The anatomy of cloud computing information technology essay

Technology, Information Technology



Cloud computing is changing the way how hardware and software are provided for on-demand capacity fulfillment. Lately there are ways for ondemand servers, storage and CDNs. These are changing the way in developing web applications and make business decisions. In 1943 Thomas J. Watson of IBM famously proclaimed " there is a world market for computers, maybe five." Today this proclamation is silly, but the statement really did hold up for approximately 10 years. Into the 1950s IBM designed computers for a possible market of 20 companies of which 5 were expected to purchase such a machine. In 1953 IBM was pleasantly surprised to find 18 of 20 companies purchased the IBM 701, provide the business of back office processing and a new division for the tabulating giant. Hardware and data operations are again consolidating towards major players. These specialist providers are building at a scale and are specialized in most web businesses. On-demand infrastructure of the cloud makes it cheaper and more efficient. Microsoft and Google are the newest entrants into the cloud computing arena. Microsoft' s Windows Azure services platform will likely be the best platform for C# and ASP. Net development as it is tuned by the creators of . Net, IIS and SQL Server. Google has similarly applied its expertise in the Python language and distributed web nodes to its Google App Engine product. The App Engine cloud is tuned by top contributors to the Python language. App Engine utilizes custom Google software, Google Front End and Megastore, for web serving and storage. Cloud developers on either platform are using a similar set of hardware and software as the proven web-scale platforms of Live. com and Google. In future there will be support for java in Google App Engine, their second major language and the most popular

language among Google's own services. Amazon's EC2 is the most wellknown cloud computing provider. The Amazon Machine Image (AMI), a machine image formatted deployment in the Amazon cloud is the basic building block of EC2 virtualization and the primary interaction point of Amazon's customers. Amazon resells premium operating system and application packages on behalf of companies such as Microsoft, IBM and Oracle, but it's possible such specializations will instead be absorbed by the software publishers themselves as they roll out their own hosted clouds. The cloud computing software stack is trending towards an integrated, managed experience maintained by some of the top contributors to each programming language and related components. More generic cloud platforms will need to stay up-to-date with managed technologies on their platform and/or establish a strong reseller relationship to more specialized cloud managers.

3. 1. 2Managed cloud stack

A description...

Fig 16: Managed Cloud Stack

Managed cloud providers handle an entire stack of infrastructure needed to deliver web applications at scale. Figure 16 shows the managed cloud stack consisting of two parts, first part consists of cache, dynamic code and attached storage and logic for the dynamic code, second part consists of stable and efficient OS, security features and business logic written using some programming language. A solid cloud computing environment abstracts the basics of a computing environment away from the implementers and lets them focus on adding value with each new

application. Managed cloud hosting providers need to offer the following basic layers to stay relevant in a web developer's world. Every managed cloud platform includes a dynamic language virtual machine and an appropriate web services gateway. Language functions too closely associated with the parent operating system and its libraries are stripped away, leaving only a pure operating environment for a machine interpreter. External dependencies such as GNU tools and custom compilers will not function within the cloud language abstraction layer. Cloud services bundle a dynamic language runtime into an easily spawned instance for standard and efficient interpretation across many application instances. Google App Engine supports most functions of the Python language with additional support for the Django framework, WebOb and PyYAML. Developers may replace these built-in libraries with newer or customized versions at an additional performance and usage cost. App Engine passes web requests into the programming language environment through the Web Server Gateway Interface. Dynamic applications persist in their application state and logic through database and file storage. In the cloud world the database and the file server are cloud services within themselves, operating in an isolated and specialized layer. This isolation makes the storage layer swappable from the rest of the cloud stack and presents new opportunities for competition. Static files fall into two major categories based on their planned consumption. Files less than 1 MB in size can be consumed by most clients in a single request, matching the expected simple request/response model of the platform. File over 1 MB in size need to be broken into parts, for a sequenced download. Static cloud storage can be broken up into differing

solutions by file size or file type, providing the best possible solution for the storage and delivery task at hand. Google App Engine offers static file storage separate from its dynamic runtime. App Engine supports up to 1, 000 files and has a limit of 10 MB HTTP response. Amazon Web Services offers static file serving through its Simple Storage Service (S3) origin server and CloudFront CDN services. Amazon allows private and public file storage and can even charge individual user of third-party services for their use through DevPay.

3.1.3Cloud consumers

The target market of a cloud computing platform will affect its stack completeness, feature sets and future support. Cloud terminology seems to be thrown around as a magical buzzword but there are major usage cases emerging. Web application developersNew web applications start small and may sometimes experience exponential growth on a worldwide basis. Web developers evaluating the cloud stack are likely starting from scratch without the concerns of switching from a legacy system or alternate implementation. Web developers prefer a cloud stack for fast web performance.

Geographically distributed dynamic instances are important at least as an upgrade option to protect a new business from a rewrite at varying levels of scale. Enterprise applications are moving out of the local server closet and into the cloud. Medium-to large-sized companies are replacing in-house maintenance of machines and applications with software and infrastructure as a service. Project management, employee tracking, payroll and many other common functions have made their way into the software-as-a-service

realm. More customized applications will migrate to cloud hosting and take their place alongside the anchor tenants of the groupware and collaboration suites. Windows Azure, Salesforce' s Force. com and Google App Engine show strong promise as integrated back office add-ons. Microsoft and Google already have a solid footing in enterprise groupware services through Exchange Online and Google Apps respectively. Force, com can be closely tied to the popular Salesforce CRM application for sales and marketing teams. More generic back office functions can operate on any cloud hosting provider with a properly maintained disk image. A new class of hosting provider operates as an abstraction layer between multiple clouds by maintaining the appropriate images and deployment scripts for any given task. Companies such as Aptana, CohesiveFT, RightScale and many others span multiple cloud hosting providers with a single management interface. Cloud management companies can monitor multiple providers and create spot pricing market for computing resources. Back office solutions represent the largest possible growth area for cloud hosting providers. Platforms with strong existing anchor tenants can add on new services combining softwareas-a-service and infrastructure-as-a-service. Generic cloud hosting providers will likely be tapped for general tasks directly or through a cloud management layer.

3. 2 Virtual Infrastructures

Today, the usage of the Internet is fundamentally changing. Internet services are constructing data centers of unprecedented scale to offer a large diversity of cloud services for research, data mining, email hosting, maps

Page 7 and other features. This evolution leads to the convergence of

communication and computation and portrays a new vision of the services that the Internet can bring to users. According to this concept, the Internet will not remain " only" a huge shared and unreliable communication facility between edge hosts enabling real time contact and data exchanges. Instead, it will become a world-wide reservoir of interconnected resources that can be shared and reserved. We envision the Internet will increasingly embed and expose its computational and storage resources, as well as its communication and interconnection capacities, to meet the requirements of emerging applications. Large-scale experimental facilities are pre-figuring this new way of sharing IT and computing resources and highlight the need for on-demand customizable infrastructures. Indeed, many computer science projects in network or distributed systems require experiments with modified operating systems and communication protocols exposed to realistic and reproducible conditions. Computer scientists need to perform distributed experiments that run on many sites at the same time. Generally the experiments are interactive and large-scale, they run on many nodes, but for a relatively short time (a few hours). This raises the need for time-limited access to customized experimental platforms. As an example, PlanetLab allows researchers to run experiments on a large scale under real- world conditions. Using distributed virtualization, every user can allocate a slice of PlanetLab's network-wide hardware resources for experiments in file sharing and network-embedded storage, content-distribution networks, routing and multicast overlays, network measurement tools, etc. Grid'5000, another experimental facility, gathers large scale clusters and gives access to 5000

CPUs distributed over 9 sites and inter-connected by 10 Gbps-dedicated networks. Grid'5000 provides a deep re-configuration mechanism allowing researchers to deploy, install, boot and run their specific software images, possibly including all the layers of the software stack. Virtualization enables an efficient separation between services or applications and physical resources. For example, the virtual machine paradigm is becoming a key feature of servers, distributed systems and grids as it provides a powerful abstraction. It has the potential to simplify the management of resources and to offer a great flexibility in resource usage. Each Virtual Machine a) provides a confined environment where non-trusted applications can be run, b) allows establishing limits in hardware-resource access and usage, through isolation techniques, c) allows adapting the runtime environment to the application instead of porting the application to the runtime environment (this enhances application portability), d) allows using dedicated or optimized OS mechanisms (scheduler, virtual memory management, network protocol) for each application, e) allows the applications and processes running within a VM to be managed as a whole. Extending these properties to the service level through the concept of " infrastructure as a service", the abstraction of the hardware enables the creation of multiple, isolated and protected virtual aggregates on the same set of physical resources by sharing them in time and space. In other words, with representation in VMs, it is possible that a physical resource (node) hosts VMs of different virtual infrastructures. The virtual infrastructures are logically isolated by virtualization and can provide customized services to each virtual infrastructure, for example in terms of bandwidth provisioning, channel

encryption, addressing protocol version. The isolation also provides a high security level for each infrastructure.

3.3 CPU Virtualization

Virtualizing a CPU is to some extent very easy. A process runs with exclusive use of it for a while and is then interrupted. The CPU state is then saved and another process runs. After a while, this process is repeated. This process typically occurs every 10ms or so in a modern operating system. It is worth nothing, however, that the virtual CPU and the physical CPU are not identical. When the operating system is running, swapping processes, the CPU runs in a privileged mode. This allows certain operations, such as access to memory by physical address, that are not usually permitted. For a CPU to be completely virtualized, there are some set of requirements, there are: Privileged instructions are defined as those that may execute in a privileged mode, but will trap if executed outside this mode. Control sensitive instructions are those that attempt to change the configuration of resources in the system, such as updating virtual to physical memory mappings, communicating with devices, or manipulating global configuration register. Behavior sensitive instructions are those that behave in a different way depending on the configuration of resources, including all load and store operations that act on virtual memory. In order for architecture to be virtualizable, all sensitive instructions must also be privileged instructions. Intuitively, this means that a hyper-visor must be able to intercept any instructions that change the state of the machine in a way that impacts other processes. CPU virtualization involves a single CPU acting, as if it were

two separate CPUs. In effect, this is like running two separate computers on a single physical machine. Perhaps the most common reason for doing this is to run two different operating systems on one machine. The CPU, or central processing unit, is arguably the most important component of the computer. It is the part of the computer which physically carries out the instructions of the applications which run on the computer. The CPU is often known simply as a chip or microchip. The way in which the CPU interacts with applications is determined by the computer's operating system. The best known operating systems are Microsoft Windows, Mac OS and various open-source systems under the Linux banner. In principle a CPU can only operate one operating system at a time. It is possible to install more than one system on a computer's hard drive, but normally only one can be running at a time. The aim of CPU virtualization is to make a CPU run in the same way that two separate CPUs would run. A very simplified explanation of how this is done is that virtualization software is set up in a way that it alone, communicates directly with the CPU. Everything else which happens on the computer passes through the software. The software then splits its communications with the rest of the computer as if it were connected to two different CPUs. One use of CPU virtualization is to allow two different operating systems to run at once. As an example, an Apple computer could use virtualization to run a version of Windows as well, allowing the user to run Windows only applications. Similarly a Linux-based computer could run Windows through virtualization. It's also possible to use CPU virtualization to run Windows on a Mac or Linux PC, or to run Mac OS and Linux at the same time. CPU virtualization should not be confused with multitasking or hyper-threading.

Multitasking is simply the act of running more than one application at a time. Every modern operating system allows this to be done on a single CPU, though technically only one application is dealt with at any particular moment. Hyper-threading is where compatible CPUs can run specially written applications in a way that carries out two actions at the same time.

3. 3. 1Working of CPU Virtualization

CPU virtualization can only be done on processors with virtualization capability. This technology also has its set of extra instructions which is known as the Virtual Machine Extensions or the VMX. This provides 10 specific instructions to the CPU for virtualization. These are the VMPTRLD, VMCLEAR, VMPTRST, VMWRITE, VMREAD, VMLAUCH, VMCALL, VMXON, VMXOFF and VMRESUME. Aside from these instructions. CPU virtualization also has two models to be used when running. One is the root operation and the other is the non-root operation. Under the root operation, it is usually the virtualization controlling software that runs over it. Meanwhile, for the nonroot operation, it is the virtual machine that runs over it. The operating system as well as software, called as " guest software," can run on top of the virtual machine. The VMXON instruction should be executed in order to enter the CPU virtualization mode. This will call the VMX software which will enter all the virtual machines through the VMLAUNCH instruction. Exiting these virtual machines can be done with the VMRESUME instruction. But if the VMX needs to shut down or exit the virtualization mode, it will have to call the VMXOFF instruction.

3. 4 Network and Storage Virtualization

3. 4. 1Network Virtualization

Network virtualization provides a powerful way to run multiple networks, each customized to a specific purpose, at the same time over a shared substrate. Network virtualization focuses on two main scenarios. First, consider the role of virtualization in running multiple experiments simultaneously in a shared experimental facility. The VINI project (radio interview) is a step in that direction, supporting experimentation with new routing, forwarding and addressing schemes on a shared facility built on top on general-purpose processors. Second, consider the role of virtualization to support multiple architectures simultaneously as a long-term solution for the future Internet. The Cabo project explores the benefits of running customized architectures, as well as how a virtualized system enables an economic refactoring of a future Internet into infrastructure providers and service providers. 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 to a particular server or device in real time. Each channel is independently secured. Every subscriber has shared access to all the resources on the network from a single computer. Network management can be a tedious and time-consuming business for a human administrator. Network virtualization is intended to improve productivity, efficiency and job satisfaction of the administrator by performing many of these tasks automatically, thereby disguising the true complexity of the network. Files, images, programs and folders can be centrally managed from a single

physical site. Storage media such as hard drives and tape drives can be easily added or reassigned. Storage space can be shared or reallocated among the servers. Network virtualization is categorized as either external, combining many networks, or parts of networks, into a virtual unit, or internal, providing network-like functionality to the software containers on a single system. Whether virtualization is internal or external depends on the implementation provided by vendors that support the technology. External network virtualizationSome vendors offer external network virtualization. in which one or more local networks are combined or subdivided into virtual networks, with the goal of improving the efficiency of a large corporate network or data center. The key components of an external virtual network are the VLAN and the network switch. Using VLAN and switch technology, the system administrator can configure systems physically attached to the same local network into different virtual networks. Conversely, VLAN technology enables the system administrator to combine systems on separate local networks into a VLAN spanning the segments of a large corporate network. Internal network virtualizationOther vendors offer internal network virtualization. Here a single system is configured with containers, such as the Xen domain, combined with hypervisor control programs or pseudointerfaces such as the VNIC, to create a " network in a box." This solution improves overall efficiency of a single system by isolating applications to separate containers and/or pseudo interfaces. 3. 4. 1. 1Components of a virtual networkVarious equipment and software vendors offer network virtualization by combining any of the following: Network hardware, such as switches and network adapters, also known as network interface cards

(NICs)Network elements such as Firewalls, Load BalancersNetworks, such as virtual LANs (VLANs) and containers such as virtual machines (VMs) and Solaris ContainersNetwork storage devicesNetwork Mobile elements such as Laptops, Tablets and Cell PhonesNetwork media, such as Ethernet and Fibre Channel

3. 4. 2Storage Virtualization

Storage virtualization is a concept and term used within computer science. Specifically, storage systems may use virtualization concepts as a tool to enable better functionality and more advanced features within the storage system. Broadly speaking, a 'storage system' is also known as a storage array or disk array. Storage systems typically use special hardware and software along with disk drives in order to provide very fast and reliable storage for computing and data processing. Storage systems are complex and may be thought of as a special purpose computer designed to provide storage capacity along with advanced data protection features. Disk drives are only one element within a storage system, along with hardware and special purpose embedded software within the system. Storage systems can provide either block accessed storage, or file accessed storage. Block access is typically delivered over Fibre Channel, iSCSI, SAS, FICON or other protocols. File access is often provided using NFS or CIFS protocols. Within the context of a storage system, there are two primary types of virtualization that can occur, i) Block Virtualization and ii) File Virtualization. Block virtualization used in this context refers to the abstraction (separation) of logical storage from physical storage so that it may be accessed without

regard to physical storage or heterogeneous structure. This separation allows the administrators of the storage system greater flexibility in how they manage storage for end users. File virtualization addresses the NAS challenges by eliminating the dependencies between the data accessed at the file level and the location where the files are physically stored. This provides opportunities to optimize storage use and server consolidation and to perform non-disruptive file migrations.

3. 4. 2. 1Block virtualization

Virtualization of storage helps to achieve location independence by abstracting the physical location of the data. The virtualization system presents to the user a logical space for data storage and handles the process of mapping it to the actual physical location. It is possible to have multiple layers of virtualization or mapping. It is then possible that the output of one layer of virtualization can then be used as the input for a higher layer of virtualization. Virtualization maps space between back-end resources, to front-end resources.

3. 4. 2. 2Benefits

Non-disruptive data migrationOne of the major benefits of abstracting the host or server from the actual storage is the ability to migrate data while maintaining concurrent I/O access. The host only knows about the logical disk and so any changes to the meta-data mapping is transparent to the host. This means the actual data can be moved or replicated to another physical location without affecting the operation of any client. When the data has been copied or moved, the meta-data can simply be updated to point to

the new location, therefore freeing up the physical storage at the old location. There are many day to day tasks a storage administrator has to perform that can be simply and concurrently performed using data migration techniques. Moving data off an over-utilized storage device. Moving data onto a faster storage device as needs requireImplementing an Information Lifecycle Management policyMigrating data off older storage devices (either being scrapped or off-lease)Improved utilizationUtilization can be increased by virtue of the pooling, migration and thin provisioning services. When all available storage capacity is pooled, system administrators no longer have to search for disks that have free space to allocate to a particular host or server. A new logical disk can be simply allocated from the available pool, or an existing disk can be expanded. Pooling also means that all the available storage capacity can potentially be used. In a traditional environment, an entire disk would be mapped to a host. This may be larger than is required, thus wasting space. In a virtual environment, the logical disk is assigned the capacity required by the using host. Storage can be assigned where it is needed at that point in time, reducing the need to guess how much a given host will need in the future. Using thin provisioning, the administrator can create a very large thin provisioned logical disk, thus the system in use thinks that it has a very large disk from day 1.

3. 4. 2. 3Risks & Complexity

Risks in Block Virtualization are: Backing out a failed implementation: Once the abstraction layer is in place, only the virtualizer knows where the data actually resides on the physical medium. Backing out of a virtual storage

environment therefore requires the reconstruction of the logical disks as contiguous disks that can be used in a traditional manner. Most implementations will provide some form of back-out procedure and with the data migration services it is at least possible, but time consuming. Interoperability and vendor support: Interoperability is a key enabler to any virtualization software or device. It applies to the actual physical storage controllers and the hosts, their operating systems, multi-pathing software and connectivity hardware. Interoperability requirements differ based on the implementation chosen. For example virtualization implemented within a storage controller adds no extra overhead to host based interoperability, but will require additional support of other storage controllers if they are to be virtualized by the same software. Switch based virtualization may not require specific host interoperability, if it uses packet cracking techniques to redirect the I/O. Network based appliances have the highest level of interoperability requirements as they have to interoperate with all devices, storage and hosts. Complexity affects several areas such as: Management of environment: Although a virtual storage infrastructure benefits from a single point of logical disk and replication service management, the physical storage must still be managed. Problem determination and fault isolation can also become complex, due to the abstraction layer. Infrastructure design: Traditional design ethics may no longer apply as virtualization brings a whole range of new ideas and concepts. The software or device itself: Some implementations are more complex to design and code - network based, especially in-band (symmetric) designs in particular, these implementations actually handle the I/O requests and so latency becomes an issue.

Performance and scalability: In some implementations the performance of the physical storage can actually be improved, mainly due to caching. Caching however requires the visibility of the data contained within the I/O request and so is limited to in-band and symmetric virtualization software and devices. However these implementations also directly influence the latency of an I/O request (cache miss), due to the I/O having to flow through the software or device. Assuming the software or device is efficiently designed this impact should be minimal when compared with the latency associated with physical disk accesses. Due to the nature of virtualization, the mapping of logical to physical requires some processing power and lookup tables. Therefore every implementation will add some small amount of latency.

3. 4. 3Using virtualization to achieve green data centers

Green IT and Virtualization are two of the hottest topics in IT and will change the current business landscape for solution providers and value-added resellers (VARs). Customers are asking about green data centers and solution providers are interested in learning more about virtualization to help meet this goal. But depending on the type of solution these providers offer, be it hardware or services, virtualization will affect them in different ways. Impact of virtualization on resellers and solution providersIt affects different resellers differently. For example, the hardware resellers, making the most public noise about being green, probably have the most to lose from the wave of virtualization that's hitting us. With server virtualization, you could possibly get from between 8 to 15 virtual systems on one physical server, so

subsequently you're going to see reduction in the hardware that's being sold. Though the hardware might be different and it might require some upgrades, the overall idea is to reduce the hardware, reduce the power consumption, the cooling, the footprints and those other pieces. At the same time, virtualization of the desktops is also popular. Thin clients are coming back and you might see thin systems replacing heavier desktop systems, so there's an uptick in sales opportunities for hardware vendors there. From a services perspective, there's a tremendous opportunity to come in and work on the virtualization pieces as part of an organization's transformation. So all have to make the plan and analyze pieces of those phases. Some people will just focus on the execution of the plan and then leave and others will attempt to do the virtualization process with the integration pieces and keep the service hosted elsewhere or on-site. Impact on cloud computing factorCloud computing will become the delivery mechanism for a lot of these services. The hardware used to be owned by the firm or the organization that's using the service. The hardware will still be sold, but it will be sold to the cloud provider. The services themselves will have a slightly different model. There'll be some integration work and the usual customization work, but if the cloud providers are doing what they promise to do, there'll be no more interaction. The upgrades, the maintenance, the optimization and tuning of those systems should be done automatically by the cloud providers, so you could see a reduction in service provider bookings because of the cloud. The biggest obstacle with green data centers is changing the hearts and minds of the humans that run those organizations, changing human behavior to understand that there's a new model, to accept it, to

embrace it, to trust it and to go with it. The reality is these humans who want to have physical proximity to their systems don't go to their systems anyway they just want the ability to go there if they needed to. Once you get folks comfortable with the concepts of virtualization and that what they're really buying is a processing service and then they can start to distance themselves from the data center. Costs and benefits for customers for green data centersThere are a number of categories they can save on, not the least of which is power and cooling for those systems that they've removed. There are also space considerations. In a virtual world, with the proper tools, they can cover more systems and the administrative costs tend to go down. So you've got power, people, physical space being reduced as well as the hardware costs. At the end of the day, when you add all those together, you can get very strong Return Of Interest (ROI) to swap out a data center and in 12 to 16 months, it will completely pay for itself. Virtualization is a rapidly expanding core technology for companies today, seeking lower operating costs and higher asset utilization. However, successful deployments require more than just consolidating a bunch of servers and then treating everything else the same. The flexibility virtualization delivers also creates many pitfalls that can eliminate all of the ROI from a deployment. If your organization is in implementing virtualization, then the following steps have to be followed: Virtualization is for more than just server consolidation. Virtualization is an extremely powerful technology that can improve the efficiency and operations of far more than just servers. Consider broadening your use of virtualization after a good understanding of server virtualization is in place. Ensure management and visibility products are in place before broadening

your use of virtualization. This is the only way to effectively scale virtualization without falling victim to the many pitfalls associated with this technology. Choose management solutions that were designed for virtual environments. The management tools that work well in a client server environment are not the ones used to manage mainframes. Similarly, the management tools used to manage virtual environments is different from the ones you have in place today.

Summary

The material in this part lays a strong foundation about virtualization. Virtualization seems to defy common sense and it can be hard to understand why there's a rush of interest surrounding it. Virtualization is the practice of running several operating systems at the same time on a single computer. Virtualization is spreading beyond servers and desktop computers and onto mobile phones, where a user can have a " work" OS alongside a " home" OS and can flip between them as need to be. Apart from this broad introduction to virtualization, this part also deals with pitfalls of virtualization (Chapter 1), how virtualization can be used in clusters and grid computing (Chapter 2) and anatomy of cloud infrastructure, how cloud computing can be benefited using virtualization (Chapter 3).