TECHNOLOGY WHITE PAPER:

CLOUD COMPUTING – *Transforming IT Enterprises through advances in Servers, Networks, and Storage.*

Chai Toh, Sr. Director & Chief Technology Advisor, ALICO Systems Inc., Torrance CA USA (Aug 2013)

WHAT AND WHY WE NEED CLOUD?

Welcome to the era of Cloud Computing. It is an era where Information Technology functions and operations will be transformed to a different level, using a different architecture and embracing new technologies.

Information technology infrastructures are found in most enterprises today. It is uncompetitive and unproductive to operate an enterprise without utilizing the automation, speed, and interactivity provided by information technology. This is evident since the days we invented the microprocessor, where computationally tedious tasks can be executed accurately, quickly, and repeatedly.

Today's information technology resources fall under 3 broad categories: (a) servers – computers, (b) storage – disk, memories, flash, etc., and (c) networks. Currently, each enterprise implements its I.T. infrastructure from scratch, which is a substantial investment for the organization. The CIO has to request a budget to implement the IT infrastructure, ensure its smooth operation, maintain it round-the-clock, perform upgrades of both software and hardware, and deal with security and scalability issues when business needs grow and changes.

So, it is clear that each enterprise (be it big or small) have to implement their own IT infrastructures. Collectively, if we group all these IT infrastructures together and bundle them as a "resource unit" – this resembles a "cloud" of immerse computational power, networked together, with substantial storage resources. However, this "virtual bundle of IT infrastructures" is not realizable because "sharing" is not permitted across enterprises. There is no agreement, let alone security issues, protection of sensitive company data, and issues of IT investments and ownership.

However, if one is able to harness this collection of IT infrastructure and computing power into one or a few consolidation points (so called cloud), then we can eliminate huge replication of IT infrastructures across millions of enterprises and save millions or billions of dollars to run IT operations. See Fig 1.

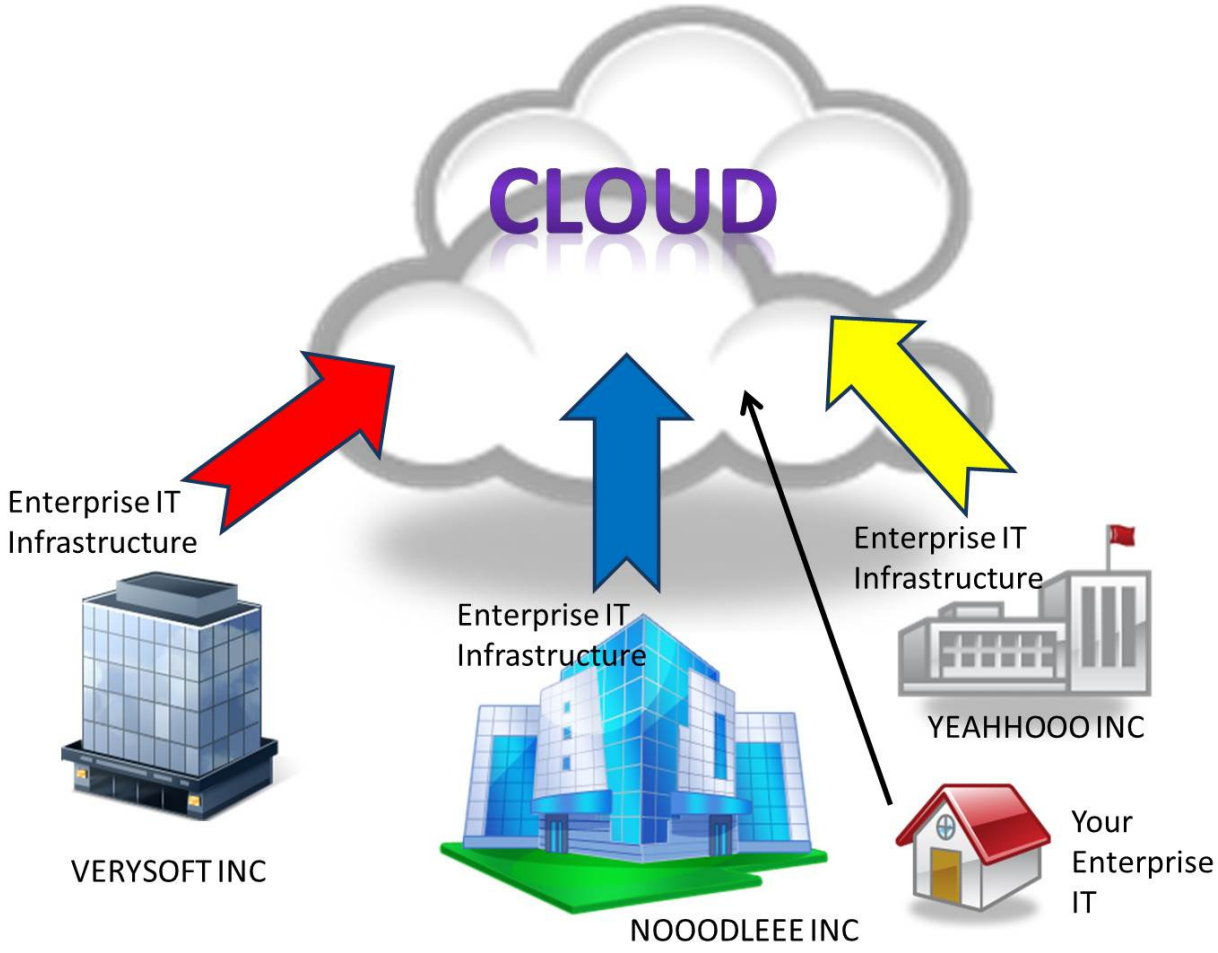

Fig 1: Imagine consolidating all the Enterprise IT power into a Cloud.

Cloud computing was born primarily to:

- Provide resources (compute, storage, networks) based on-demand (pay as you go)
- Provide users ease of usage through user-friendly web interface
- Provide users free from pain and risks of IT investment, installation, and maintenance
- Provide developers quick environments to develop applications without an in-house IT
- Avoid replication of IT infrastructure for the same enterprise with different branches
- Provide superior performance without huge IT investments (especially useful for SME)
- Provide scalability when business needs changes or grow
- Provide location independence – resources can be accessed anywhere via the internet
- Provide reliability – no single point of failures through redundancy and disaster recovery
- Provide choice – users can chose a cloud provider based on its preference – no monopoly

A cloud is essentially ***a group of IT resources*** that is centrally managed by a dedicated provider that install, maintain, upgrade, and lease its resources to external users. The resources here include servers (web servers, ERP servers, CRM servers, Email servers, etc.), storage (disks, memories, etc.), and networks (VPNs, VLANs, etc.). Virtualization technology can be used to provide Virtual Machines, allowing the support of multiple operating systems running different applications.

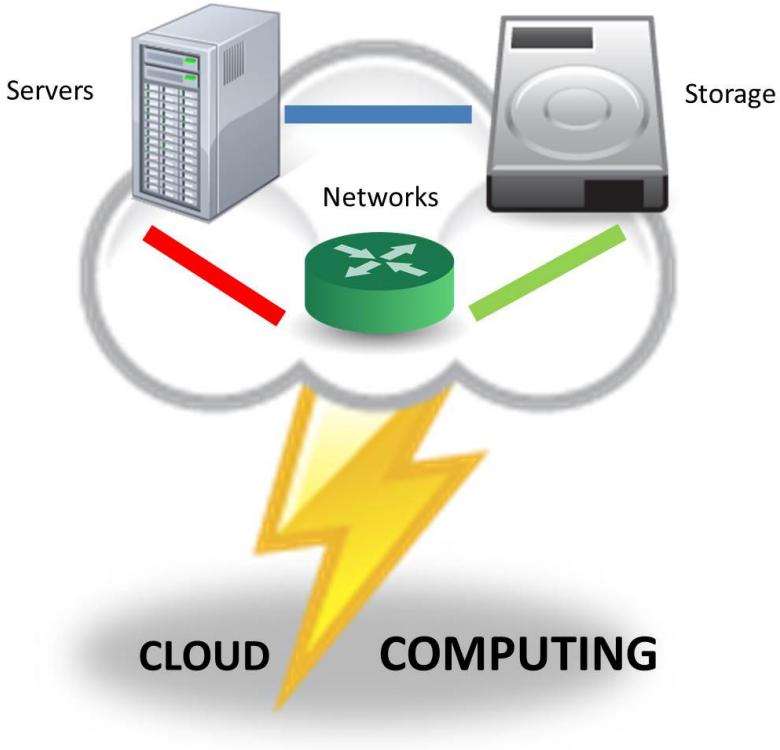

Fig 2: Cloud is chosen rather than the Sky.

Cloud is above us all, and it symbolizes an assembly point of IT resources.

A **Cloud System** must have the ability to perform backup of client data, provide high speed reliable networks, accurate accounting and billing, security (firewalls, IDS, tenant isolation, etc.), user-friendly client web access, round-the-clock technical support, recovery from crashes, provide 100% availability, replication, application acceleration, and load balancing.

Currently, there exist different types of cloud: (1) Private Cloud, (2) Public Cloud, (3) Enterprise Cloud, (4) Hybrid Cloud, and (5) Micro Cloud. Before discussing these various clouds, let's try to understand what makes cloud computing possible at the present time. In particular, a few key technologies that are part of the cloud ecosystem made cloud computing possible.

ADVANCED TECHNOLOGIES ENABLING THE CLOUD

Cloud uses several advanced technologies in existence today, namely – the **innovations found in servers, storage, and networks.** Virtualization is a "computing power" abstraction, creating a software entity capable of fully functioning like the physical entity while at the same time is dynamic, controllable, and scalable. Virtualization (for computing) is possible because of the advances made in hardware (CPU MIPS, multi-core processors, computer clusters, memory capacity, speed of data access, disk capacity, etc.) and software (hypervisors, modular operating systems, APIs, libraries, etc.).

SERVER Virtualization

- **VMs – Virtual Machines** are software abstractions of a physical computer, with access to CPU computation capability, memories, I/O, networks, etc. It is the hypothetical computer that is an independent entity – dynamic, powerful, and transparent enough to execute software applications on it. Hence, from a common hardware infrastructure, multiple VMs can be executed with each VM running different or same OS, executing different applications. VMs provide the multi-machine multi-user experience, coupled with versatility and scalability.

- **Hypervisors –** HYP – also known **as *Virtual Machine Monitor (VMM)*** creates and executes VMs. Two types of hypervisors exist. In hypervisor type 1, the hypervisor runs directly on top of the hardware to directly access the hardware and control guest operating systems which run above it. Hence, the hypervisor is the manager and controller of multiple guest OSs. Examples are Xenserver, VMware ESX, KVM, Microsoft Hyper-V, etc. Type 1 hypervisor is also known as native hypervisor or "bare metal".

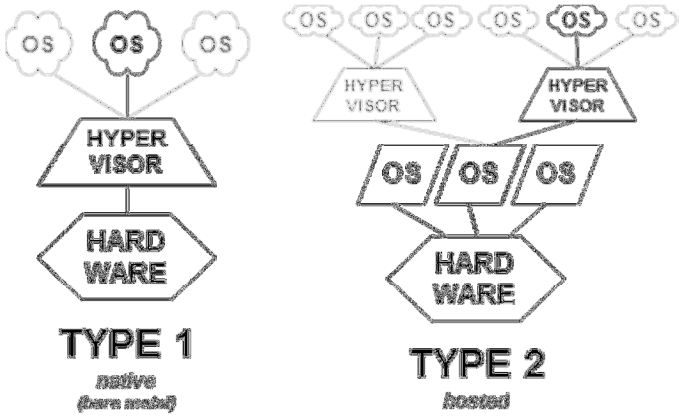

Fig 3: Hypervisors: Native versus Hosted. (source wiki)

Type 2 hypervisors run on top of a "base OS" (Fig 3) which in turn, supports the execution of multiple guest OSs and multiple VMs. Multiple base OSs and multiple hypervisors can exist. Examples are VMware workstation and VirtualBox hypervisors .

ADVANCES in CPUs

- Hypervisors do need powerful hardware resources before it is capable of supporting multiple VMs and OSs. INTEL and AMD have both produced specialized chips INTEL-VTx and AMD-V, exploiting advanced memory management (such as extended page tables and page table virtualization), privileged instructions (Virtual Machine Extensions - VMX), wider bus (64 bits and more), and rapid virtualization indexing. Hence, severs today are constructed using these advanced multicore virtualization-enabled CPUs.

NETWORK Virtualization

- **VPN – Virtual Private Networks** have been around for some time. Although not a new technology, it is a very useful technology as it creates a private network over the public internet securely through virtual point-to-point connections.

- **NFV – Network Function Virtualization** refers to the creation of software-based network functions that can be instantiated from anywhere without the need to install new hardware. This is possible because more and more functions for networking can be implemented in software, be it routing, naming, addressing, etc. In 2012, an ETSI industry specification group for NFV was created to address this area.

- **NV – Network Virtualization** commonly refer to virtual networks, including Virtual LANs. VLANs allow the creation of several different networks on the same physical local area networks, improving the efficiency of large corporate networks.

- **Network-in-a-Box –** VLANs are viewed as a form of "external network virtualization". Network-in-a-Box achieves internal network virtualization by providing containers for multiple VMs to share resources and exchange data, all on a single machine. This network-in-a-box concept is analogous to a tightly coupled system and that of a motherboard bus where all devices interconnect.

IMPROVED Storage

Storage technologies have also advanced considerably since the age of floppy disk and SCSI drives. While flash memories and solid state drives are now in use, large capacity drives capable of storing PetaBytes of data require new approaches. Three storage technologies worth mentioning are NAS, SAN, and RAID.

- **NAS – Network Attached Storage** is a file system technology where the disk storage appears as a networked element accessible by others in the network. Hence, the disk device can have an

Ethernet connection to Ethernet switches and to the rest of the computer network. NAS operates as a file server and a disk no longer needs to be tightly coupled to the CPU within a computer box. The file server now is a network node. NAS supports distributed file sharing using protocols such as NFS (Network File Systems), and CIFS (Common Internet File Systems). DELL, HP, and NETApp sell NAS products.

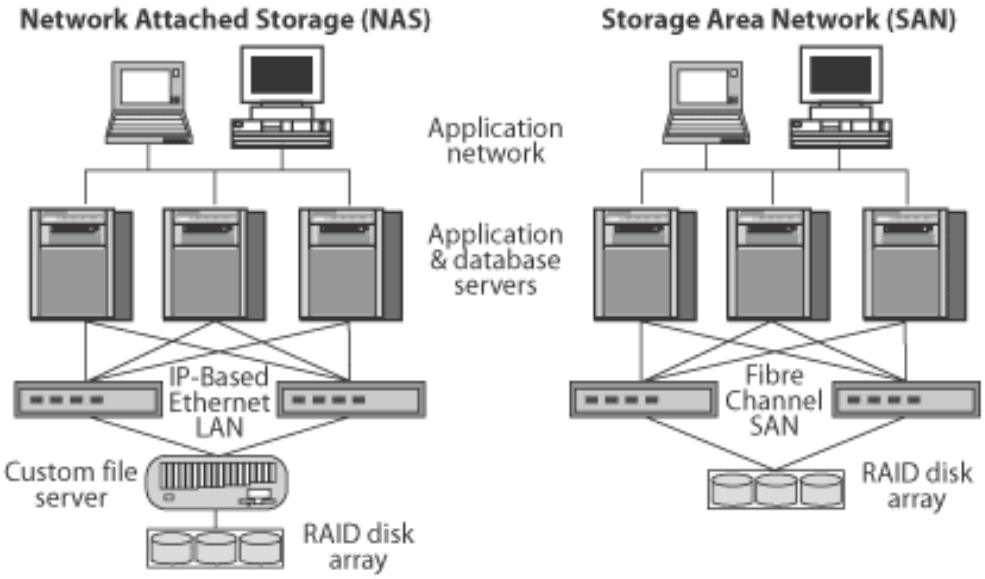

Fig 4: NAS versus SAN (source: wiki)

- **SAN – Storage Area Network i**s different from NAS in the sense that a dedicated network of block data storage in created and this storage network is not accessible through LANs by other devices. SAN is implemented over Fiber Channel, not Ethernet. SAN is supposed to make storage appears local to the operating system. SAS provides high speed block-level access to hard drives for email servers, databases, etc. Producers of SAS include IBM, Aztech, EMC, Hitachi, etc.

- **RAIDs – Redundant Array of Independent Disks** allows the creation of logical disk units, where data can be replicated and distributed across drives, depending on attributes such as level of redundancy and performance. It is one way of building fault-tolerant computer systems where storage is controllable and made to be reliable. Host clusters of data center todays use RAID for primary storage. RAIDs, therefore, refer to physical disk arrays.

- **iSCSI – Internet Small Computer System Interface** – is an interface specification using IP-based approach to link data storage facilities. SCSI commands are carried over IP networks, including LANs and WANs. It enables location-independent data storage and retrieval. SAN uses iSCSI and Fiber Channel.

- **NFS – Network File System** is a distributed file system protocol developed originally by Sun Microsystems (1984). NFS uses RPC to allow a client to access (read/write) files over a network in a way that appears that the storage is local. NFS is very common and used many operating systems including BSD UNIX, Linux and Windows. Hence, NAS uses NFS.

Hence, innovations in servers, networks, and storage have allowed cloud computing to be possible – by enabling scalability, and performance and supporting on-demand access and usage with data reliability and security.

Cloud Operation Characteristics:

Many companies have become cloud providers, hosting cloud services and providing infrastructures locally and globally. To successfully operate a cloud, cloud providers must fulfill certain functions, such as:

- Service Level Agreements (Performance, data portability, location, accessibility, etc.)
- Multi-user Admin
- High Availability
- Robustness and Security
- Support for change in design of Data Center – optimization

Classification of Cloud

When cloud computing was born, people were baffled as to what it meant specifically and how to define or categorize it. Over the years, the IT community has gradually converged to the following classification of clouds:

- **Private Cloud –** As the name implies, this means an organization has sole ownership and control of the cloud which is implemented within the organization, i.e., in house. Considerable expenditure and support have to be provided in this scenario. Private cloud can be used to create VMs and VLANs to support a large organization with external branches.

- **Public Cloud -** refers to cloud established by cloud providers which offer their cloud computing services to enterprise globally, through the internet. Users subscribe to the cloud provider for services by paying a fee and specify the type of service (service level agreement) they desire. Users do not own the cloud in this scenario.

- **Hybrid Cloud –** refers to the scenario where private and public clouds are used by an organization to fulfill its IT objectives. Some CIOs believe that certain infrastructures have to be in-house and certain platforms and software environment can be derived from public cloud.

- **Enterprise Cloud –** refers normally to a public cloud that offers software and services needed by business organizations to operate. Enterprise cloud clients are normally business users and would desire higher quality performance, redundancy, security, and availability.

- **Converged Cloud –** this has similar meanings as hybrid cloud and it allows migration of workload between private and public clouds.

- **Micro Cloud –** refers to a minimum cloud setup created on a desktop or computer. It runs a single virtual machine on a computer, providing mostly PaaS.

Types of Cloud Services

Cloud services can be viewed and offered via a "layered" approach. As shown in Fig 5, it include: (a) IaaS, (b) NaaS, (c) PaaS, and (d) SaaS. Most people would regard NaaS as part of IaaS.

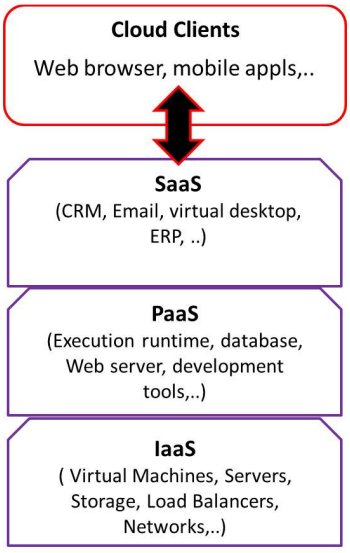

Fig 5: Cloud Services Defined.

IaaS – Infrastructure-as-a-Service: Infrastructure here refers to compute, network, and storage resources. The IaaS provide houses the infrastructure and allows each user to use these resources through the creation and allocation of "instances" and each instance runs on a virtual machine, configured with specified attributes (such as memory, storage size, computation speed, etc.). Hypervisors are used to support large number of virtual machines and hence the ability to scale services

up and down according to demands. IaaS also allows the creation of VLANs (virtual local area networks), and access to file or block based storage.

NaaS – Network-as-a-Service: Network virtualization has enabled the ability to have "Network-in-a-box" implementation and export of network services (addressing, routing, naming, partitioning, VLANs creation, etc.). Network functions are also increasing software-defined, and hence can be viewed as a form of offered services on the cloud.

PaaS – Platform-as-a-Service: PaaS is an upper layer of software suite that provides the environment needed for developers to create software applications. PaaS rides on top on IaaS in order to access to computing, networking, and storage services. PaaS environment is hosted in the cloud (by a PaaS provider) and accessed by the subscriber via a web interface through the internet. Developers need to incur additional cost to setup new IT environment or wait for IT staff to install and implement those environments. This allows developers to achieve quick turn-around time for developing and testing their new applications. PaaS offering can include: (a) operating system, (b) server-side scripting environment, (c) database management system, (d) server software, (e) storage, (f) network access, (g) tools and frameworks for design and development, and (h) hosting.

SaaS – Software-as-a-Service: Currently, IT departments still buy software application packages and install them into their machines. And depending on the operating system platforms of their machines, different versions of the application packages have to be installed. Installation of software applications takes time and sometimes problems occurred during installation, which slows down work flow and hinder productivity. For example, these software applications can be for office work (such as word processing software, drawing or graphic software, project planning software, payroll software, CRM – customer relationship management, ERP – enterprise resource planning, etc.) that once installed, can be used by multiple users in an organization. Normally, software vendors limit the number of users to the software – something controlled by so-called users' license. A large company, therefore, has to pay a lot more to install the same piece of software to be used by the large number of workers. Although this group-users' license approach is better than paying for one-user-per-software approach, it is still considered costly.

The Software-as-a-Service approach is one where the software vendor runs the software in-house and is fully responsible for its installation, upgrade, and maintenance. The application software can be "sold" (a better term is lease) to users based on pay-as-you-go basis. This software-on-demand approach is more scalable as it offloads IT departments from most of the tasks of software installation, upgrade, debugging, and maintenance. In addition, multiple users can execute the application on the cloud (via internet connectivity) and preserve their data either locally or in the cloud or both. SaaS providers include companies like SalesForce.com, SAP, Oracle, Microsoft, Netsuite, etc.

WHO ARE THE CLOUD VENDORS?

With the definition of possible cloud services, cloud vendors have started to take shape. Infrastructure makers (such as DELL, HP, etc.) starts to offer IaaS services, while network equipment makers enter the NaaS market. Operating system companies (such as Microsoft, and Amazon) have their places in PaaS. Finally, software application development companies (such as Oracle, SAP, Microsoft) have their feets on SaaS. Table 1 list some of the companies offering these cloud services.

IaaS	NaaS	PaaS	SaaS
Amazon	Tata Comms	Google	SalesForce.com
IBM	Aerohive	Salesforce.com	SAP
Cisco	Aryaka	Microsoft	Oracle
HP	Pertino	IBM	Microsoft
Dell		CloudBees	Netsuite
Google		Amazon	Google
Microsoft		MegaPath	RackSpace
EMC		NTT Comms	RightScale
Juniper		Openshift	Verizon
TegoData		Pivotal	Appcore
Critix		F5 Networks	
Akamai		AT&T	
Savvis		Metacloud	
Hitachi		RedHat	

Table 1: A List of Cloud Vendors & Service Providers.

CLOUD OPEN STANDARDS – OpenStack & CloudStack

OpenStack – OpenStack is the Cloud OS that provides the platform to manage, compute, store, and network resources in the cloud. OpenStack is open source, where numerous persons can modify, improve on the software. OpenStack uses REST APIs. Over 180 companies participated in OpenStack. As the name implies, it refers to standardized cloud software and specifically cloud operating system that under the Apache-licensed agreement (widely and freely available) allows the creation of massively scalable cloud environments.

OpenStack is sometimes referred to an open forum or open source community and project that allow contributions and enhancements to its code, which essentially consist of 3 parts:

- OpenStack Compute or known as **Nova**,
- OpenStack Object Storage or known as **Swift**, and
- OpenStack Image Service or known as **Glance**.

OpenStack focuses a lot on supporting, creating, and maintaining large networks of virtual machines (VMs). And it also allows one to create secure and reliable storage (through Swift) and finally cataloging and managing massive libraries of server images (through Glance). OpenStack is written in Python and supports both Linux and Windows. OpenStack's other components include:

- Dashboard – **Horizon**: provides a GUI for administrators to control and manage cloud resources.
- Networking – **Quantum**: allow for managing networks and IP addresses
- Block Storage – **Cinder**: provides block-level storage devices to servers
- Identity Service – **Keystone**: provides access control and authentication services

OpenStack component architecture includes an API server, Cloud Controller, Network Controller, Compute Controller, Volume Controller, Authentication Manager, etc. OpenStack has a larger group of code contributors while CloudStack is more production grade, ease to install, and provide better user interface. eBay's X commerce infrastructure uses OpenStack.

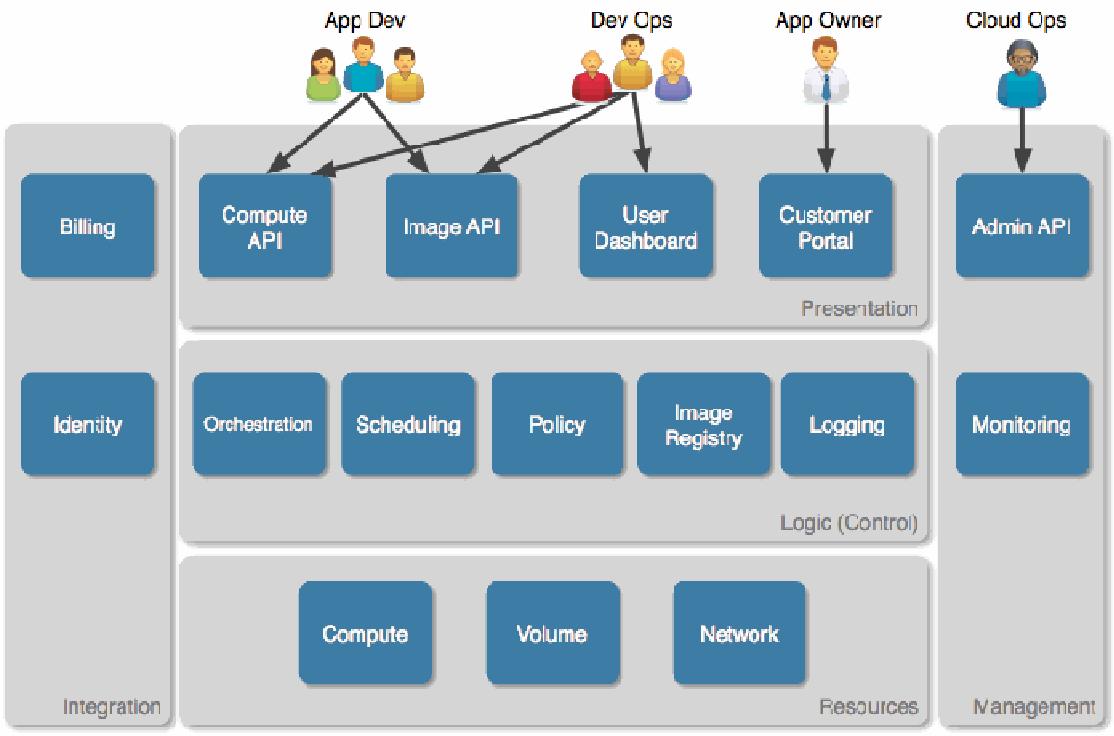

Fig 6: OpenStack Conceptual Architecture (Source: OpenStack)

CloudStack – is an open source cloud computing OS to support the creation, management, and deployment of IaaS cloud services. CloudStack is a widely available software released under Apache License and is written in Java. It supports cross platform, including Linux and Windows. Its most recent version is 4.0.2. (April 2013).

Cloudstack was originally developed by Cloud.com – a company that was acquired by Critix Inc. in 2011. CloudStack uses existing hypervisors (KVM, vSphere, XenServer) for virtualization, and provides its own API, which interoperates with AWS's (Amazon Web Services) API. CloudStack's features include:

- compute orchestration,
- user and account management,
- Network-as-a-Service,
- open native API,
- resource accounting,
- web based user interface,
- multi-tenancy
- account separation.

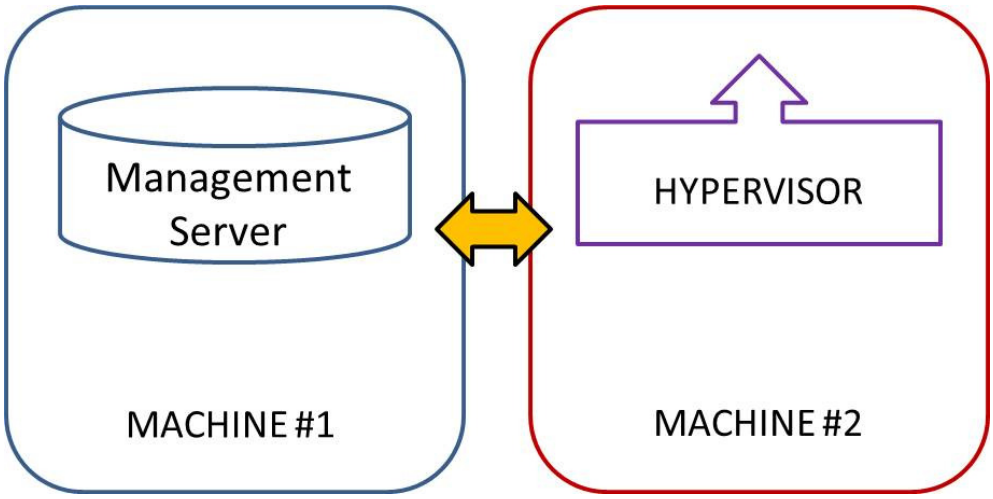

Fig 7: Simplified View of a Basic Cloud Deployment

Cloudstack uses a management server (MS) to manage resources such as hosts, storage devices, and IP addresses. Usually, the MS can be residing on one machine while another machine acts as the cloud infrastructure, running the hypervisor and supporting the creation of VM instances. Cloudstack organizes infrastructure in terms of zones, pods, clusters, hosts, primary, and secondary storage. All hosts within a cluster must be homogeneous. Hosts contain high speed virtualization-enabled CPUs, memories, NICs, etc. Cloudstack management server supports MySQL. CloudStack uses SSH keys for authentication, in addition to the normal username and password. CloudStack can support a variety of hypervisors, including VMware vSphere, Citrix XenServer and KVM Hypervisor.

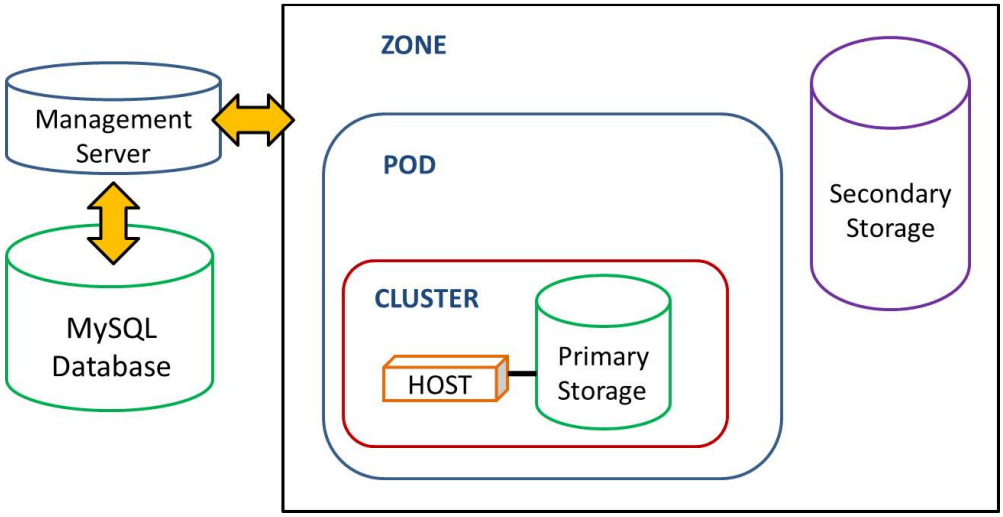

Fig 8: Conceptual View of A Basic Deployment

CLOUD SMALL SCALE DEPLOYMENT ARCHITECTURE

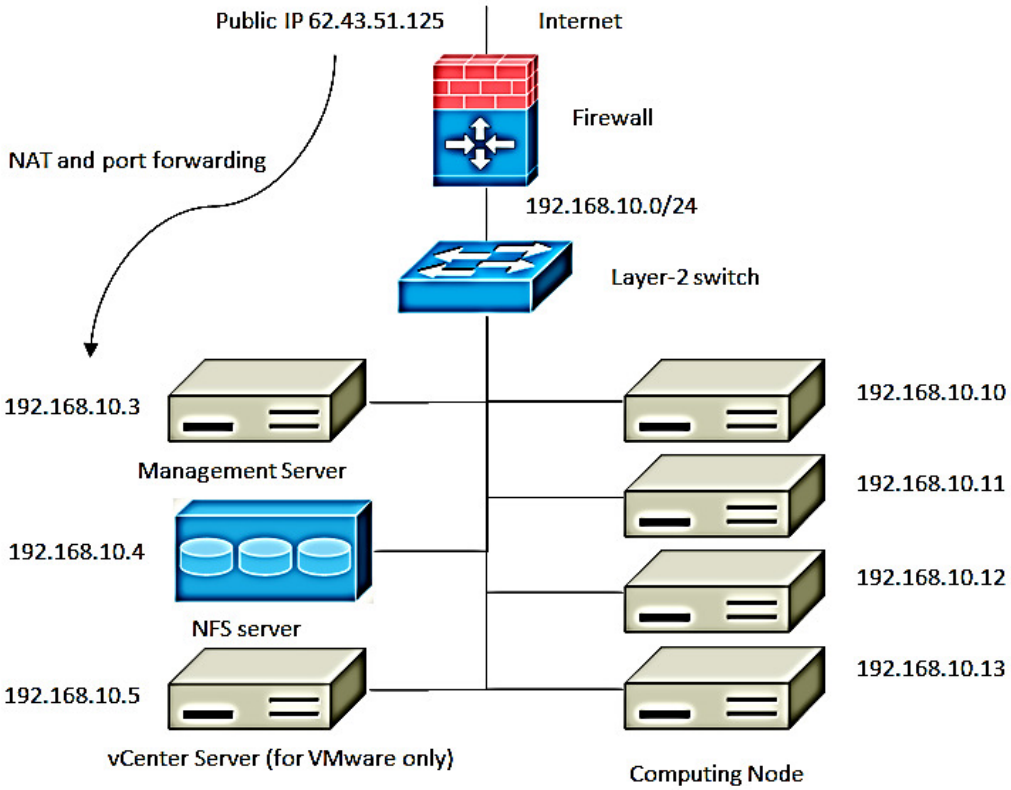

Fig 9: Small Scale Cloud Architecture (source: CloudStack)

In a small scale cloud deployment, a single Layer 2 switch connects all hosts, the management server, and the primary storage, all within the same subnet. This is a single cluster with no zones. A single NIC is present in each host and a firewall provides protection from the external world. This is a tightly coupled and minimum setup.

CLOUD LARGE SCALE DEPLOYMENT ARCHITECTURE

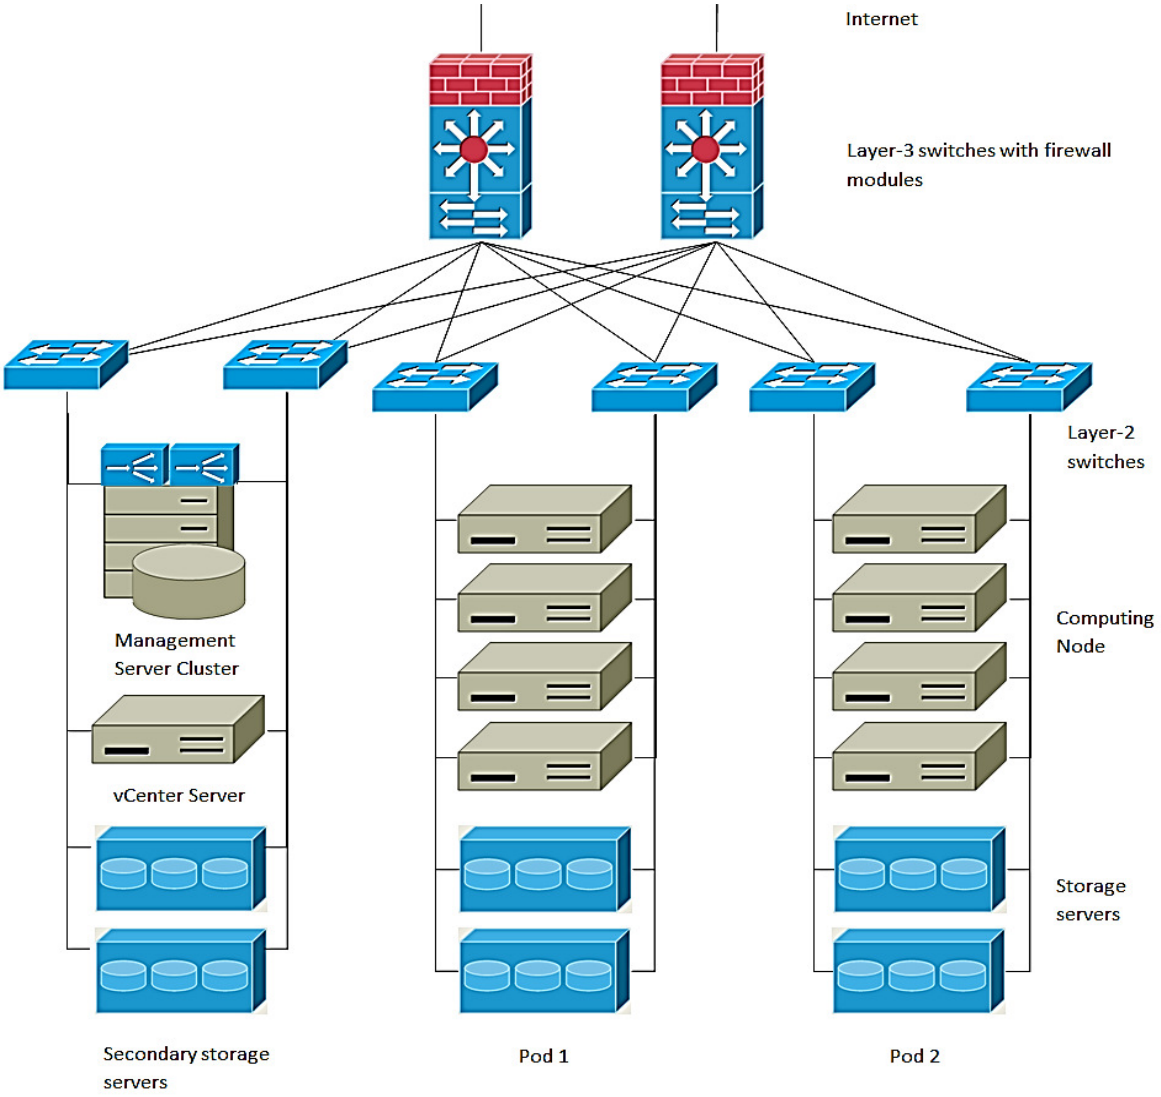

Fig 10: Large-scale Redundant Deployment (Source: CloudStack)

In the large scale deployment architecture, multiple pods are connected to routers which run VRRP (Virtual Router Redundancy Protocol) protocol to ensure redundancy is built into the cloud infrastructure. Firewalls are configured in NAT mode. Multiple Layer 2 switches are connected to each pod to provide extra redundancy in the interconnection of hosts and storage using MPIO – multipath I/O. Load balancers are connected to the management server to redirect load to appropriate pod to

handle the requests. To enable disaster recovery, additional management server and MySQL replication can be used.

CLOUD MULTI-SITE DEPLOYMENT ARCHITECTURE

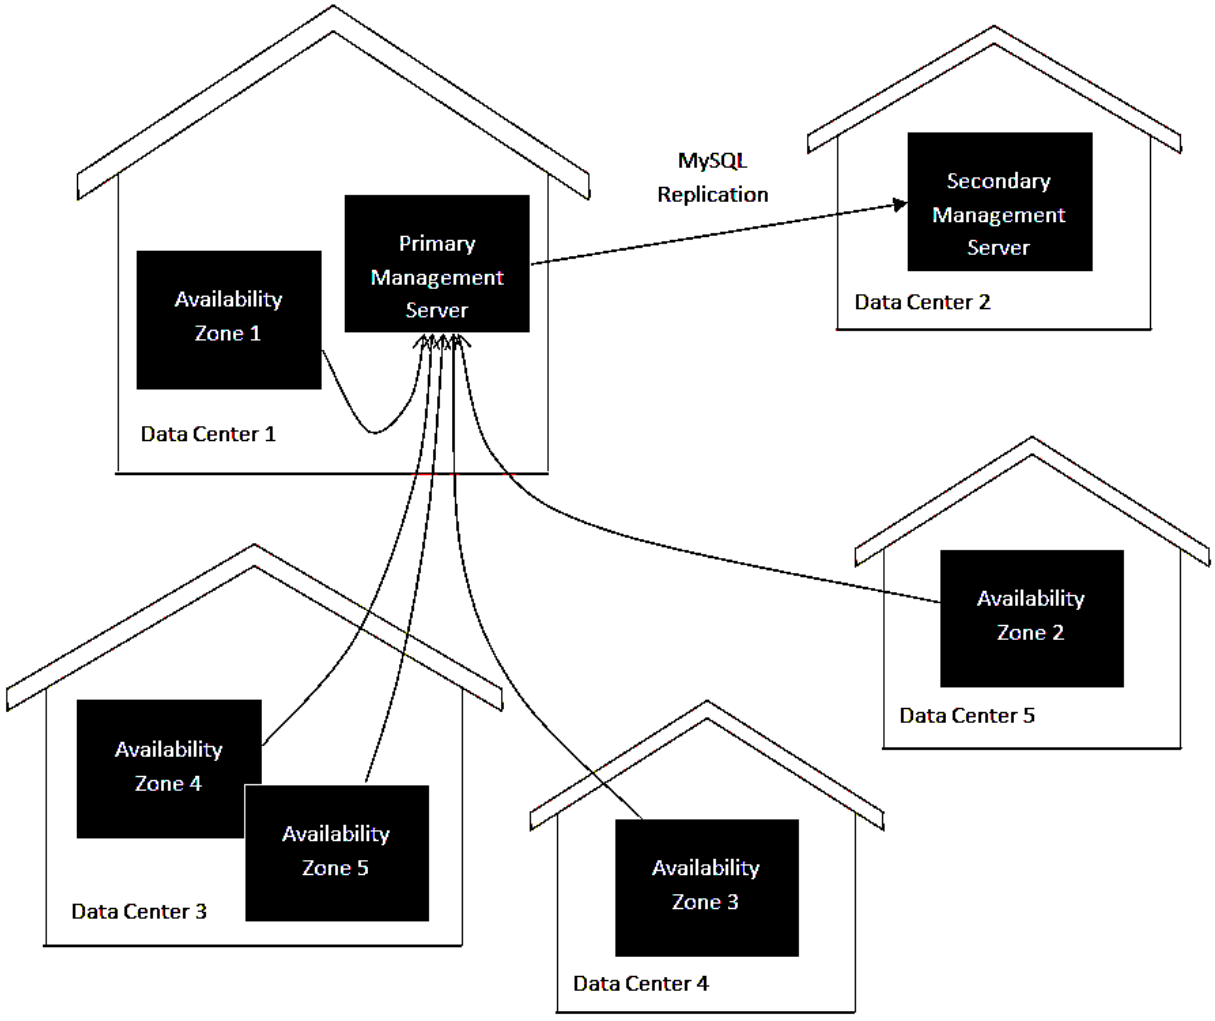

Fig 11: Example of a Multi-Site Deployment (source: CloudStack)

The multi-site scenario can be established through zones. Each zone contains clusters of pods and secondary storage. Multi-site zones can be managed through a primary management server. This allows the establishment of multiple data centers, with the primary data center housing the primary management server. A secondary management server on a different site is used for replication and backup in case the primary management server fails.

CLOUD FOUNDRY

While OpenStack and CloudStack are open source projects for IaaS, Cloud Foundry is the open source for PaaS. Developed by VMware and released via Apache License, Cloud Foundry is written in Ruby. Other companies offering PaaS can also use non-Cloud-Foundry software, such as Amazon Web Services, and Windows Azure. But customers using Windows Azure platform cannot switch to another cloud provider (such as AWS) without having to rewrite their code.

CLOUD APIs – Application Programmable Interfaces

CLOUD OS services and configurations are accessed and modified through a set of APIs (Application Programmable Interfaces). In Amazon AWS, this would be in the form of EC2 commands. In SOAP API, it could be a series of function calls and in CloudStack, it would a series of API calls. The API will cover the following services:

- IP address mapping
- Zone mapping
- Image Attributes mapping
- Instances mapping
- Key pairs mapping
- Security groups mapping
- Snapshots mapping
- Volume mapping

CLOUD SECURITY

Most people would be skeptical about having their own personal data or company sensitive data residing not in-house but with an external party. There is a great issue about security and trust. Nevertheless, one needs not always need data at hand or data in the house to feel secured. Let me give some examples.

Take the case where your money (perhaps millions or billions) are not with you in your pocket or your house, but are deposited in banks (external parties) and investments (external parties too). To date, many people have done so, putting their monies with these external parties (and let these parties safeguard their monies) and they still feel secured. They can remove their monies whenever they wish to. Comparing this to data, data are stored, backup and replicated by cloud providers, and they too have to offer their pledge on safeguarding your data and destroying them whenever the data are no longer desired by the clients. Any breach of this data protection contract can constitute to a criminal offence that is fully enforceable and prosecutable by law.

Another example is the case of email - where millions and billions of users are accessing their emails via YAHOO or GOOGLE or HOTMAIL over the Internet every day. In this scenario, users' email information is stored in disks owned by external parties and users can access their emails over a web interface. Obviously, billions of people have put their trusts into the free and publicly available email system, which is a form of cloud SaaS service.

Hence, with the same mindset, we can make more and more cloud services available through external third parties (cloud services providers) without too much worry about data loss or theft or espionage. Having said that, it does not mean that cloud providers need not provide lip-tight security. In fact, it is expected that data, network, storage, and computer security need to be provided. These include the use of IPSec, TSL, SSL, DNSec, AAA, DES, SSH, data encryption, IDS, firewalls, multi-factor authentication, etc. Bromium, a cloud security startup, proposes a security-focused hypervisor that provides hardware level isolation to prevent undetectable attacks. For government agencies, they must comply with regulatory statutes, such as the Health Insurance Portability and Accountability Act (HIPAA), the Sarbanes–Oxley Act of 2002 (SOX), and the Federal Information Security Management Act (FISMA). Also, the location of a cloud provider's data center can affect the privacy laws that apply to that area of jurisdiction.

Cloud security companies include Bromium, CipherCloud, Cloud Passage, High Cloud Security, and HyTrust. McAfee, TrendMicro, and Symantec are also getting into this new market space.

SUMMARY

The cloud computing era has arrived and will gradually transform how information technology infrastructures evolve in commercial, government, and educational enterprises throughout the world. Cloud computing uses several innovative technologies found in servers, networks, and storage. The establishment of IT infrastructures will change from local to cloud, and services can be paid for based on demand and business needs. Less maintenance and fewer IT staffs are needed in house. For enterprises, IT services can be accessed in a location-independent and scalable manner virtually anywhere in the world through the Internet. Millions of dollars will be saved and productivity will be improved as a result of faster accessibility and availability of information.

Author: Chai Toh is a 15+ year veteran in computing and networking. He has hacked kernels, drafted new architectures, designed new protocols, and wrote client/server software. He architects new systems involving networks, storage, and servers. Working with executives, he defines product strategies, visions, technology roadmaps, and leads product management.