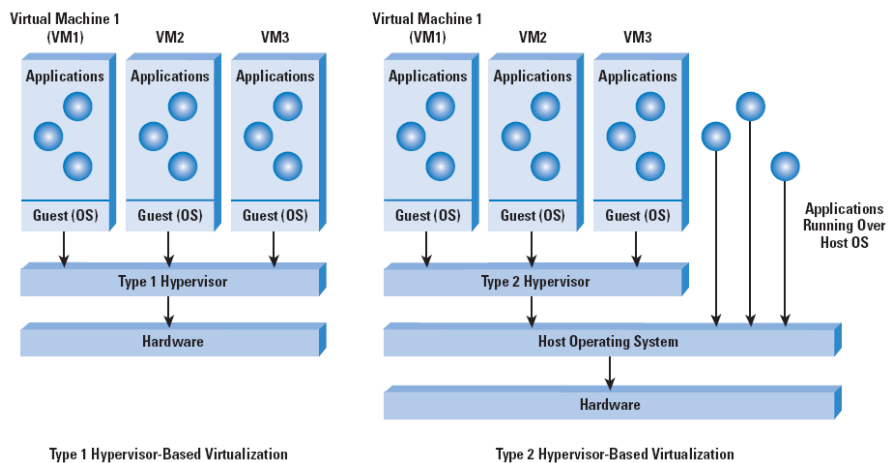


- Hypervisor implementation: (figure on next slide)
 - **Type 1 hypervisor**: directly running over the hardware
 - **Type 2 hypervisor**: running over an operating system
- Support the running of **multiple VMs**, **schedule** the VMs, provide a **unified** and **consistent access** to the CPU, memory... resources on the physical machine.
- A VM runs an **operating system** and **applications**.
- The OS inside the VM may be virtualization-aware and require modifications—a scheme known as **paravirtualization** (as opposed to **full virtualization**).

Hypervisor Implementation



Figure 2: Hypervisors in Virtualization



Mobile Cloud Computing

Cloud Computing 13

VM Migration



- **VM migration** allows you to move an entire VM from one machine to another and **continue operation** of the VM on the second machine.
- This advantage is unique to virtualized environments.
- Can migrate after **suspending** the source VM, **moving** its attendant information to the target machine and **starting** it on the target machine.
- Can also migrate while the VM is **running** (aka. "**live migration**") and **resuming** its operation on the target machine after all the state is migrated.

Mobile Cloud Computing

Cloud Computing I 14

Benefits of Virtualization



- **Elasticity** and **scalability**: Firing up and shutting down VMs involves less effort as opposed to bringing servers up or down.
- **Workload migration**: Can carry out workload migration with much less effort as compared to migration across physical servers at different locations.
- **Resiliency**: Can isolate physical-server failure from user services through migration of VMs.

Virtualization and Cloud



- Virtualization is **not a prerequisite** for cloud computing.
- However, virtualization provides a valuable toolkit and enables significant **flexibility** in cloud-computing deployments.
- Therefore, it is almost adopted by all cloud platforms.

Cloud Computing Models



- Some popular models of cloud computing are offered today as services.
 - **Software as a Service (SaaS)**
 - **Platform as a Service (PaaS)**
 - **Infrastructure as a Service (IaaS)**
- Variations and add-ons:
 - Data Storage as a Service (DaaS)
 - Business Process as a Service (BPaaS)
 - ...?



Software as a Service



- Instead of obtaining desktop and/or licenses for software products, an enterprise can obtain the same functions through a **hosted service** from a provider (known as SaaS provider).
- The interface is usually through a **web browser**.
- **Save** the **complexity** of software installation, maintenance, upgrades, and patches.
- Services can be provided in a **multitenant** model.
- Examples: Google Docs, Webmail, Dropbox
- **Salesforce.com** is an example of a SaaS provider.

SaaS Pros and Cons



- Simple and easy access. Fire up a browser, log in, and go.
- The development, maintenance, updates, backups and so on are the responsibility of the provider.
- Pay-per-use
- The development, backups, updates and so on are *the responsibility of the provider*. You have no control.
- Data security
- Data portability

Mobile Cloud Computing

Cloud Computing I 19

Platform as a Service



- Provide a **software platform** on which users can build their **own applications** and **host them** on the PaaS provider's infrastructure.
- It is used as a **development framework** to build, debug, and deploy applications.
- It often provides **middleware-style services** such as database and component services.
- The **elasticity** and **scalability** is guaranteed transparently by the PaaS platform.
- Examples: Google GAE, Force.com from Salesforce.com

Mobile Cloud Computing

Cloud Computing I 20

PaaS Pros and Cons



- Applications do not need to worry about the **elasticity** and **scalability** issues.
- **Greater** degree of **user control** than SaaS
- Pricing can be on a **per-developer license** and on a **hosted-seats** basis
- Applications need to follow **specific API** and be written in **specific languages** (likely to change in the near future)
- The concerns about **lock-in**
- Not easy to **migrate** existing applications to a PaaS environment

Infrastructure as a Service



- An IaaS provider offers you “raw” computing, storage, and network infrastructure so that you can load your own software, including operating systems and applications, on to this infrastructure.
- Amazon **Elastic Computing Cloud (EC2)** service lets you rent servers with a certain CPU speed, memory, and disk capacity along with the OS and applications that you need.
- Pricing for the IaaS can be on a **usage** or **subscription** basis.

IaaS Pros and Cons



- Offers the **greatest** degree of **control**
- Infrastructure that can **dynamically scaled**
- Much **less cost** than having to build the infrastructure yourself

- **Scaling** and **elasticity** are your—not the provider's—responsibility.
- You need to know the **resource requirements** for your specific application to exploit IaaS well.
- A mini **do-it-yourself data center** that you have to configure to get the job done.

Different Clouds



- **Public clouds** – both the infrastructure and control of these clouds is with the service provider.
- **Private clouds** – the cloud provider is responsible only for the infrastructure and not for the control.
 - A section of a shared data center is partitioned for use by a specific customer.
- **Internal clouds** – cloud services are provided by the IT department of an enterprise from the company's **own** data centers.
 - Better security and control
 - Resiliency, scalability, and workload migration

When Does CC Make Sense?



- For **startup**, you can **focus on your core business** without having to set up and provision your IT infrastructure.
- As your company grows, the cloud-provided IT environment can **scale** along with it.
- When an IT department needs to “**burst**” to access additional IT resources to fulfill a short-term requirement (known as **cloud bursting**).
- **Consistent** and **universal access** to services
- May prove to be **good to the environment** as well.

When not to use Clouds?



- **Regulation** and **legal considerations** may dictate that the enterprise house, secure, and control data in a specific location or geographical area.
- **Access** to the data might need to be **restricted** to a limited set of internal applications.
- When application **response time** is critical, companies might be better off keeping such demanding applications in house.

Cloud Computing Infrastructure



- About the **data center**, the **interconnection** of data centers, and their **connectivity** to the users.
- A data center is similar to a corporate data center but at a different scale to support **multiple tenants** and provide **scalability**, **elasticity** and **virtualization**.
- Google supports the **MapReduce** computing paradigm which takes a set of input key-value pairs, processes it, and produces a set of output key-value pairs.
- Google employs **commodity servers** and **disks** running **Linux** interconnected by **Ethernet switches**.

Cloud Computing Infrastructure



- **Jobs** (set of **tasks**) are scheduled and mapped to the available machine set.
- Implemented by a **Master** machine and **Worker** machines which are scheduled by the Master to implement **Map** and **Reduce tasks**.
- The **topology** and **task distribution** among the servers is optimized for the MapReduce applications.
- SaaS vendors can **partition** their cloud data center according to **load**, **tenant**, and **type** of services.
- May have to redirect the traffic to a different data center

Storage Infrastructure



- **Storage** plays a major part in the data center.
- Can be locally attached or accessible through a net.
- The most popular storage network technologies being **Fibre Channel (FC)** and **Ethernet** with FC or Ethernet **adapters** connected to FC or Ethernet **switches**.
- **Network Attached Storage (NAS)** devices with Ethernet interfaces also have a strong presence.
- **Internet Small Computer System Interface (iSCSI)** is quite popular among smaller data centers and enterprises (SCSI + TCP/IP-over-Ethernet)

Mobile Cloud Computing

Cloud Computing I 29

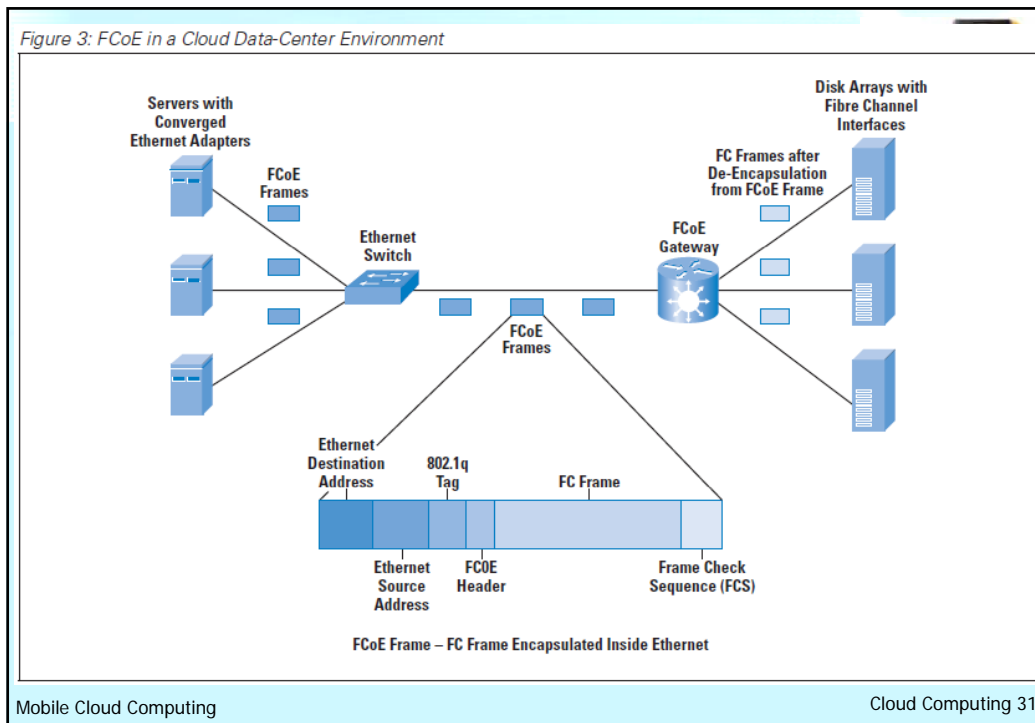
Storage Infrastructure



- **Fibre Channel over Ethernet (FCoE)** removes the need for each server to have an FC adapter.
- FC traffic is **encapsulated** inside an **Ethernet frame** and sent across to a **FCoE gateway** that provides Ethernet-to-FCoE termination to connect to FC storage arrays. (next slide)
- Some storage products provide FCoE functions.
- An adapter on the server that provides both “classical” Ethernet and FCoE functions is known as a **Converged Network Adapter (CNA)**.

Mobile Cloud Computing

Cloud Computing I 30



CC Effect on the Network



- **Network** is also a big part of cloud computing.
- The cloud is accessible through a **public network** (the Internet) or through a **private network**.
- Response-time guarantees depend upon this connectivity.
- Some vendors offer **dedicated links** and appropriate **SLAs** (and charges) for uptime or response time.
- Others might take a **best-effort** scheme but provide **tools** for **monitoring** and **characterizing** application performance and response time.

Network in Data Center



- Consists mainly of servers in **racks** interconnected through a **Top-of-Rack (TOR)** Ethernet switch which, in turn, connects to an **aggregation switch**, sometimes known as an **End-of-Rack (EOR)** switch.
- The **aggregation switch** connects to other aggregation switches.
- A **core switch** connects to the various aggregation switches and provides connectivity to the outside world.
- (Figure 4 on next slide)


Mobile Cloud Computing

Cloud Computing I 33

Data Center Net Architecture



Figure 4: Example Data-Center Switch Network Architecture



Mobile Cloud Computing

Cloud Computing 34

New Standards for Clouds

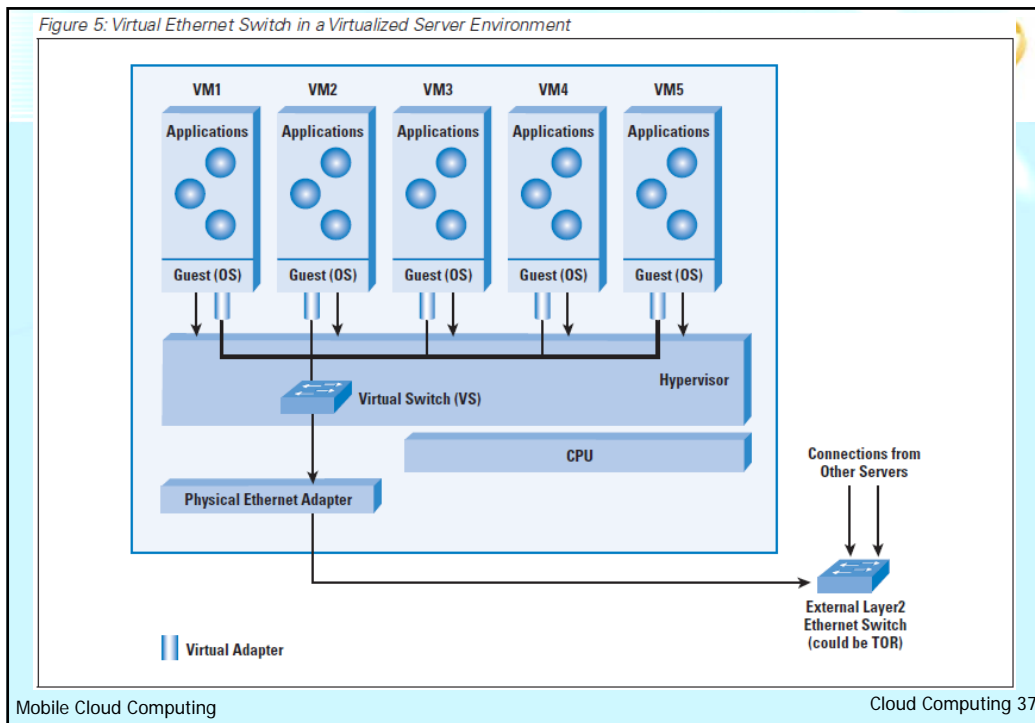


- IEEE is working on new standards for **Convergence Enhanced Ethernet (CEE)**.
- The **Data Center Bridging (DCB) Task Group (TG)** is a part of the IEEE 802.1 Working Group.
- Some examples of new standards:
 - 802.1Qbb – Priority Flow Control (PFC) for FCoE
 - 802.1Qau – end-to-end congestion notification
 - 802.1aq – shortest-path bridging
 - 802.1Qaz – Enhanced Transmission Selection
- <http://www.ieee802.org/1/pages/dcbbridges.html>

Virtualized Network Equipment Functions



- With a virtualized environment, you can move network functions to reside inside virtualized servers.
- You can use the **Virtual Switch** to switch between VMs inside the same physical server and aggregate the traffic for connection to the external switch.
- The Virtual Switch is often implemented as a **plug-in** to the hypervisor.
- The VMs have **virtual Ethernet adapters** that connect to the Virtual Switch.
- (Figure 5 on next slide)



Virtualized Network Equipment Functions



- It is possible to implement a **virtualized firewall** as a self-contained **VM** instead of as a plug-in to the hypervisor. (known as a **firewall virtual appliance**)
- Network packets destined for any of the VMs pass through the **firewall VM** for validation before being passed to other VMs.
- Can also serve as a **front end** to the physical servers in the data center.
- Performance can be an issue (s/w not h/w).

Management Issues in Clouds

- Several facets: **billing**, application-response **monitoring**, **configuring** network resources, and workload **migration**.
- Can manage the network equipment through the **Simple Network Management Protocol (SNMP)** and a network management console.
- The virtualization vendor often offers a **framework** to manage and monitor VMs.
- Several vendors offer products to act as management front ends for public clouds.

Cloud Comp: Common Myths

- **Myth:** Cloud computing should satisfy **all** the requirements specified: scalability, on demand, pay per use, resilience, multitenancy, and workload migration.
- **Facts:**
 - Cloud-computing deployments seldom satisfy all the requirements.
 - Depending upon the type of service offered (SaaS, IaaS, or PaaS), the service can satisfy specific subsets of these requirements.

Cloud Comp: Common Myths

- **Myth:** Cloud computing is useful only if you are outsourcing your IT functions to an external service provider.
- **Facts:**
 - You can use cloud computing in your own IT department for on-demand, scalable, and pay-per-use deployments.
 - You can build your own internal cloud.

Cloud Comp: Common Myths

- **Myth:** Cloud computing requires virtualization.
- **Facts:**
 - Virtualization is not a requirement for cloud computing.
 - But it is likely to see increased usage in cloud deployments.
- **Myth:** Cloud computing requires you to expose your data to the outside world.
- **Facts:**
 - Internal clouds can be used.
 - Can put front ends in the cloud and backend in local.

Cloud Comp: Common Myths

- **Myth:** Converged networks are essential to cloud computing.
- **Facts:**
 - Cloud computing is possible without converged networks.
 - Use of converged networks results in cost efficiencies, but it is not a requirement.

Cloud Comp: Gaps and Concerns

- **Security:**
 - a significant concern for enterprise IT managers
 - Cloud provider must guarantee **data isolation** and **application security** (and **availability**) through isolation across multiple tenants.
 - **authentication** and **authorization** of cloud users and **encryption** of the “network pipe” are other factors.
- **Network concerns:**
 - What to do when cloud bursting is involved?
 - Networking across multiple cloud data centers?

Cloud Comp: Gaps and Concerns



■ Cloud-to-cloud and Federation Concerns:

- When an enterprise uses two separate cloud service providers, how do they **interoperate**.
- What about **VM migration** is in federation?

■ Legal and regulatory concerns:

- Especially important for cases involving **storing data in the cloud**.
- It could be that the **laws** governing the data are not the laws of the **jurisdiction** where the company is located.