Volume: 2 Issue: VII Month of publication: July 2014

DOI:

INTERNATIONAL JOURNAL FOR RESEARCH IN APPLIED SCIENCE AND ENGINEERING TECHNOLOGY (IJRASET)

Cloud Computing: A Beginners Primer

K.Vijesh[1], P.Santhadevi[2]

[1]*Head – School of Computer Science and Information Technology*
DMI-St. John the Baptist University, Mangochi, Malawi.
[2]*Assistant Lecturer - School of Computer Science and Information Technology*
DMI-St. John the Baptist University, Mangochi, Malawi.

Abstract: Information Technology (IT) has embarked on a new paradigm — cloud computing. Cloud computing provides alternative services to deliver computer resources and has sparked a revolution in the way organizations provide information and service. It is an Internet-based computing solution where shared resources are configured to work together with various applications utilizing a collective computing power in the delivery of data. Previously IT services structures- websites and server-based applications-were executed on a specific system. With the advent of cloud computing, resources are used as an aggregated virtual computer. This merged configuration provides an environment where applications are executed independently with little regard for particular configurations. This article reflects work-in-progress information on cloud computing.

Keywords: Cloud computing, infrastructure, models, information technology environment, networks, servers, elasticity, scalability, resiliency, multi-tenancy.

I. INTRODUCTION

Cloud computing is an emerging area that affects IT infrastructure, network services, and applications. This article introduces various aspects of cloud computing, including the rationale, underlying models, and infrastructures and specific technologies and scenarios.

The term "cloud computing" has different connotations for IT professionals, depending upon their point of view and often their own products and offerings. As with all emerging areas, real-world deployments and customer success stories will generate a better understanding of the term. This discussion employs the National Institute of Standards and Technology (NIST) - America.

Definition prescribed for cloud computing by NIST: "Cloud computing is a [information technology environment]model for enabling convenient, on-demand network access to a shared pool of configurable computing resources including networks, servers, storage, applications, and services that can be rapidly provisioned and released with minimal management effort or service provider interaction."

A list of cloud-computing environment can be identified by the following characteristics: (a) elasticity and scalability; (b) pay-per-use; (c) use-on–demand; (d) resiliency;(e) multi-tenancy and (f) workload movement; However, not all characteristics may be present in a specific cloud solution.

1.1 Elasticity and scalability

Cloud computing offers the customer the ability to expand and reduce resources according to its specific service requirement. For example, a customer may need a large number of server resources for the duration of a specific task. The server resources are released after a specific task is completed.

1.2 Pay-per-use

INTERNATIONAL JOURNAL FOR RESEARCH IN APPLIED SCIENCE AND ENGINEERING TECHNOLOGY (IJRASET)

A customer can pay for cloud services when it hires the facility for varied time period. The hire period can be temporary or long term. For example, for CPU time or for cloud-based storage or vault services.

1.3 On demand Usage

Cloud computing is not a permanent part of a business It infrastructure because the user can invoke cloud services only when need arises. This is a significant advantage for cloud use as opposed to internal IT services. With cloud services, there is no need to have dedicated resources waiting to be used, as is the case with internal services.

1.4 Resiliency

The resiliency of a cloud service offering can 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.

1.5 Multi-tenancy

Public cloud services providers can host cloud computing services for multiple users within the same infrastructure. Server and storage isolation may be physical or virtual depending on specific user requirements.

1.6 Workload movement

This characteristic is related to resiliency and cost considerations. Cloud-computing providers can migrate workloads across servers inside the data centre and across data centres to different geographic area. This migration might be necessitated by cost. It may be less expensive to run a workload in a data centre in another country based on time of day or power requirements or efficiency considerations including network bandwidth. A third reason could be regulatory considerations for certain types of workloads.

1.7 Cloud computing Functions and Structure.

Cloud computing involves shifting the bulk of the costs from capital expenditures (CapEx), or buying and installing servers, storage, networking, and related infrastructure to an operating expense (OpEx) model, where the user only pays for usage of cloud computing resources. Figure 1 provides a context diagram for cloud computing.

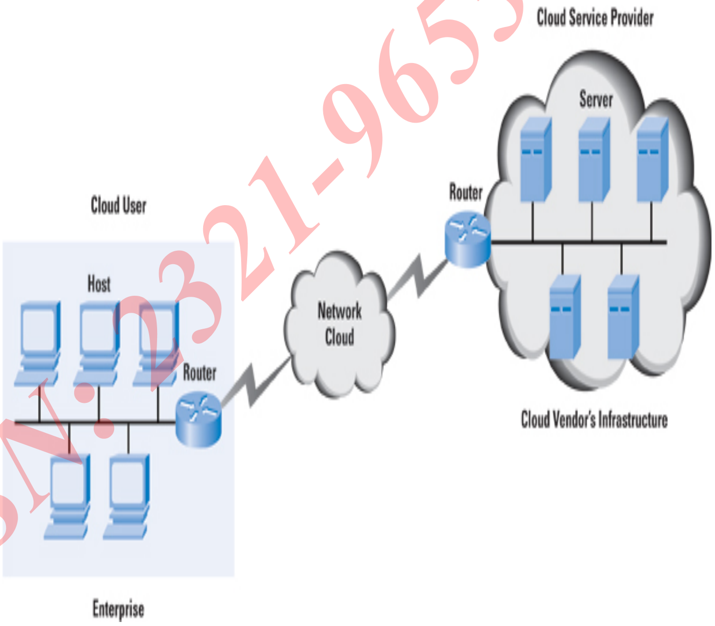

Figure 1: Cloud Computing Context

II. RESEARCH METHODOLOGY

2.1 Exploration on Virtualization and its effects on Cloud Computing

What is Virtualization?

Cloud computing has accelerated because of the popularity and adoption of virtualization, specifically, server virtualization. Virtualization was invented and popularized by IBM in the 1960s for running multiple software contexts on its mainframe computers. It regained popularity in the first decade of 2000 in data centers because of server usage concerns.

Data centers and web farms consisted of multiple physical servers. Measurement studies on these server farms noted that individual server usage was often as low as 15 percent, the idleness of some server applications. The consequence of this server sprawl with low usage was a large financial outlay for CapEx and OpEx, such large investment in extra machines, related power, their cooling infrastructure and hosts -real estate.

Virtualization is software used to run multiple Virtual Machines (VMs) on a single physical server to provide the same functions as multiple physical machines. The virtualization software is known as a hypervisor. This virtualization software performs the abstraction of the hardware and serves data to individual VMs.

Enter Virtualization.

A hypervisor is implemented on a server either directly running over the hardware (a Type 1 hypervisor) or running over an operating system (OS) (a Type 2 hypervisor). The hypervisor supports the running of multiple VMs and schedules the VMs along by providing a unified and consistent access to the CPU, memory, and I/O resources on the physical machine. A VM typically runs an operating system and applications. The applications are not aware that they are running in a virtualized environment, so they do not need to be changed to run in such an environment. Figure 2 depicts these scenarios. The OS inside the VM may be virtualization—aware and require modifications to run over a hypervisor—a scheme known as para-virtualization (as opposed to full virtualization).

Virtual Machine Migration: an advantage of migration

Vendors have implemented VM migration in their virtualization solution. It is an advantage for application up-time in a data centre. What is VM migration? Consider the case of a server with a hypervisor and several VMs, each running an OS and applications. If one needs to bring down the server for maintenance such as adding more memory to the server, one has to shut down the software components and restart them after the maintenance window—significantly affecting application availability. VM migration allows you to move an entire VM with its contained operating system and applications from one machine to another and continue operation of the VM on the second machine. This advantage is unique to virtualized environments because you can take down physical servers for maintenance with minimal effect on running applications.

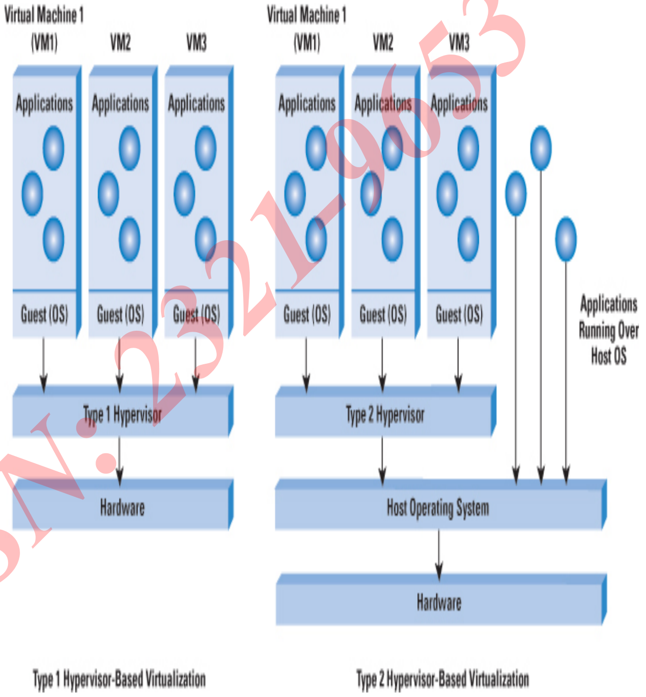

Figure 2: Hypervisors in Virtualization

One can perform this migration after suspending the VM on the source machine, moving its attendant information to the target machine and starting it on the target machine. To lower the downtime, one can perform this migration while the VM is running, hence the phrase “live migration" and resuming its operation on the target machine after all the state is migrated

2.2 Benefits of Virtualization

The following are some of the benefits of virtualization in a cloud-computing environment:

- Elasticity and scalability: Firing up and shutting down VMs involves less effort as opposed to bringing servers up or down.
- Workload migration: Through facilities such as live VM migration, you can carry out workload migration with much less effort as compared to workload migration across physical servers at different locations.
- Resiliency: You can isolate physical-server failure from user services through migration of VMs.

It must be clarified that virtualization is not a prerequisite for cloud computing. In fact, there are examples of large cloud service providers using only commodity hardware servers (with no virtualization) to realize their infrastructure. However, virtualization provides a valuable toolkit and enables significant flexibility in cloud-computing deployments.

2.3 Cloud Computing Major Models

Whilst there are variations based on specific vendor offerings, some models of cloud computing are offered today as services. These models are discussed hereunder.

Software as a Service Consider the case of an enterprise with its set of software licenses for the various applications it uses. These applications could be in human resources, finance, or customer relationship management, to name a few. Instead of obtaining desktop and server licenses for software products it uses, an enterprise can obtain the same functions through a hosted service from a provider through a network connection. The interface to the software is usually through a web browser. This common cloud-computing model is known as Software as a Service (SaaS) or a hosted software model; the provider is known as the SaaS Provider.

Software as a service (SaaS) saves the complexity of software installation, maintenance, upgrades, and patches (for example, for security fixes) for the IT team within an enterprise, because the software is now managed centrally at the SaaS provider's facilities. Further, the SaaS provider can provide this service to multiple customers and enterprises, resulting in a multi-tenant model. The pricing of such a SaaS service is typically on a per-user basis for a fixed bandwidth and storage. Monitoring application-delivery performance is the responsibility of the SaaS provider.

Salesforce.com is an example of a SaaS provider. The company was founded to provide hosted software services, unlike some of the software vendors that have hosted versions of their conventional offerings.

2.4 Platform as a Service

Platform as a Service (PaaS) provides a software platform on which users can build their own applications and host them on the PaaS provider's infrastructure. The software platform is used as a development framework to build, debug, and deploy applications. It often provides middleware-style services such as database and component services for use by applications. PaaS is a true cloud model in that applications do not need to worry about the scalability of the underlying platform (hardware and software). When enterprises write their application to run over the PaaS provider's software platform, the elasticity and scalability is guaranteed transparently by the PaaS platform.

The platforms offered by PaaS vendors like Google (with its App-Engine) or Force.com (the PaaS offering from Salesforce.com) require the applications to follow their own Application Programming Interface (API) and be written in a specific language. This situation is likely to change but is a cause for concerns about lock-in. Further, it is not easy to migrate existing applications to a PaaS environment. Consequently, PaaS sees the most success with new applications being developed specifically for the cloud. Monit¬oring application-delivery performance is the responsibility of the PaaS provider. Pricing for PaaS can be on a per-application developer license and on a hosted-seats basis. The PaaS has a greater degree of user control than SaaS.

2.5 Infra-sturacture as a Service

An Infrastructure as a Service (IaaS) provider offers business “raw" computing, storage, and network infrastructure so that it can load its own software, including operating systems and applications, on to this infrastructure. This scenario is equivalent to a hosting provider provisioning physical servers and storage and letting business install your own OS, web services, and database applications over the provisioned machines. Amazon lets businesses rent servers with a certain CPU speed, memory, and disk capacity along with the OS and applications that they need to have installed on them. Amazon is the first and major proponent of Infrastructure as a Service (IaaS) through its Elastic Computing Cloud (EC2) service. It provides some "canned" software for the OS and applications known as Amazon Machine Images [AMIs], as a starting point. However, you can also install your own OSs (or no OS) and applications over this server infrastructure.

Infrastructure as a Service (IaaS) offers businesses the greatest degree of control of the three models. Businesses need to know the resource requirements for their specific application to exploit IaaS well. Scaling and elasticity are businessess’ — responsibility. In fact, it is a mini do-it-yourself data centre that you have to configure to get the job done. Interestingly, Amazon uses virtualization as a critical underpinning of its EC2 service, so businesses actually get a VM when they ask for a specific machine configuration,. Pricing for the IaaS can be on a usage or subscription basis. CPU time, storage space, and network bandwidth (related to data movement) are some of the resources that can be billed on a user-pay basis.

These are three common models for cloud computing. They have variations and add-ons, including Data Storage as a Service (providing disk access on the cloud), communications as a service. For example, a universal phone number through the cloud.

2.6 Public Private and Internal Clouds

We have focused on cloud service providers whose data centres are external to the users of the service (businesses or individuals). These clouds are known as public clouds—both the infrastructure and control of these clouds is with the service provider.

A variation on this scenario is the private cloud.The cloud provider is responsible only for the infrastructure and not for the control. This setup is equivalent to a section of a shared data centre being partitioned for use by a specific customer. Note that the private cloud can offer SaaS, PaaS, or IaaS services.IaaS appears to be a more natural fit.

An internal cloud is a relatively new term applied to cloud services provided by the IT department of an enterprise from the company's own data centres. This setup might seem counterintuitive at first—why would a company run cloud services for its internal users when public clouds are available? Doesn't this setup negate the advantages of elasticity and scalability by moving this service to inside the enterprise?

The internal cloud model is very useful for enterprises. The biggest concerns for enterprises to move to an external cloud provider are security and control. Central Intelligence Organizations (CIO) is naturally cautious about moving their entire application infrastructure and data to an external cloud provider, especially when they have several person-years of investment in their applications and infrastructure as well as elaborate security safeguards around their data. However, the advantages of the cloud—resiliency, scalability, and workload migration—are useful to have in the company's own data centers. IT can use per-usage billing to monitor individual business unit or department usage of the IT resources and charge them back. Controlling server sprawl through virtualization and moving workloads to geographies and locations in the world with lower power and infrastructure costs are of value in a cloud-computing environment. Internal clouds can provide all these benefits.

This classification of cloud’ computing as public, private, and internal is not universally accepted. Some researchers see the distinction between private and internal clouds to be a matter of semantics. In fact, the NIST draft definition considers a private cloud to be the same as an internal cloud. However, the concepts are still valid and being realized in service provider and enterprise IT environments today.

III. ANALYSIS AND RESULTS

3.1 When does Cloud computing make sense?

Outsourcing businesses entire IT infrastructure to a cloud provider makes sense if your deployment is a "green field" one, especially in the case of a start-up Business can focus on your their core business without having to set up and provision your IT infrastructure, especially if it primarily involves basic elements such as e-mail, word processing, collaboration tools, and so on. As the company grows, the cloud-provided IT environment can scale along with it.

Another scenario for cloud usage is when an IT department needs to "burst" to access additional IT resources to fulfill a short-term requirement. For example, testing of an internally developed application to determine scalability, prototyping of "nonstandard" software to evaluate suitability, and execution of a one-time task with an exponential demand on IT resource.. The term cloud bursting is used to describe this scenario. The cloud resources may be loosely or tightly coupled with the internal IT resources for the duration of the cloud bursting. In an extremely loosely coupled scenario, only the results of the cloud bursting are provided to the internal IT department. In the tightly coupled scenario, the cloud resources and internal IT resources are working on the same problem and require frequent communication and data sharing.

Cloud computing sometimes does not make sense for an enterprise. 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 applications, all of which need to be internal.

Another situation where cloud computing is not always the best choice is when application response time is critical. Internal IT departments can plan their server infrastructure and the network infrastructure to accommodate the response-time requirements. Although some cloud providers provide high-bandwidth links and can specify Service-Level Agreements (SLAs) as in the case of SaaS for their offerings, companies are better off keeping such demanding applications in house.

A variation of these scenarios is when companies outsource their web front ends to a cloud provider whilst keeping their application and database servers internal to the enterprise. This setup is useful when the company is ramping up its offerings on the web but is not completely certain about the demand. It can start with a small number of web servers and scale up or down according to the demand. Further, acceleration devices such as Application Delivery Controllers (ADCs) can be placed in front of the web servers to ensure performance. These devices provide server load balancing, Secure Sockets Layer (SSL) front ends, caching, and compression. The deployment of these devices and the associated front-end infrastructure can be completely transparent to the company. It only needs to focus on the availability and response time of its application behind the web servers.

3.2 Cloud Computing Infrastructure

The most significant discussion on cloud computing infrastructure is related to the data centre, the interconnection of data centres, and their connectivity to the users of the cloud service.

A simple view of the cloud data centre is that it is similar to a corporate data centre but at a different scale because it has to support multiple tenants and provide scalability and elasticity. In addition, the applications hosted in the cloud as well as virtualization (when it is used) also play a part.

A case in point, is the MapReduce computing paradigm that Google implements to provide some of its services (other companies have their own implementations of MapReduce). The MapReduce scheme takes a set of input key-value pairs, processes it, and produces a set of output key-value pairs. To realize the implementation, Google has an infrastructure of commodity servers running Linux interconnected by Ethernet switches. Storage is local through inexpensive Integrated Drive Electronics (IDE) disks attached to each server.

Jobs, which consist of a set of tasks, are scheduled and mapped to the available machine set. The scheme is implemented through a Master machine and Worker machines. The latter are scheduled by the Master to implement Map and Reduce tasks, which themselves operate on chunks of the input data set stored locally. The topology and task distribution among the servers is optimized for the application (MapReduce in this case). Although Google has not made public the details of how the back-end infrastructure is implemented for Google Apps and Gmail, we can assume that the physical and logical organization is optimized for the tasks that need to be carried out, in a manner similar to what is done for Map Reduce.

Structure as a service vendor can partition their cloud data centre according to load, tenant, and type of application that they will offer as a service. In some cases they might have to redirect the traffic to a different data centre, based on the load in the default data centre. IaaS provides the greatest degree of control for the user, as discussed earlier. Even here, the topology and load assignment can be based on the number and type of servers that are allocated.

3.3 Storage Infrastructure

Storage plays a major part in the data centre and for cloud services, especially in environments with virtualization. Storage can be locally attached or accessible through a network—the most popular storage network technologies being Fibre Channel and Ethernet. For such network access of storage, servers are equipped with Fibre Channel or Ethernet adapters through which they connect to a Fibre Channel or Ethernet switch. The switch provides the connectivity to storage arrays. Fibre Channel is . Ethernet interfaces also have a strong presence in the data centre.

Another Ethernet-based storage option is the Internet Small Computer System Interface (iSCSI), which is quite popular among smaller data centres and enterprises because of the cost benefits. This technology involves running the SCSI protocol on a Transport control Protocol/Internet Protocol (TCP/IP)-over-Ethernet connection.

Fibre channel connections to the storage network necessitate two types of network technologies in the data centre: Ethernet for server-to-server and server-to-client connectivity and Fibre Channel for server-to-storage connectivity.

A recent initiative in data-centre technology is a converged network, which involves the transport of Fibre Channel over Ethernet (FCoE). FCoE removes the need for each server to have a Fibre Channel adapter to connect to storage. Instead, Fibre Channel traffic is encapsulated inside an Ethernet frame and sent across to a FCoE gateway that provides Ethernet-to-FCoE termination to connect to Fibre Channel storage arrays (refer to Figure 3). Some storage products provide FCoE functions, so the Ethernet frame can be carried all the way to the storage array. An adapter on the server that provides both "classical" Ethernet and FCoE functions is known as a Converged Network Adapter (CNA). Cloud-computing environments can reduce the data-centre network complexity and cost through this converged network environment.

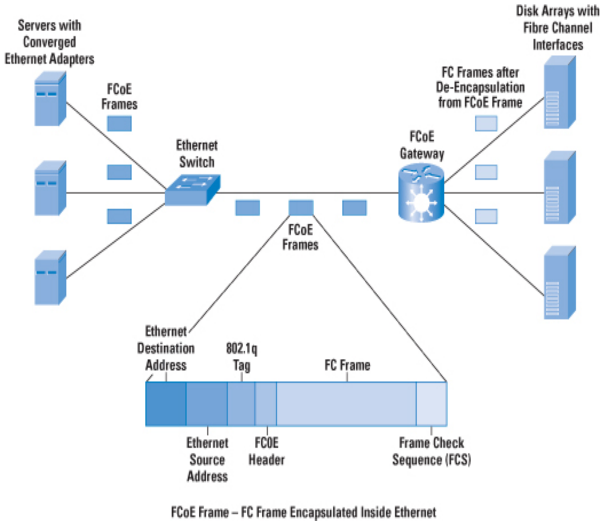
Servers with Converged Ethernet Adapters
FCoE Frames
Ethernet Switch
FCoE Frames
FCoE Gateway
FC Frames after De-Encapsulation from FCoE Frame
Disk Arrays with Fibre Channel Interfaces
Ethernet Destination Address
802.1q Tag
FC Frame
Ethernet Source Address
FCOE Header
Frame Check Sequence (FCS)
FCoE Frame – FC Frame Encapsulated Inside Ethernet

Figure 3: FCoE in a Cloud Data-Centre Environment

3.4 Cloud Computing Effects on the Network

Discussions elsewhere indicated that the network is a big part of cloud computing. A cloud user connects to the network to access the cloud resources, as indicated earlier in Figure 1. The cloud is accessible through a public network (the Internet) or through a private network (dedicated lines or Multiprotocol Label Switching [MPLS] infrastructure,). Response-time guarantees depend upon this connectivity. Some cloud vendors offer dedicated links to their data centers and provide appropriate SLAs (What is SLA?) for uptime or response time and charge for such SLAs. Others might implement a best-effort scheme but provide tools for monitoring and characterizing application performance and response time, so that users can plan their bandwidth needs.

The most significant effect on the network is in the data centre, as indicated previously.. The most common network architecture for enterprises is the three-layer architecture with access, aggregation or distribution, and core switches. The data centre requires a slightly different variation to this layering, as proposed by some vendors. The data centre 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 (Figure 4).

The aggregation switch connects to other aggregation switches and through these switches to other servers in the data centre. A core switch connects to the various aggregation switches and provides connectivity to the outside world, typically through Layer 3 (IP). It can be argued that most of intra-data centre traffic traverses only the TOR and the aggregation switches. Hence the links between these switches and the bandwidth of those links need to account for the traffic patterns. Some vendors have proposed a fat-tree or a leaf-spine topology to address this anomaly, though this is not the only way to design the data-centre network. Incidentally, the fat-tree topology is not new—it has been used in Infiniband networks in the data centre.

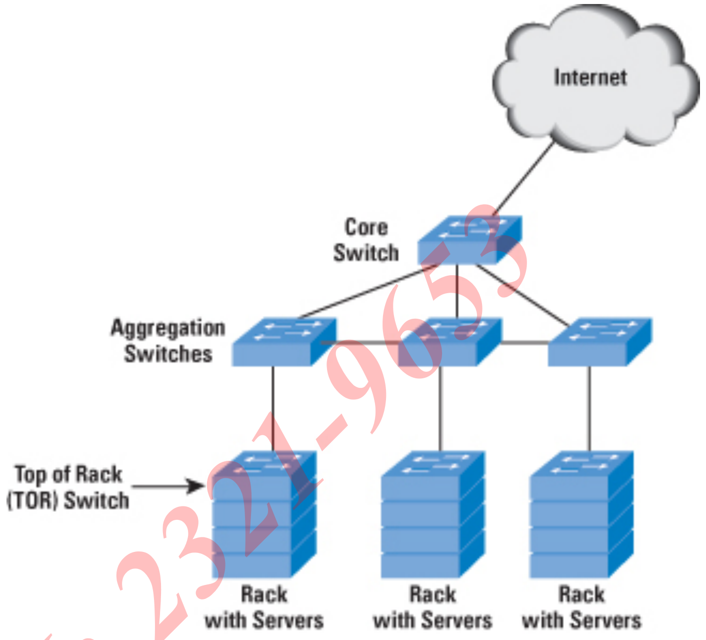

Figure 4: Example Data-Centre Switch Network Architecture

The presence of virtualized servers adds an extra dimension. Network connections to physical servers will need to involve "fatter pipes" because traffic for multiple VMs will be multiplexed onto the same physical Ethernet connection. This result is expected because busineses have effectively collapsed multiple physical servers into a single physical server with VMs. It is quite common to have servers with 10-Gbps Ethernet cards in this scenario.

IV. SUGGESTION AND CONCLUSIONS

4.1 New Protocol for Data Centre Networking

Numerous initiatives and standards regulatory bodies are addressing the standards related to cloud computing. From the networking side, the IEEE is working on new protocols and the enhancement of existing protocols for data centers. These enhancements are particularly useful in data centres with converged networks—the area is known as Convergence Enhanced Ethernet (CEE).

A previous section indicated the importance of FCoE for converged storage network environments. The IEEE is working to enable FCoE guarantees (because Fiber Channel is a reliable protocol as compared to best-effort Ethernet) through an Ethernet link in what is known as "Lossless Ethernet." FCoE is enabled through a Priority Flow Control (PFC) mechanism in the 802.1Qbb activities in the IEEE. In addition, draft IEEE 802.1Qau provides end-to-end congestion notification through a signaling mechanism propagating up to the ingress port, that is, the port connected to the server Network Interface Card (NIC). This feature is useful in a data-centre topology.

A third draft IEEE 802.1aq defines shortest-path bridging. This work is similar to the work being done in the IETF TRILL (Transparent Interconnect of Lots of Links) working group. The key motivation behind this work is the relatively flat nature of the data-centre topology and the requirement to forward packets across the shortest path between the endpoints (servers) to reduce latency, rather than a root bridge or priority mechanism normally used in the Spanning Tree Protocol (STP). The shortest-path bridging initiative in IEEE 802.1aq is an incremental advance to the Multiple Spanning Tree Protocol (MSTP), which uses the Intermediate System-to-Intermediate System (IS-IS) link-state protocol to share learned topologies between switches and to determine the shortest path between endpoints.

The fourth draft 802.1Qaz is also known as Enhanced Transmission Selection (ETS). It allows lower-priority traffic to burst and use the unused bandwidth from the higher-priority traffic queues, thus providing greater flexibility.

4.2 Virtualized Network Equipment Functions

Though cloud computing does not depend upon virtualization, several cloud infrastructures are built with virtualized servers. In an environment with physical servers, switches are used to connect servers to other servers. Firewalls and application-delivery controllers are other types of equipment that you can use in a data centre on the connection to external clients. With a virtualized environment, you can move some or all of these functions to reside inside a server.

When one studies the case of the software-based Virtual Switch as shown in Figure 5, one 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, which in turn connects to the physical Ethernet adapter on the server and to the external Ethernet switch.

To the network manager, the virtual switch can appear as a part of the network to the network manager. Unlike physical switches, the Virtual Switch does not necessarily have to run network protocols for its operation, nor does it need to treat all its ports the same because it knows that some of them are connected to virtual Ethernet ports. It can avoid destination address learning on the ports connected to the VMs. It can function through appropriate configuration from an external management entity.

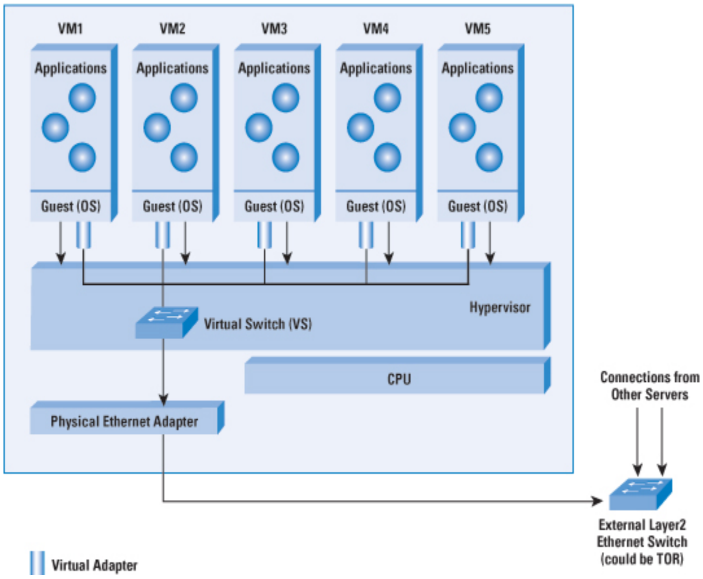

Figure 5: Virtual Ethernet Switch in a Virtualized Server Environment

It is possible to implement a virtualized firewall as a VM instead of a plug-in to the hypervisor. These VMs are self-contained, with an operating system along with the firewall software. The complete package is known as a firewall virtual appliance. These VMs can be loaded and configured so that network packets destined for any of the VMs pass through the firewall VM, where they are validated before being passed to the other VMs. Another use of the firewall VM is as a front end to the physical servers in the data centre. The disadvantage of a virtual appliance is the performance hit due to its implementation as a software function in a virtualized environment.

4.3 A strategically planning management

Management has several facets in a cloud-computing environment: billing, application-response monitoring, configuring network resources (virtual and physical), and workload migration. In a private cloud or tightly coupled environment, management of the applications may have to be shared between the internal cloud and the private cloud.

Businesses can manage cloud-computing environments in several ways, depending upon the specific area. They can manage the network equipment (physical and virtual) through the Simple Network Management Protocol (SNMP) and a network management console. In a virtualized environment, the virtualization vendor often offers a framework to manage and monitor VMsSeveral vendors offer products to act as management front ends for public clouds including Amazon, whose products act as brokers and management consoles for applications deployed over the Amazon cloud offering.It is clear that this area of management for cloud computing is still evolving and needs to be revised for a unified management view.

4.4 Common myths about Cloud Computing

This section outlines common myths about cloud computing.

- Myth1: Cloud computing should satisfy all the requirements specified: scalability, on demand, pay per use, resilience, multi-tenancy, and workload migration.

  In fact, 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. There is, however, value in trying to satisfy most of these requirements when businesses are building a cloud service.

- Myth 2: Cloud computing is useful only if you are outsourcing your IT functions to an external service provider.

  This is not true. You can use cloud computing in your own IT department for on-demand, scalable, and pay-per-use deployments. Several vendors offer software tools that you can use to build clouds within your enterprise's own data centre.

- Myth 3: Cloud computing requires virtualization.

  Although virtualization brings some benefits to cloud computing, including aspects such as efficient use of servers and workload migration, it is not a requirement for cloud computing. However, virtualization will experience increased usage in cloud deployments.

- Myth 4: Cloud computing requires businesses to expose their data to the outside world.

  With internal clouds you will never need to expose your data to the outside world. If data security and privacy are concerns, you can develop a cloud model where web front ends are in the cloud and back-end data always resides in your company's premises.

- Myth 5: Converged networks are essential to cloud computing.

  Although converged networks (with FCoE, for example) have benefits and will see increased adoption in data centers in the future, cloud computing is possible without converged networks. In fact, some cloud vendors use only Fibre Channel for all their storage needs today. Use of converged networks in the future will result in cost efficiencies, but it is not a requirement today.

4.5 Cloud Computing: Issues and Challenges

Cloud-computing technology is still evolving. Various companies, standards bodies, and alliances are addressing several remaining issues and challenges. These include :

Security: Security is a significant issue for enterprise IT managers when they consider using a cloud service provider. Physical security through isolation is a critical requirement for private clouds, but not all cloud users need this level of investment. For those users, the cloud provider must guarantee data isolation and application security (and availability) through isolation across multiple tenants. In addition, authentication and authorization of cloud users and encryption of the "network pipe" from the cloud user to the service provider application are other factors to be considered.

Network concerns: When cloud bursting is involved, should the servers in the cloud be on the same Layer 2 network as the servers in the enterprise? Should a Layer 3 topology be involved because the cloud servers are on a network outside the enterprise? How would this work across multiple cloud data centers?

Cloud-to-cloud and Federation concerns: Study a case where an enterprise uses two separate cloud service providers. Compute and storage resource sharing along with common authentication (or migration of authentication information) are some of the problems with having the clouds "interoperate." For virtualized cloud services, VM migration is another factor to be considered in federation.

Legal and regulatory concerns: These factors become important especially in those 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.

This article introduced the still-evolving area of cloud computing, including the technologies and some deployment concerns. Definitions and standardization in this area are a work in progress, but there is clear value in cloud computing as a solution for several IT requirements.

REFERENCES

1. "The Wisdom of Clouds," James Urquhart's blog on Cloud Computing, http://news.cnet.com/the-wisdom-of-clouds/Accessed on (04.06.2014)

2. "A New Approach to Network Design When You Are in the Cloud," The Lippis Report,121,http://lippisreport.com/2009/03/a-new-approach-to-network-design-in-the-cloud/ Accessed on (04.06.2014)

3. “What cloud computing really means?” http://www.infoworld.com/d/cloud computing/what-cloud-computing-really-means-031. Accessed on (23.06.2014)

4.”How a cloud computing works?” http://computer.howstuffworks.com/cloudcomputing/cloud-computing.htm. Accessed on (23.06.2014)

5. “What is cloud computing?” http://www.pcmag.com/article2/0,2817,2372163, 00.asp/Accessed on (25.06.2014)