apachecloudstack™
open source cloud computing

Solution Brief:

Delivering Public Cloud Services with Apache CloudStack

Challenges

Cloud has become ubiquitous: not just by SMEs looking to increase agility and avoid prohibitive CAPEX, but also by enterprise and governmental departments. The early march of the hyperscalers in the IaaS space led many smaller service providers to move away from providing their own IaaS offerings, instead looking to work commercial models with those hyperscalers. But the relentless demand, not just for commodity cloud, but also for more specialised services around hybrid, geo-spanned deployments, edge computing and data sovereign environments, has driven many service providers to re-assess this strategy.

Only by delivering services from their own datacenters can Service Providers fully capitalise on this market that is expected to reach $189.52 billion in 2027. The correct choice of the underlying technology stack is essential: often the difference between commercial success and failure. Balancing functionality, longevity, cost of ownership and ease of operation is pitted with many potential challenges.

The key challenges cloud service providers need to overcome are:

- Building secure and reliable cloud service
- Ensuring cost-efficiency to stay competitive
- Achieving good profit margins
- Increasing data center costs and competition
- Complex cloud management requires a larger team to support it
- Control over the provisioning, de-provisioning, and operations of infrastructure

Over 50% of SME's technology budget will go to cloud spend in 2023

71% of organizations expect their cloud spend to increase

Solution

apachecloudstack
open source cloud computing

Traditional cloud deployments require a large hardware footprint. Separate hardware is allocated for the cloud's control plane services, compute resources and virtual machine hypervisor and storage. A highly available cloud service deployment requires an even larger number of servers for redundancy and to be able to deliver cutting-edge cloud solutions. All this complex hardware setup can be managed easily and efficiently with a powerful cloud management platform like Apache CloudStack.

Apache CloudStack is a highly scalable IaaS cloud computing platform that provides a cloud orchestration layer, automating the creation, provisioning and configuration of IaaS components (such as virtual servers). It turns existing virtual infrastructure into a cloud-based Infrastructure as a Service (IaaS) platform. Because CloudStack leverages existing infrastructure, the costs and time for the organization to build a multi-tenant IaaS platform are significantly reduced.

Among the most significant advantages CloudStack is the simplicity and ease of use it brings, even for large-scale and geo-distributed environments. With CloudStack, you can orchestrate hosted public and private clouds, on-premise private clouds and hybrid environments without the need to engage a huge operations team to support them in the long term.

CloudStack can manage tens of thousands of physical servers installed in geographically distributed data centers. The management server scales near-linearly eliminating the need for cluster-level management servers. Maintenance or other outages of the management server can occur without affecting the virtual machines running in the cloud.

Apache CloudStack powers some of largest clouds globally and is used by British Telecom, KDDI, Fraunhofer-Gesellschaft, Pearson Vue, SAP, Ticketmaster, NTT Data, and thousands of more organizations around the world.

CloudStack Architecture

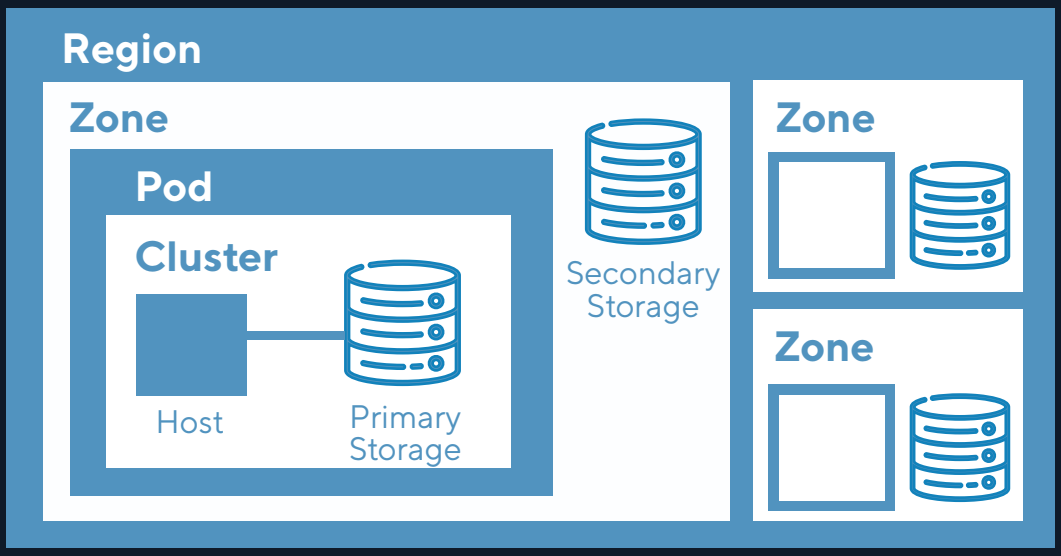

CloudStack Architecture

Zones

A Zone is the largest organisational unit within a CloudStack deployment, many providers refer to Zones as Availability Zones typically, a data center (DC) implementation will contain a single Zone, but there are no hard and fast rules, and a DC can contain multiple Zones. An Availability Zone consists of at least one Pod, and Secondary Storage which is shared by all Pods in the Zone.

Pods

A Pod relates to a discrete rack in a data center. Pods contain one or more Clusters and a Layer 2 switch architecture which is shared by all Clusters in that Pod. End users are not aware of and have no visibility of Pods. Pods are usually the primary unit of scale in a CloudStack deployment.

Clusters

A Cluster is a group of identical compute servers running a common Hypervisor. Each Cluster has dedicated Primary Storage array which is where the virtual machine instance disks are hosted.

Primary Storage

Primary Storage is usually unique to each Cluster (although it could also be used Zone-wide) and is used to host the virtual machine instance disks. CloudStack is designed to work with all standards-compliant (iSCSI, HBA or NFS) and Software-defined Storage (SDS) solutions supported by the underlying Hypervisor.

Secondary Storage

Secondary Storage stores virtual machine Templates, ISO images and Volume Snapshots. The storage is available to all PODs in a Zone. Secondary Storage uses the Network File System (NFS) or an S3 object storage to ensure it can be accessed by any Host in the Zone.

CloudStack Networking

CloudStack is extremely versatile when it comes to networking, and it supports different types of guest network setups for instances to communicate with each other and the public internet. Guest Networks in CloudStack provide the services for Instances (guests) to communicate with each other. Guest Networks in CloudStack coexist next to the storage and management networks, but the last two networks are internal and thus invisible to users of the environment.
There are two main types of guest networks to discern, Isolated Network and Shared Network. Also, there is the option to create a Virtual Private Cloud (VPC), grouping instances in separate tiers with their own IP subnet from a larger CIDR. Networks leverage layer 2 MAC addressing to establish networking links between the available devices, both in Shared and Isolated networks. It is possible to use VLAN-based or supported Software-defined Network (SDN) to isolate the guest traffic.

Storage

CloudStack is designed to work with various commodity and enterprise-grade storage systems. It can also leverage the local disks within the hypervisor hosts if supported by the selected hypervisor. Storage type support for guest virtual disks differs based on hypervisor selection.
CloudStack defines two types of storage: primary and secondary. Primary storage can be accessed by either HBA, iSCSI, NFS or supported SDS storage. Additionally, direct attached storage may be used for primary storage. Secondary storage is accessed using NFS or S3 protocols. In addition, many storage providers such as LINBIT, StorPool, NetApp Solidfire and Dell PowerFlex have additional enhanced integrations with CloudStack.

Hypervisor Support

CloudStack is hypervisor agnostic, supporting VMware, KVM, XenServer and Xen Cloud Platform, XCP-ng and Hyper-V. This gives CloudStack a unique position in the market: it is not proprietary technology and is not tied to any specific infrastructure.

User Interface

CloudStack offers an administrator's web interface used for provisioning and managing the cloud, as well as an end-user's Web interface used for running VMs and managing VM templates. The UI can be customized to reflect the desired service provider or enterprise look and feel and is decoupled from the CloudStack core. It relies on the API for managing CloudStack resources.

Restful API and CLI

Users can manage their cloud with an easy to use web interface, command line tools, and/or a full-featured RESTful API. The CloudStack UI helps the CloudStack administrator provision, view, and manage the cloud infrastructure, domains, user accounts, projects, and configuration settings. It is modern, well-organised and easy to use.

In addition, Apache CloudStack is integrated with leading open source cloud infrastructure automation projects such as Terraform, Ansible, ClusterAPI and Packer enabling developers to easily integrate existing CI/CD workflows into the platform.

CloudStack Large-Scale Redundant Deployment for Service Providers

The following logical diagram shows a typical Service Provider deployment of Apache CloudStack. This deployment shows a single zone deployment with a mixed hypervisor environment (KVM and VMware). The CloudStack Management servers and their underlying database are deployed into a Management Cluster, which provides the orchestration plane for the environment. Secondary storage is usually centralised per zone. Each Pod (usually a single rack) will have it own leaf switch. A Pod can consist of multiple clusters, each having its own Primary Storage and running a specific hypervisor. The Pod is the usual scale point in service provider environments. Further Zones can be added, usually at different data-centres – orchestrated from the initial Management Cluster.

The actual architecture chosen is usually dependant on the specific business objectives of the service provider - Apache CloudStacks flexible architectural constructs allow for a wide range of deploy-ment models.

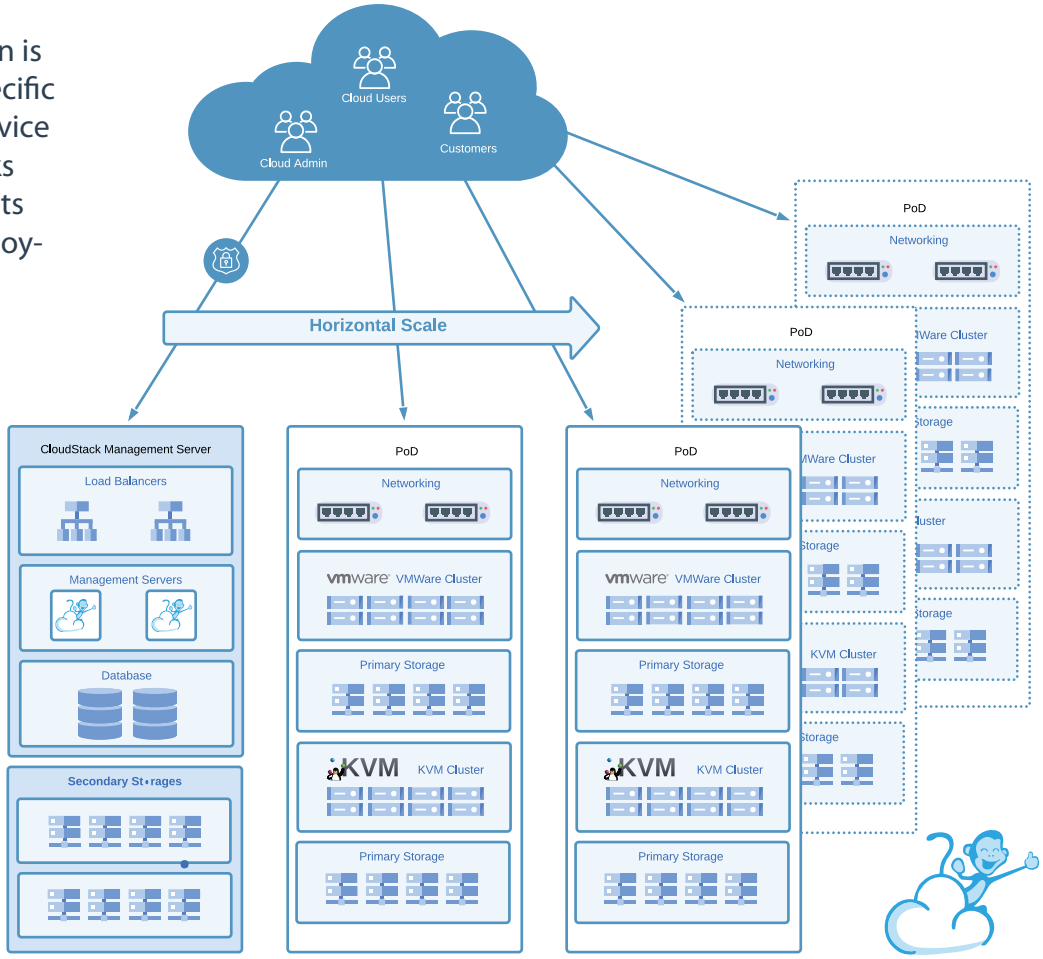

Project Sustainability

Apache CloudStack is a full-featured IaaS cloud computing platform supporting a wide range of integrations. Its global community constantly develops new features and supports new technol-ogies with a clearly defined, evolving roadmap guided by users and the community. There are no different levels of support or versions, and although there are vendor distributions available, no vendor has a dominant influence over the project and most organizations run the freely avail-able, open-source version in production.
Being open-source ensures that CloudStack follows the needs of its user and enables them to build future-proof technology solutions.

Leading companies trust Apache CloudStack

BT KDDI ikoula leaseweb reliable hosting
CloudOps NTT DATA Telia

apachecloudstack™
open source cloud computing

Apache CloudStack supports a wide range of integrations.
Its global community constantly develops new features and supports new technologies with a clearly defined, evolving roadmap guided by users and the community. There are no different levels of support or versions, and although there are vendor distributions available, no vendor has a dominant influence over the project and most organizations run the freely available, open-source version in production.

Being open-source ensures that CloudStack follows the needs of its user and enables them to build future-proof technology solutions.

https://cloudstack.apache.org/

SCAN ME