Supplementary Whitepaper
Detailing the Structure of CloudStack
Smart Grid

John Reynolds, CEO, Agile Fractal Grid

10-01-2020

The high-level arrangement for CloudStack resources is shown in the following diagram.

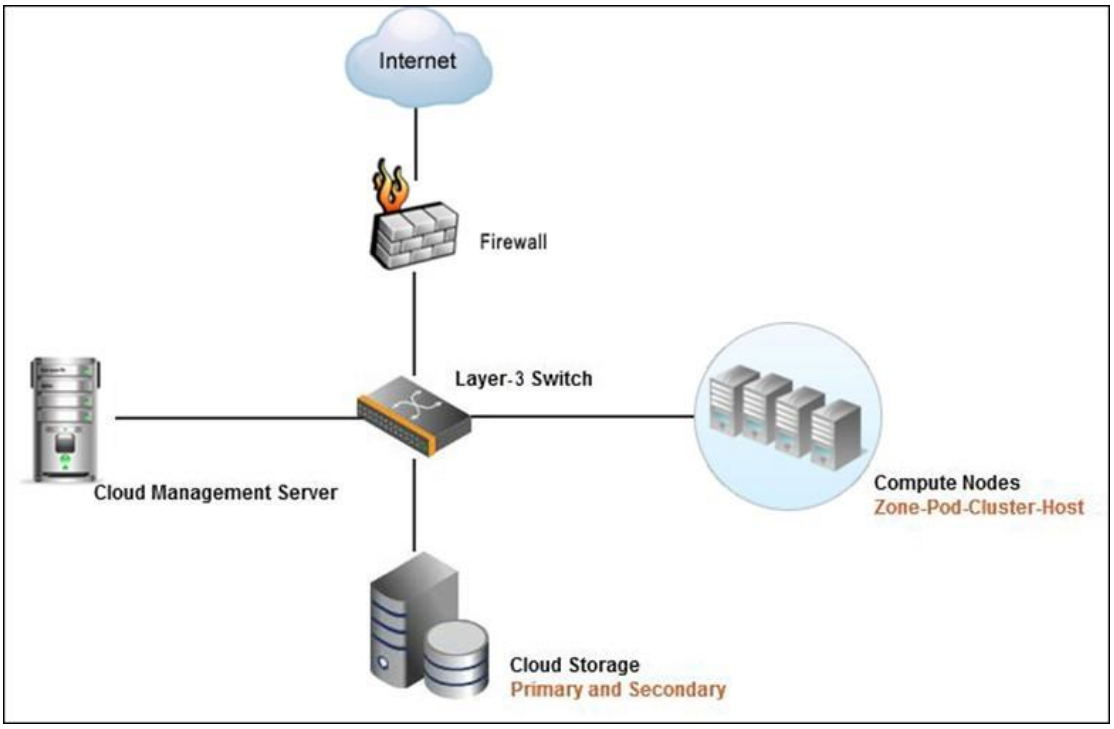

-1-

The general CloudStack approach includes resources for communications, compute nodes, persistent storage, and also management services for controlling the distributed resources of the cloud computing service.

A closer look at the architecture for the CloudStack environment shows that it is organized such that the various computing, storage, and communications operations can be distributed and managed over several locations, too. This is excellent given the structure of the 4-tiered arrangement of the fractal structure where computing resources reside at every node at every layer.

The hierarchy of computing compartmentalization for a CloudStack arrangement is shown in the following diagram.

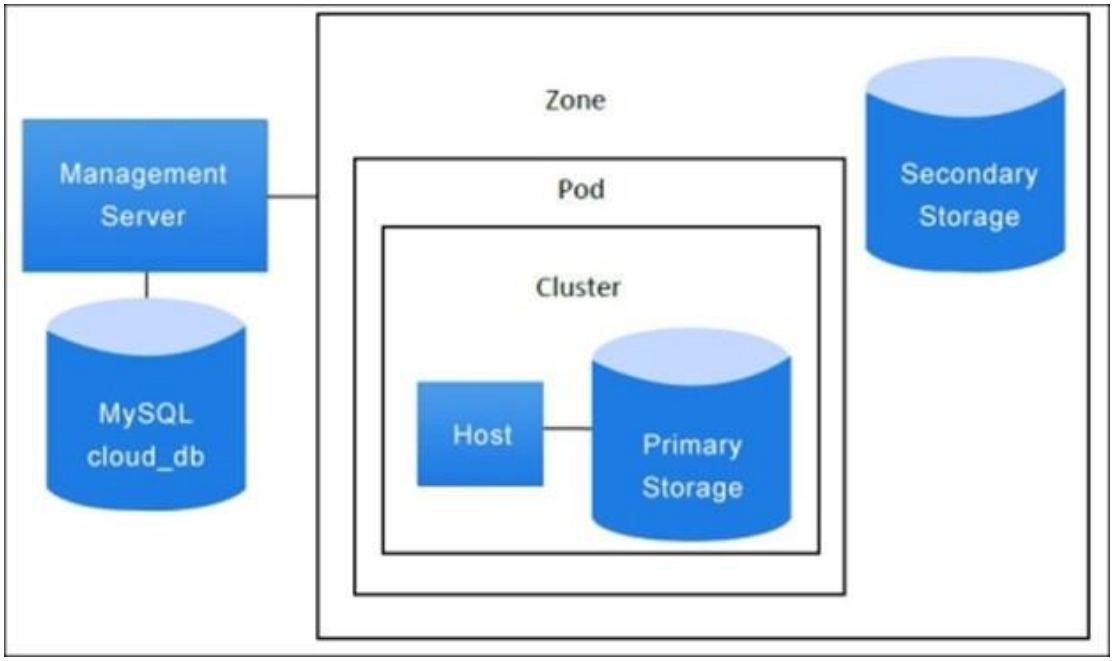

-2-

One could easily see simply declaring that each Tier-2 Cluster Node would equate to what CloudStack calls a “Pod.” Each Tier-3 Regional Energy Operations Center could take on the role of what CloudStack calls a “Zone.” And the overall administration of the distributed computing resources could operate at a Tier-3 Energy Operations Center in what CloudStack refers to as a “Zone.” In practice, should the use of the computing resources of the power grid grow to become used on a massive basis, one could see having a multitude of Zones as AWS and

Azure have done across a national domain to keep from serializing the management of the entire nation down to a single point of failure. Such is the advantage of the fractal pattern as there really is no limit to the expansion of the scope of operations as long as one continues to use the massively parallel processing structure using the fractal patterns.

The CloudStack deployment model covers the basic components of CloudStack, using which CloudStack provides all the functionalities; it also covers the logical segregation of resources to help better manage them.

Zones

- There are multiple zones, each with their own Management Server and Secondary Storage facility – all operating in synchronized high availability mode for disaster recovery. These operate at Tier-3 Energy Operations Centers.
- A Pod is a district rack at a Tier-2 enclosure.
- A Cluster is a subset of processors at a Tier-2 Cluster Node for a district.
- (Tier-1 Fractals are also end nodes using Stratus compute and storage components.)

Zones are the highest level of hierarchy in which the distributed datacenter is logically partitioned to provide physical isolation and redundancy. A CloudStack zone may be a complete datacenter that has a geographic location with its own power supply and network uplinks. The zone can also be distributed as we are using it across various geographic locations. Typically, a unitary datacenter contains a single zone but it can also contain multiple zones. The logical division into zones enables the instances and the data stores at a specified location to be compliant with an organization's data storage policies, and other compliance policies, and so on. (This is important as the popularity of the Agile Fractal Grid expands to international locations and each country will need to operate its own set of zones.)

A zone is divided into various other logical units. A zone comprises of following:

- At least one or more pods. A pod is the second level of hierarchical unit discussed later which consists of one or more clusters. We will discuss this in detail next.
- Secondary storage, this storage is shared by all the pods in a zone.

End users can select a zone while requesting a guest VM.

A zone can be associated with country domains; a zone can be public or private. *A public zone is visible to all the users whereas private zones can be reserved for some specific domain which means that the users that belong to that domain or its subdomain have the permissions to create guest VMs in that private zone.* For instance, the power grid will operate in its own zone and be sequestered from other forms of private cloud processing.

The pod(s) in a zone consist of one or more clusters which in turn contains host(s).

The hosts in a zone can access and be accessed by other hosts in that zone if there's no restriction from any firewall and a flat network is defined, but if a host in a zone tries to access hosts in other zones, they have to follow VPN tunnels which are configured with firewall rules. Among all the other traffic, only key management traffic is allowed in between the hosts in different clusters.

The zones must be configured by the cloud administrator with the number of pods in the zone, number of clusters in a pod, and number of hosts in each of the clusters. The clusters also contain primary storage servers so the cloud administrator must also decide upon the number and the capacity of primary storage servers that are to be placed in each of the clusters. As a minimum configuration, each zone must have at least one pod, each pod must have at least one cluster with at least one host, the cluster must have at least one primary storage. The capacity of the secondary storage that is used by all the hosts in all the pods of a zone is also to be configured by the administrator.

Once these components are configured, they are consumed by the management server to provide cloud services to the customers. A zone is hypervisor specific—a zone will only consist of hosts with the same hypervisor only. Hosts with different hypervisors are added to different zones.

Pods

A zone is further logically divided into one or more entities known as pod. A pod is the second level of hierarchy in the deployment of the CloudStack that occurs at a Tier-2 Cluster Node enclosure. A pod can be assumed to be like a physical rack in a datacenter in which all the hosts reside in the same network subnet. The pod defines the management subnet of the system VMs (discussed later), this network is used by the CloudStack management server communication. A pod contains one or more clusters (or Tier-2 servers) and a cluster in turn contains one or more hosts (Tier-2 server and all Tier-1 Edge nodes within its scope). All the hosts that are inside a pod are configured to be in the same subnet. There are one or more primary storage servers in a Tier-2 pod. A pod can be invisible to users—the users do not have to know which pod their machine is being provisioned in to. The logical division into pods is for administrative purposes only. As with the zone, all the hosts in a pod will have hosts with same hypervisor type as defined during the creation of zone and these hosts are in the same subnet.

Clusters

A cluster is the third level of hierarchical division in the deployment of CloudStack inside the Tier-2 pods. The hosts are grouped into a logical group called clusters inside the pod. This cluster can be a Xen Server pool, VMware cluster that is preconfigured in the VCenter, or a set of KVM servers.

Hosts within a cluster are accessible to each other and guests residing on a host in a cluster can also be "live migrated" to another host in the same cluster using a shared storage device. The live migration occurs without any interruption to the users and thus provides a high availability option. In order to support live migration of VMs between hosts in a cluster, those hosts should have similar properties such as using the same hypervisor version and the same hardware. They should also have access to the same primary storage servers that are shared among the hosts in a cluster. A cluster contains at least one primary storage server.

Storage

In this section we will discuss the various storage options available in the CloudStack system.

Primary storage

A cluster contains at least one primary storage server that is shared among all the hosts in that cluster. This storage server is among the critical component of the cluster and is used to host the guest virtual machine instances. The primary storage is unique to each of the clusters inside the pods.

The primary storage can be a SAN such as iSCSI, NAS such as NFS or CIFS, or it can be a DAS such as VMFS or EXT file systems. The primary storage should be supported by the underlying hypervisor. Building primary storage on high performance hardware with multiple high-speed disks increases the performance. The primary storage should also be reachable from all the hosts in the cluster. This storage is used to host the guest virtual machines in the cluster stores. The volume of these virtual machines and the allocation of guest virtual disks to the primary storage managed by CloudStack.

The primary storage server is basically a machine with a large quantity of disk space, the capacity of which is dependent on the users' need. Primary storage can be a shared storage server or local storage and hosts the storage for guest virtual machines.

The primary storage pool can be created using any of the technologies mentioned earlier for any given cluster. If it is created using iSCSI LUN or an NFS share, the CloudStack management server asks each of the hosts' hypervisors in the cluster to mount the NFS share or the LUN. The hypervisor then communicates with the primary storage pool as it is presented as a datastore in case of VMware, storage repository in case of a Xen Server or a mount point in case of KVM.

The hosts are recommended to have a dedicated management interface for communication with the primary storage pool. The mechanism for making an additional interface for the host for primary storage traffic is to create a management interface.

The primary storage pool provides storage for the root volumes and the disk spaces for the guest VMs. When a guest virtual machine is created, its root volume is automatically created from this primary storage. When the VM is deleted, the data volumes associated with it are disabled and this VM can also be recovered afterwards. It thus provides security by ensuring no data is shared or made available to a new guest in a multi-tenant environment by deleting the data on deletion of a virtual machine.

The primary storage can also be a pool of local storage in case of vSphere, Xen Server, OVM, and KVM. CloudStack supports multiple primary storage pools in a cluster. It also supports dynamic addition of storage servers as your requirements increase.

As primary storage is a critical component, its capacity must be monitored by the administrator and additional storage space must be attached to it when needed. This can be done by creating a storage pool that is associated with the clusters. Additional capacity can be added by adding additional iSCSI LUN to the storage when the primary storage is iSCSI or an additional NFS server can be added to the primary storage when the first one reaches its size limit. Thus, CloudStack supports multiple storage pools in a cluster as well as a single SAN or a storage server can also be used to provide primary storage to multiple clusters.

Volumes

Volumes are the basic unit of storage in CloudStack. All the storage space that is provided to the guest instance is provided in the form volumes. These volumes are created from the primary storage servers that are described as above.

The additional storage space and the storage space for the root disk drives of the VMs are provided by volumes. These volumes are dependent upon the type of hypervisor because the disk image format for different hypervisors are different. So, the volume that has been created for a guest of a hypervisor type cannot be attached to a guest VM using another type of hypervisor.

The guest VMs' root disk is provided storage in the form of volumes from the primary storage that contains all the files for booting the OS or additional storage for storing data. There can be multiple additional volumes mounted to a guest VM. The users are provided with multiple disk offerings (discussed later) that are pre-created by the administrator and which users can select to create different

types of volumes. In addition, a volume can be used to create templates. A volume can be detached or attached to any instance but they must be of same hypervisor type.

The volumes are provided to the virtual machines from the primary storage. CloudStack doesn't provide the functionality of backing up the primary storage but it does provide the functionality for backing up of individual volumes in primary storage using snapshots.

Secondary storage

This storage space is used to store the templates, ISO images and snapshots that can be used to deploy IT infrastructure using CloudStack. Every zone has at least one secondary storage server and this secondary storage(s) is shared by all the Pods in that zone.

CloudStack also provides the functionality to automatically replicate the secondary storage across zones so that one can implement a disaster recovery solution by backing up the application across zones, allowing easy recovery from a zone failure. Thus, CloudStack provides a common storage solution across the cloud. Unlike primary storage, secondary storage only uses **Network File System** (**NFS**) as it is to be accessed by all the hosts in the clusters across the zones irrespective of the hypervisors on the hosts.

The secondary storage is used to store templates that are basically OS images that are used to boot VMs with some more additional configuration information and installed applications.

Secondary storage also stores ISO images that are disk images used for booting operating system and disk volume snapshot, used for backing up the data of VMs for data recovery or for creating new templates. These items are available to all the hosts in one zone.

The administrators can change the secondary storage server afterwards or it can be replaced with a new one after implementation; to achieve this one just needs to copy all the files from the old one to the new one.

CloudStack Management Server

The CloudStack management server helps us to manage the IT infrastructure as defined in the previous sections. The CloudStack management server provides a single point of configuration for the cloud using a Stratus ztc fault-tolerant Edge Node.

Basic functionalities of the management server include:

- The web user interface for both the administrator and the users
- The APIs for the CloudStack
- The assignment of guest VMs to specific hosts
- The public and private IP addresses to specific accounts
- Storage to guests as virtual disks
- The functionalities such as snapshots, templates and ISO images, and their replication across multiple datacenters
- It acts as a single point of configuration for the regional cloud.

The management server provides an easy to use web user interface for administrators and the users who can leverage it to manage and request IT infrastructure on demand. It also provides the users and administrators with CloudStack APIs. The management server can be configured to be fault tolerant to prevent single point of failure. This is achieved by deploying it in a multi-node Stratus configuration and placing it behind a load balancer so that multiple requests must be served by multiple nodes acting as the management server. High availability can be ensured for the CloudDB by setting up a MySQL cluster. Setting up the CloudStack in this fashion, with a fault tolerant management server and highly available CloudDB provides no single point of failure.

The CloudStack Management server is comprised of various parts which handle the functionalities of the cloud.

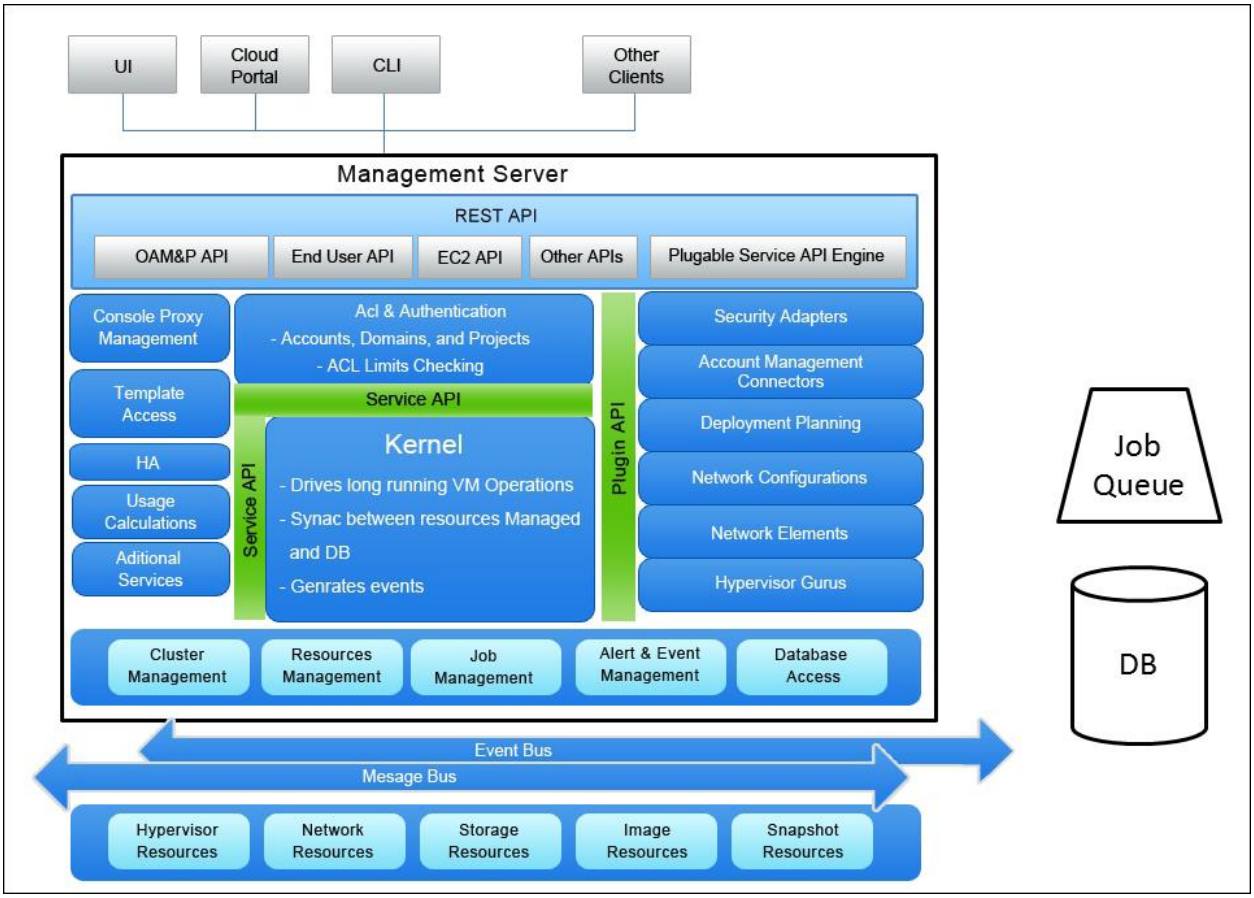

-3-

The basic units of CloudStack Management Server are:

- **Interface**: The User Interface and Application Program interface. The management server provides different types of interfaces for both administrators and users. The user interface is the Management Server console which the admins and end users use for configuring, requesting and provisioning IT infrastructure. The API is used for programmatic access to the server's functionalities.
- **Business Logic**: This part takes care of the business logic and sits below the interfaces. When the user or admins requests any kind of operation, it is processed through the business logic and then passed down to the Orchestration engine to perform the operation. For example, if a user requests a VM using the API or UI, the request is passed through the

business logic section to process the necessary information such as the host the VM is to be deployed to, the workflow process, the required user authentication to perform that operation, the availability and/or applicability of required network configuration, and so on. The request is then passed down to the orchestration engine that performs the operation to create that virtual machine.

- **Orchestration Engine**: The orchestration engine is critical to the CloudStack deployment and is used for configuring, provisioning and scheduling any operations. After the request or command is processed by the business logic, the request is passed to the orchestration engine for processing. The orchestration engine is responsible for performing that action, for example provisioning virtual machine and configuring them, allocating storage and so on. The orchestration engine also helps in scheduling operations.
- **Controllers**: The controllers are basic underlying parts of the CloudStack management server that talks to the underlying hypervisor or hardware for compute, network, and storage resources. These controllers, also known as "Gurus" or "Providers" help in provisioning resources that are requested by the admin or the user and are used to build the guest virtual machine as per the request of the user.

On a more granular level, the basic functions of the management server are divided among the server's various modules, which are described as follows.

API layer

The API layer is the top-most layer of the CloudStack management server that the management server listens to. It is basically the call to the functional components of the CloudStack. It passes on this call to the concerned component of the CloudStack. The various API calls can be among OAM&P API, EC2 API, End User API or it can be any other pluggable service API engine. *The translation of third-party APIs such as Amazon Web Services* is done using the CloudBridge (discussed later). These API calls are fired as per the request by the users or the admin for the execution of some task; the CloudStack management server is responsible for executing the given task as per the request.

Access control

The access control component of the management server is responsible for the access control and authentication of the users requesting services. This layer is the second layer, just beneath the API layer, which cross-checks the authorization of the users requesting the action. The user must authenticate, and the access control component maps the users to the domain, project, and other groups to which the user belongs. The request action should always have an authentication token that authorizes the user for the action and also specifies the permissions, which indicate whether he has the rights to perform the action that he is requesting. This is also recorded in the logs of the actions performed.

Kernel

The kernel is made up of several different components performing tasks in silos. It is the central component for distributing, integrating, and handling the tasks and operations going in and out. The kernel distributes the tasks among the various other components of the CloudStack and drives long-running VM operations, performs sync between the resources, and the database, generates events that are caught by different components and performs actions based on it.

- ***Virtual Machine Manager***: The virtual machine manager is responsible for the management of the virtual machines in the CloudStack environment. All the virtual machine requests, requirements, and states are handled by the virtual machine manager. It is responsible for the management of the resources allocated to the virtual machines in the CloudStack environment. It also manages the live migration, and other actions that are to be performed on the virtual machines such as start, stop, delete, assign IP addresses, and so on. In addition, it ensures that resources are allocated to the virtual machines, as per their needs or specifications.
- ***Storage Manager***: This component of the management server is responsible for the management of storage space attached to the CloudStack as resource. It creates, allocates, or deletes the storage volumes or space as per the end users' request. The storage manager is responsible for all the actions that are concerned with the storage from the users or virtual machines. The storage manager evaluates requests from users and

performs the specific action to the storage resources or generates an error if necessary.

- ***Network Manager***: Network manager handles all the networking of the virtual machines in the CloudStack environment. The network manager is responsible for managing the network configurations of the virtual machines and any other resources in the environment. It has functions such as IP address management, load balancing, firewall configuration, and others that are performed as per the user's request. These configuration operations are predefined in the services that the user chooses. The users are unaware of the operations that are being performed in the back end by such managers.
- ***Snapshot Manager***: The snapshot manager is the component responsible for managing the snapshot of the virtual machines or any other resources in the environment. When a user requests an action for creating a snapshot or any other operations based on snapshots, such as creating a virtual machine using a snapshot, this component takes care of the request. Snapshots are used for backups and restoration. They are taken on the primary storage and then moved onto the secondary storage. There are basically two types of snapshot; **incremental snapshot**, where the snapshot of only modified data is taken since the last snapshot and **full snapshot**, where full snapshot of the service is taken every time.
- ***Async Job Manager***: The jobs that take place in the CloudStack can be synchronous or asynchronous. This component manages the asynchronous jobs that are requested and are to be performed. Commands are usually executed asynchronously; the manager schedules the jobs as per the priority.
- ***Template Manager****:* The template manager is responsible for handing templates and their operations. Whenever there is a request for creating template, creating VM from a given template, deleting template, and such other tasks. The template manager is notified of the same and it handles all the operations pertaining to it.

Below all the components lie some more core components that provide the end-to-end interaction possible in CloudStack. Some of these are discussed as follows:

- ***Agent Manager***: Agents are very critical resource to the Cloud architecture. Agents are deployed in all the resources that are managed in the CloudStack environment. They provide communication channel between the management server and the resources. They provide information, as well as assist in performing operations on the resources.
- ***Resource Layer***: The resource layer is the layer which provides fuel to our engine, i.e. the resources. The resources can be of many kinds such as Xen Server resources, KVM resources, vSphere resources, F5 resources, and so on.

The CloudStack management server is the core component that is responsible for managing the actions that make the whole deployment a success. Let's take a close look into the process flow within the CloudStack—when a new request comes in, how is it fulfilled?

CloudStack Operations

The user is presented with a number of options with respect to the interface through which he or she can submit his/her request or demand. There are user interfaces such as CloudPortal, Command Line Interface, API calls, or any other clients. The user submits the request using such a console.

When a request is submitted by the user, that request is authenticated and the access rights are checked to confirm the user's right to perform the specified request. If the user's authentication fails, the access control layer of the CloudStack management server denies further processing of the request. If the user's request is successfully authenticated, the request is passed by the Access layer to the kernel.

Security Check

After all the security checks are passed, the request is passed down to the kernel and the kernel distributes the tasks to its different components for execution one by one. Let's take an example of a user requesting a new virtual machine service with some software packages installed on it. After passing through the security checks, the request is passed to the virtual machine manager.

The Virtual Machine Manager

The virtual machine manager is responsible for the deployment plan for provisioning the virtual machine such as the host to which this machine is to be created on, to which cluster the host belongs, the pod and the zone of the cluster. The virtual machine manager then starts with the creation of virtual machine by allocating the resources from the host to the virtual machine.

The virtual machine manager initiates the creation of virtual Network Interface Cards (NICs) that must be attached to the virtual machine and assigns this task to the network manager. The network manager takes care of the preparation of NICs and the assignment of IP addresses that are to be attached to the virtual machine that is being created.

The virtual machine manager also triggers the storage allocation to the virtual machine and requests the storage manager to create volume specified. This volume is then attached to the new VM.

If the virtual machine is to be created using a template, then the template manager is contacted for the details and other resources such as the OS, packages, and other resources to be associated with the request.

Server Resources

After all the resources are allocated, the request is passed on to the deployment planner and the server resources. The server resource helps in translation of the CloudStack commands to the resources' APIs. The resource APIs perform tasks as per the instructions and create the virtual machine.

Installation

Once the virtual machine is created, the OS and the software packages are installed on it. The software packages also include the agent that is to be deployed inside the virtual machine, which helps the virtual machine and the CloudStack management server to communicate with each other.

Job Result

After the request is provisioned, the job status is reported back to the user. The user gets the details of the virtual machine that has been newly created upon his request. He can now log in to the virtual machine and use it as he wishes. The job

results and the properties of the virtual machine are stored in the databases. The logs of the process are stored for reference and the database records for the resource are updated.

CloudDB

The CloudDB is the primary MySQL database that is used by management server in the CloudStack deployment for storing all the configuration information. The CloudDB can be installed on the same server or on a different server. The CloudDB can also be set up as a MySQL cluster for high availability. The CloudStack management server communicates with the master database and the replicas are in sync with the master continuously. The administrator must configure the MySQL database or CloudDB with a username and password for security.

The administrator can also provide MySQL replication in order to provide manual failover in case of database failure.

CloudDB is a critical component for the working of CloudStack as it contains information about

- the offering
- hosts' profiles
- accounts credentials
- configuration information
- network information
- and so on.

The database is accessed for information on an on-demand basis, which shows the criticality of the management server and database and the need to configure them properly.

CloudStack Networking Architecture

There are two main types of network configurations that can be deployed using CloudStack: the basic and advanced types.

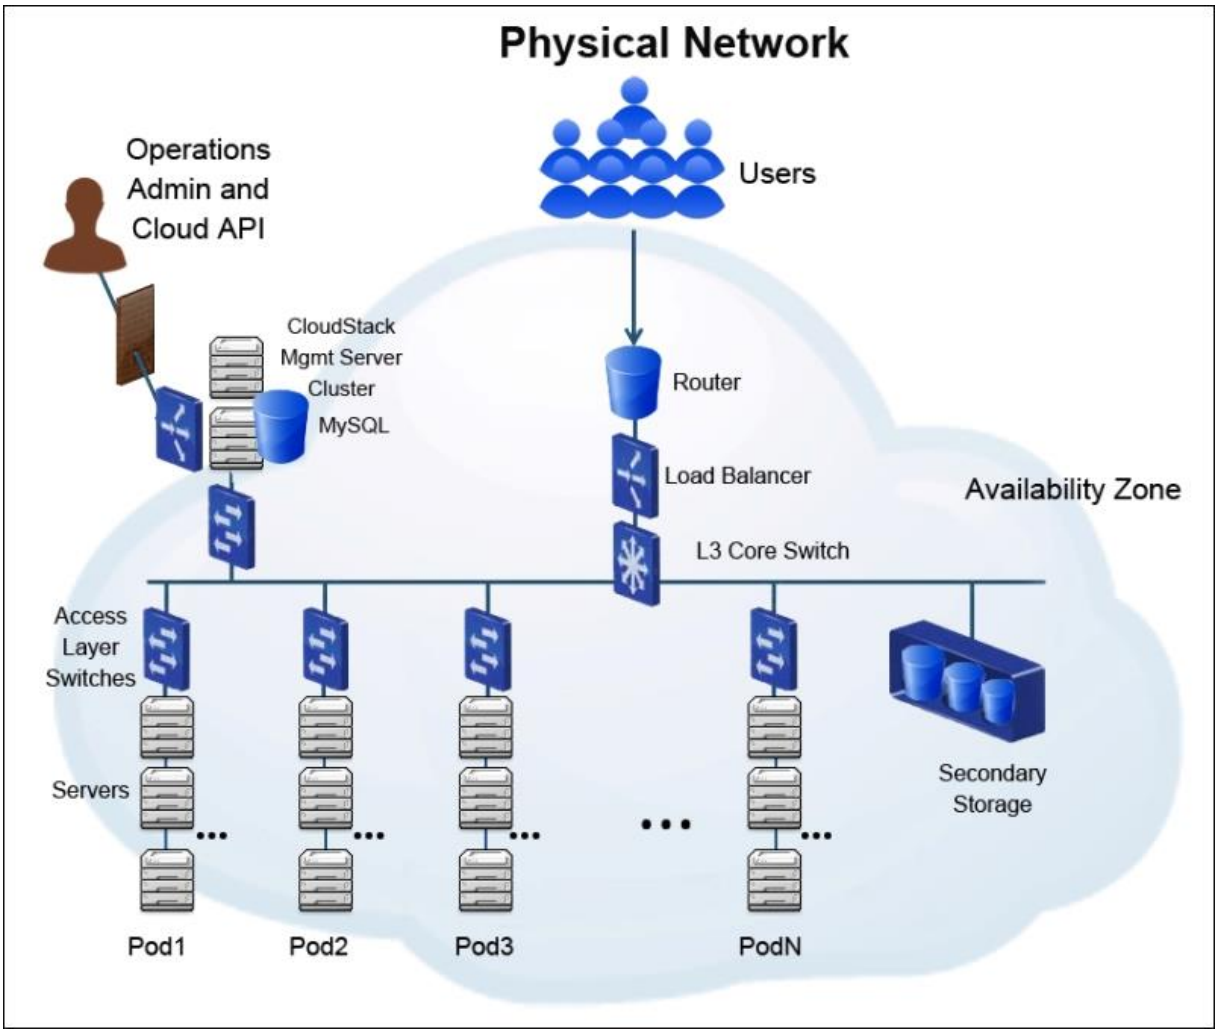

The basic type is similar to AWS level 3 isolation using security groups. The security groups assure the isolation and the egress and ingress of the traffic. The IPAM system manages the IP addresses and tenants don't usually get a contiguous IP address or subnet, the assignment of IP addresses to tenants is

basically random from the IPAM system. This configuration has the capability to scale to configure millions of virtual machines on thousands of hosts.

The advanced type offers full level 3 subnets where security and isolation are provided using VLANs, **Stateless transport tunneling** (**STT**), and **Generic Routing Encapsulation** (**GRE**). Other features such as NAT, VPN, and so on can also be configured.

Network Service Providers

CloudStack network services are made possible with the help of a network service provider, which is basically a network element, hardware, or a virtual appliance. It can be a Cisco or Juniper device(s) that provide(s) firewall services in the same physical network or a F5 load balancer which provides load balancing for the virtual machines registered with it, it can also be a CloudStack virtual router which provides networking configuration for VLANs or overlay network, which helps in the division of a network into multiple tiers.

There can be single or multiple network service providers which are used to provide network services for a single network. There can be multiple instances of the same service provider in a single network. In the case where various network service providers are configured to provide network services, the users have the option to select from the several network offerings that are created by the administrator.

CloudStack Network Offerings

CloudStack provides various network offerings. These network offerings are a group of network services such as firewall, DHCP, DNS, and so on, that are provided as an offering to the users. These network offerings also provide the specifications of the service providers and are tagged to specific underlying network.

The cloud administrator can define new network offerings which can be segregated based on tags. It is up to the administrator to determine the network offering they want to provide throughout their entire cloud offering. The users are allowed to access the network offering based on their tags. The network offerings group together a set of network service such as a firewall, DHCP, and DNS.

The administrator can also choose specific network service providers to be provided as an offering. The network offerings can be any of the three states—**Enable**, **Disable** or **Inactive**. By default, they are in the Disable state when created. Some of the network offerings are used by the system only and the users don't have their visibility. The tags that are used with each network offering cannot be updated but the physical network tags can be updated or deleted.

CloudStack is deployed with three default network offerings for the end users, a virtual network offering, a shared network offering without a security group, and a shared network offering with security group. Furthermore, new network offerings can be created by the administrator to suit the environment needs and include various networking services. These network offerings include different networking services based on the configuration defined.

The shared network offerings are created when the user provisions a VM using that network. The users can also create networks from these network offerings. A set of network offerings can be DHCP, DNS, Source NAT, Static NAT, port forwarding, load balancing, firewall, VPN, or any other optional service provider-based network offering Some of these services are provided using third party hardware equipment such as Juniper or Netscaler.

Types of Network in CloudStack

CloudStack provides various types of network services for end users. CloudStack also supports multiple network services from third parties.

Physical network

A zone in the CloudStack deployment can be associated with one or more physical networks. A physical network can be used to carry one or more types of network traffic. A zone can use the basic network configuration or advanced network configuration, which will decide the type of network traffic that flows through the physical networks.

In a zone with basic network configuration, only one physical network can be present. There are basically three types of network traffic that are allowed. They are:

- ***Guest Network traffic***: This is the traffic flowing over the guest network for communication between the guest VMs when they are running. All the

guest networks which are of type isolated share the same subnet which is set at the zone level. Guest traffic of a VM within one zone is carried in one network, VMs in different zones cannot communicate with each other. In order for the VMs in different zone to communicate, they must do it via a router through a public IP address.

- ***Management traffic***: This traffic is generated by the internal resources of CloudStack. This basically comprises of the traffic between the hosts in the clusters, system VMs (these VMs perform various tasks by CloudStack in the cloud). The administrator must configure the IP ranges of the system VMs. This type of network traffic is usually untagged. The management traffic is should be isolated from the other traffic. The management traffic contains all the UDP traffic for heartbeats. It is highly recommended to isolate the management traffic from the other network traffic.
- ***Storage traffic***: This traffic is the traffic flowing between the primary and secondary storage servers. These can be the VM templates which are placed on the secondary storage and when the user requests to create a VM based on some template, that template data has to flow from secondary storage server to the primary storage server. Another example would be when a user creates a snapshot; the snapshots are stored in the secondary storage, so this snapshot data has to flow to the secondary storage. The storage network traffic is generally configured to be on a separate NIC to ensure better performance. In a zone with advanced network traffic types, there are additional network traffics that flow apart from the traffic flow in zone with basic network traffic. In the basic type of zone, VM traffic is publicly routable by default, whereas in advanced zone type, public label network traffic is exposed.
- ***Public network traffic****:* This kind of traffic flows between VMs and the Internet; this requires the VM to have a public IPs which can be assigned to the VM through the UI. In the case of an advanced network zone, one public IP is assigned to per account to be used as the source NAT. Using hardware devices such as Juniper SRX firewall, a single public IP can be used across all the accounts. Users can also request additional public IPs. This public IP can also be used to implement and configure a NAT instance between the guest and public networks.

All these types of network traffic can be multiplexed in the same underlying physical network using VLANs. It is up to the admin how they configure the network traffic and maps these network types to the underlying physical network and configures the labels on the hypervisor. These can all be done using the admin user interface of the CloudStack. Once the zone is created, the traffic labels can be changed from the user interface, whereas if we need to change the physical networks, some database entries are to be changed as well.

Virtual Network

In order to enable multi-tenancy on a single physical network, the physical network has to be logically divided into several logical constructs, each logical construct is known as virtual network. All the information about the virtual networks and their setting are configured and stored in CloudStack. These settings are activated only when the first VM is started and assigned to this network and the virtual network is also deleted or garbage collected when all the VMs are removed from that network. Thus, CloudStack helps in preserving the network resources and optimizing wastage. CloudStack allows the virtual network to be shared or isolated. The various types of virtual networks are discussed in the following sections.

Isolated Networks

These networks, as the term suggests, are isolated and can be accessed only on virtual machines of a single account except for the domain administrators. The resources such as VLAN are allocated to these types of networks and the garbage collection is done dynamically. The isolated network can be upgraded or downgraded only if it is done for the entire network because it is unique for the entire network.

Shared Networks

As the name suggests, a shared network can be accessed by the VMs of different accounts. But if one wants to attain isolation on this type of network, it can be achieved by using security groups as per CloudStack 4.0. These networks are created by the administrator who can also designate the shared network to a certain domain. The administrator has the responsibility to designate the network resources such as VLAN and the physical network it is mapped to. This network should be pre-created before the guest VM is provisioned on it.

Layer 3 (L3) Network Configuration in CloudStack

The core switches provide L3 network isolation. The core switch connects to the various access switches in various pods. The various features provided in the L3 network configuration in CloudStack are:

- *The IP Address Management systems IPAM:* IPAM provides the IP address management for the network. The virtual router configured can act as the DHCP server to provide IP addresses to the guest VMs in the environment. The tenants are provided with the facility to enable more than one NIC and assign multiple IP addresses to the guest VMs, so an instance can lie in different networks at one time.
- ***Gateway***: The gateway provides routing between multiple subnets. This means if two subnets are to be connected, the traffic flows through the gateway.
- ***Remotely accessible VPN***: CloudStack allows configuration of remote access VPN between multiple remote sites to be connected over a public network. This is facilitated by security features such as IPSec which uses PSK.
- ***Firewall***: Firewall can be configured in CloudStack which is based on source **Classless Inter-Domain Range** (**CIDR**) IP range and egress and ingress of network traffic to the instances.
- ***Network Address Translation (NAT)***: There can be two types of NAT services defined in a CloudStack network—source NAT, and static NAT. The source NAT can be configured per network where the virtual router provides the NAT service to the guest VMs connected to it. Static NAT can be configured by assigning an Elastic IP address to any instance in the environment.
- ***VPN***: The VPN facility can also be configured where a device like the virtual router or core switch can be configured to provide site to site VPN access to connect to remote sites including security features such as IPSec.

The core switches also allow multiple protocol label switching (MPLS), which enhances the speed of communication and prevents unnecessary lookups in the route table. It encapsulates packets of various network protocols and data

packets are assigned labels which are used for forwarding the packets without even examining the packets.

Access Switches or L2 Switches

The access switches used in CloudStack provide L2 network isolation and are present at the POD level. There can be one or more access switches in a POD. The access switches are connected to the L3 core switch present in the zone that connects to all the pods' access switches in that zone. Core switches provide routing and connectivity between the access switches in the various pods in a zone. The access switches facilitate various network traffics such as management traffic, storage traffic, VM traffic, and the public traffic. The access switch has different types of networks connected to its different port groups and thus provides L2 network isolation to these networks.

It also allows configuring an overlay network and VLANs to configure various guest private networks over a single network. The access switches also provide physical isolation based on network labels. This type of configuration allows multiple NICs to be attached to a particular instance so that an instance can belong to more than one network at any time. These multiple NICs can be attached or detached from the VMs over time as per the need.

The administrator can also configure the monitoring of the various traffics that helps in providing updated information of the traffic congestion and the resource utilization. The access switches provide the facility of traffic monitoring of various network traffic and also provides various security configuration such as anti-spoofing and other features.

CloudStack Virtual Router

The virtual network in the deployment of CloudStack consists of various virtual networks which can be configured as per the demands by the administrator. This is achieved by using the CloudStack virtual router. Virtual routers are deployed in the basic type of networking where they are used as a shared service among the multiple tenants and provides features such as DHCP, DNS, and so on. By default, only one virtual router can be deployed per network in an account in the advanced type of networking where there is one virtual router unique to the isolated guest private network. The administrator can raise this limit by increasing the quota.

A virtual router has three NICs, one is connected to the isolated guest network used for advanced VLAN, and is assigned the first IP in the CIDR range and will also act as DHCP, DNS, and gateway for instances in the private guest network, the second NIC is used for a local link network (only for KVM and Xen Server), the management network, and for configuration of the virtual router. The link local interface is created on all the VMs on a host for dedicated communication between the guests and the host, but it is not supported in case of VMware. The third and the last NIC of the virtual router resides on the public network and is assigned a public IP that is used to provide NAT services to the guest VMs connected to it. The source NAT service is already configured on the virtual router in the default isolated mode which forwards the outbound traffic for the entire guest VMs and manages the incoming traffic as per the rules defined by the users.

Networking using CloudStack Virtual Router

The virtual router is used by CloudStack to provide various networking services. The machines connected to guest virtual networks are connected to the outside world using public IP addresses through the virtual router and communicate with the other guests using the virtual router. All the communications from these servers behind the virtual router to the outside world are routed by the virtual router. The CloudStack virtual router is a device which is actually a virtual machine that provides networking services such as DHCP server which assigns IP addresses to the machines behind it, a DNS server which manages the host name to IP address mapping of the machines, NAT, a VPN gateway and many other.

There can be multiple guest virtual networks in the CloudStack environment. These networks are created using some network offering. Most of the networking services that are provided in the network offerings used in the creation of a network are provided by the CloudStack virtual router. The machines can have more than one network interface cards that can be assigned multiple IP addresses to the guest machines. This means that a virtual machine connected to one of the guest virtual network can have two NICs attached to it. One of the NICs has a private IP address and the other is connected to another guest virtual network of different subnet facilitated by another CloudStack virtual router.

The machine with two NICs, one attached to the guest virtual network 1 can be used to host the web server which has incoming public traffic of requests and the other NIC dedicated to the communication to databases server which can be

present on the another VM attached to the second guest virtual router, and it can be accessed only through the web servers and does not have connectivity to the outside world.

The guest virtual network 2 has different subnet or a different VLAN tag and is private to the guest virtual network 1, which means any communication to the guest virtual network 2 has to be routed through the guest virtual network 1. This type of networking architecture is known as a multi-tier network.

In a multi-tier network, there can be multiple CloudStack virtual routers present between different networks of different VLAN tags or different overlay networks, and the communication between these private networks is facilitated by the CloudStack virtual router that acts as the gateway to the communication between the two or more networks, this router can also act as the DHCP server, DNS server, NAT, and for site-to-site VPN access.

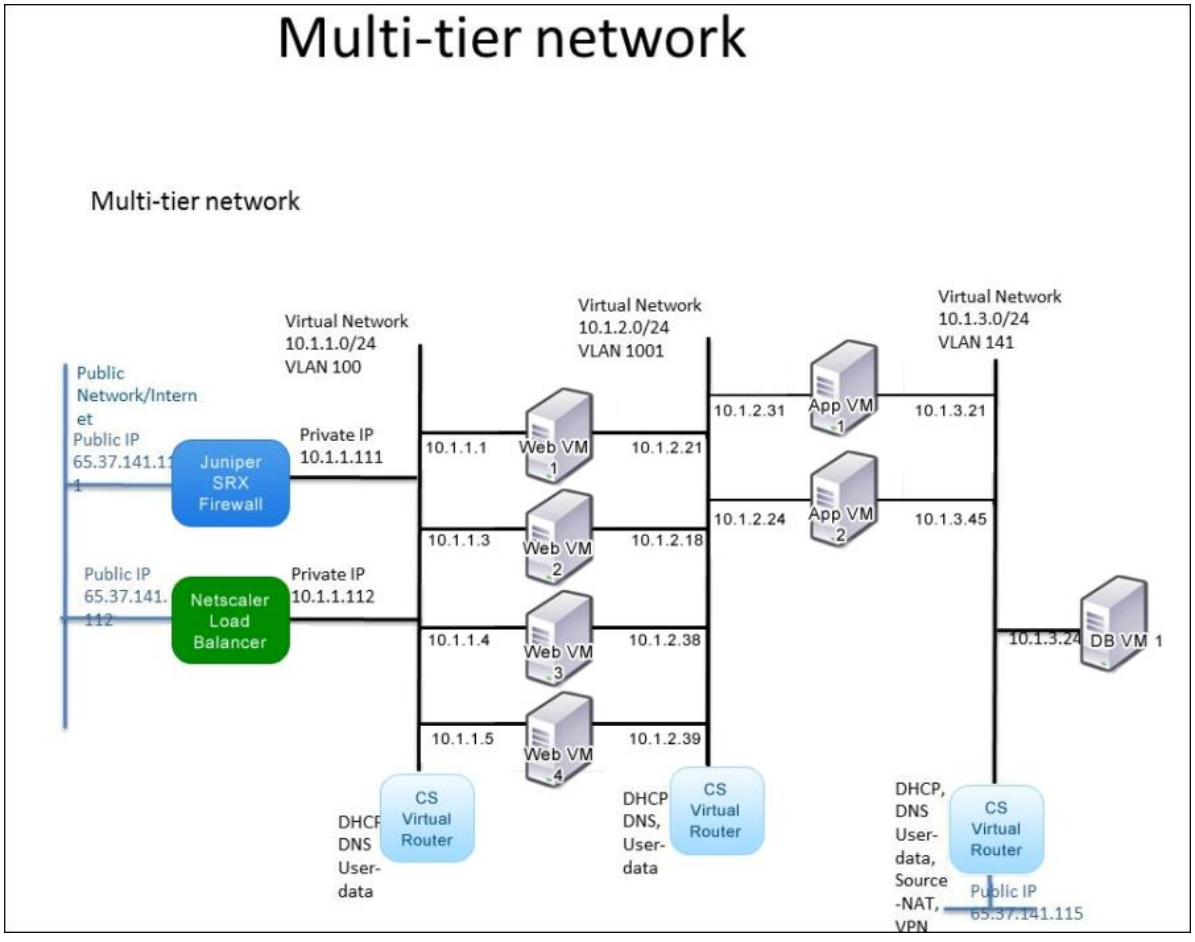

Firewall and F5 Load balancer

The device between the access switch and the Vswitch of the hypervisor hosts can also be a firewall which provides the firewall protection to the virtual machines connected to the guest virtual network. The CloudStack administrator can install the Juniper SRX firewall between the public network and the guest virtual private network.

It can also be a F5 load balancer that provides load balancing services to the application hosted on machines attached to the guest virtual private network. Administrator installs the static route, which can be manual or automated configuration to point to his routing VM. The routing VM also provides VPN connectivity between two sites and this is configured directly on routing VM, not by CloudStack. The load balancer configuration allows architecting a highly available solution where there are multiple instances on different servers hosting same application and registered to the load balancer. In case one of the instances goes down, the other continues to serve users' request. The load balancer also provides the distribution of requests to the different instances registered to it which helps in providing better performance.

Security Groups

Security groups can be attached to any particular instance in the CloudStack environment. The security groups act as the firewall to allow or deny the egress and ingress of network traffic. The rules defined in security groups decide whether communication of some protocol, to some port of the instance from a particular source can be allowed or denied. The security groups can also be used to define rule for the outgoing traffic.

The security groups offer an extra level of security or firewall that can be applied to the instances for restricting the incoming and outgoing traffic. For example, the web servers on one VLAN can have security group to allow traffic from anywhere on the Internet to its particular port serving the users' requests, whereas a backend database server can have security group configured to allow traffic only from the security group of the web servers to its database port and deny any other traffic from anywhere. Thus, the security of the database server is maintained by restricting the access only from the web server. Security groups

provide firewall configuration at the instance level. There can be one security group with a particular setting attached to multiple instances. Security groups allow a scalable network configuration in cloud over VLANs, where the numbers of VLANs that can be created on a Vswitch are restricted.

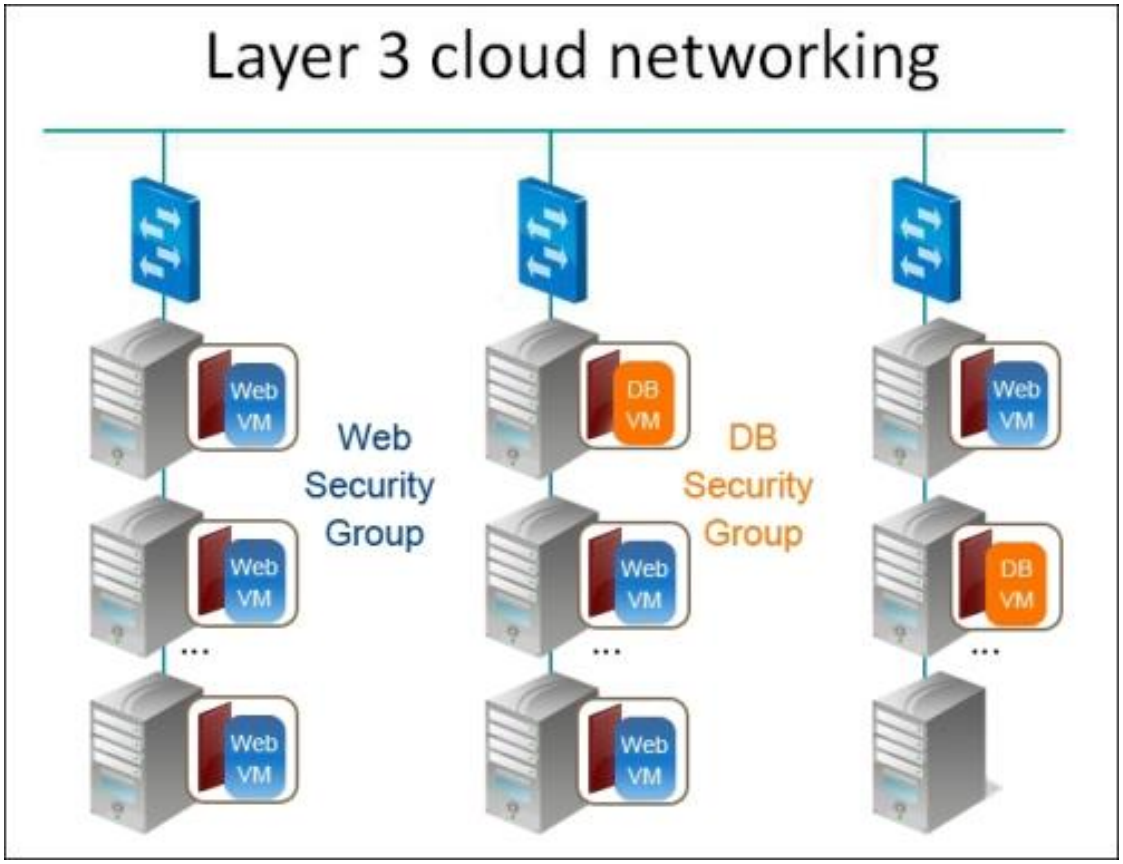

-6-

CloudStack is built using the above defined components and these components are used for different functionalities in cloud. The users can configure these components differently to configure the cloud as they want. CloudStack can be used to deploy a public cloud where people can access it over the Internet on the go, from anywhere; or it can be used to configure a private cloud that is private to an organization and can also be extended to be used as a hybrid cloud where it offers limited public access.