

Dell PowerFlex Integration with Apache CloudStack

January 2024

H19906

White Paper

Abstract

This white paper describes the integration of Dell PowerFlex storage with the Apache CloudStack Infrastructure-as-a-Service platform. This integration facilitates the deployment of reliable and easy-to-manage cloud infrastructure with KVM hypervisor using the PowerFlex storage.

Dell Technologies Solutions

PowerFlex Engineering

Validated



Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2024 Dell Inc. or its subsidiaries. Published in the USA 01/24 White Paper H19906.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of Contents

Executive summary	4
Introduction	5
Revision	5
Audience	5
Terminology	5
We value your feedback	6
Product Overview	6
PowerFlex family	6
PowerFlex software components	6
PowerFlex deployment architectures	8
PowerFlex consumption options	9
Apache CloudStack	11
Solution overview	12
Overview	12
Logical architecture	12
CloudStack architecture	13
Network architecture	15
Deployment of CloudStack	17
Overview	17
Management Server installation	17
CloudStack Infrastructure: KVM Host setup	20
Accessing the CloudStack UI	23
CloudStack configuration	23
CloudStack and PowerFlex integration	34
Conclusion	37
References	38
Dell Technologies documentation	38
Apache CloudStack documentation	38
Appendix A	39
Configuration details	39

Executive summary

In the era of cloud computing and emerging data, storage is not just a necessity, it is the bedrock upon which cloud services are built. As data grows exponentially, the demand for reliable, scalable, and cost-effective storage has increased.

Apache CloudStack is an open-source Infrastructure-as-a-Service (IaaS) platform that manages and orchestrates pools of storage, network, and computer resources to build a public or private IaaS compute cloud.

Within the CloudStack environment, storage assumes a pivotal role to ensure the seamless functioning of instances and the accessibility of vital resources. This infrastructure hosts instance both root disks and additional data disks. For optimal CloudStack deployment and operations, a reliable and high-performance enterprise storage platform is indispensable. It must offer attributes such as reliability, low latency, and high IOPS to guarantee the optimal performance and operations of instances within CloudStack. PowerFlex, with its software-defined approach, stands as an ideal platform to meet these requirements, especially in the context of large-scale IaaS deployments.

PowerFlex also delivers stable and predictable performance for all workloads. Furthermore, PowerFlex Manager, a unified management toolset for PowerFlex systems, simplifies IT Operations Management (ITOM) and Life Cycle Management (LCM) tasks by providing extensive automation and a rich out-of-the-box toolset that includes PowerFlex REST API, PowerFlex Ansible modules, and Dell Container Storage Modules (CSM) and CSI drivers to enhance DevOps productivity and IT agility.

CloudStack can use these REST APIs to enhance functionality, facilitate streamlined provisioning, effective data management, robust snapshot management, comprehensive data protection, and seamless scalability making the combination of PowerFlex storage and CloudStack a robust choice for modern IaaS environments.

This integration allows you to enable instant snapshots, volume provisioning, VM provisioning, and other settings that can be changed in-flight to address changes in user requirements.

This white paper provides guidance for the deployment of CloudStack and configuring Dell PowerFlex as the underlying software-defined infrastructure. This paper also gives an overview of a few CloudStack operations such as VM operations, templates, volumes, and snapshots using the PowerFlex storage.

Introduction

Revision

Date	Description
January 2024	Initial release

Audience

The audience for this paper includes system engineers, field consultants, IT administrators, technical architects, and anyone else interested in configuring and deploying Apache CloudStack with Dell PowerFlex as the underlying software defined infrastructure.

Readers are expected to have an understanding and working knowledge of Red Hat Enterprise Linux, KVM Hypervisor, Apache CloudStack and Dell PowerFlex software-defined infrastructure.

Terminology

The following table provides definitions for some of the terms that are used in this document.

Table 1. Terminology

Term	Definition
API	Application Programming Interface
CLI	Command Line Interface
DNS	Domain Name System
IaaS	Infrastructure-as-a-Service
KVM	Kernel-based Virtual Machine
MDM	Meta Data Manager for PowerFlex
RHEL	Red Hat Enterprise Linux operating system
SDC	Storage Data Client for PowerFlex
SDS	Storage Data Server for PowerFlex
VM	Virtual Machine

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by [email](#).

Author: Punitha HS

Content editor: Ripa Bhagawati

Note: For links to additional documentation for this solution, see [Dell Technologies Solutions Info Hub for PowerFlex](#).

Product Overview

PowerFlex family PowerFlex software-defined infrastructure enables broad consolidation across the data center, encompassing almost any type of workload and architecture. The software-defined architecture offers automation and programmability of the complete infrastructure and provides scalability, performance, and resiliency to enable effortless adherence to stringent workload SLAs.

The PowerFlex family provides a foundation that combines compute and high-performance storage resources in a managed unified fabric. PowerFlex comes in flexible deployment options (rack, appliance, or custom nodes and in the public cloud) that enable independent (two-layer), HCI (single-layer), or mixed architectures. PowerFlex is ideal for high-performance applications and databases, building an agile private/hybrid cloud, or consolidating resources in heterogeneous environments.

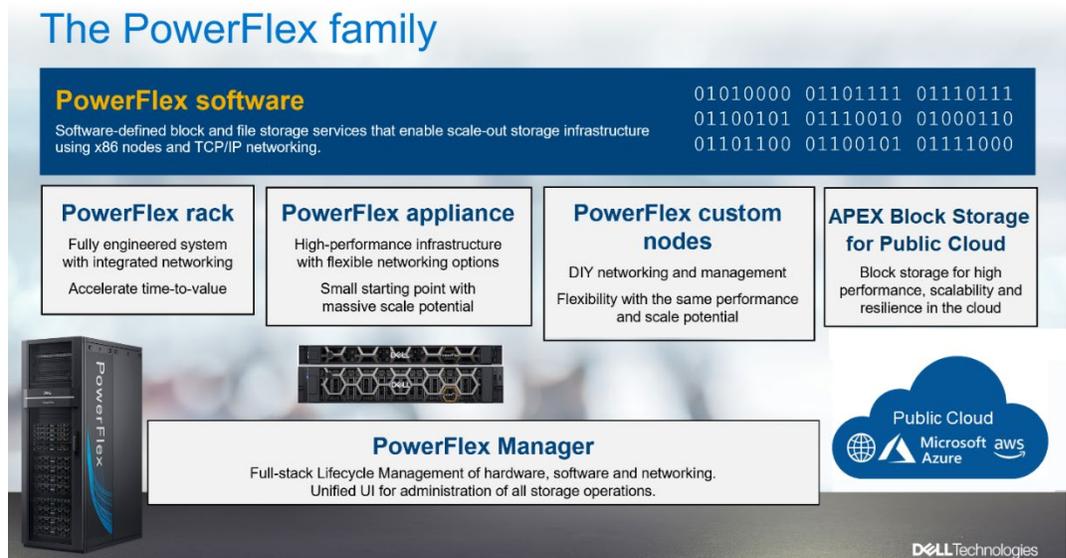


Figure 1. PowerFlex family

PowerFlex software components

Software is the key differentiation in the PowerFlex offering. PowerFlex software components not only provide software-defined storage services, but also help simplify infrastructure management and orchestration. This enables comprehensive IT Operational Management (ITOM) and Life Cycle Management (LCM) capabilities that span compute as well as storage infrastructure, from BIOS and Firmware to nodes, software, and networking.

PowerFlex

PowerFlex is the software foundation of PowerFlex software-defined infrastructure. It is a scale-out block and file storage service that is designed to deliver flexibility, elasticity, and simplicity with predictable high performance and resiliency at scale.

PowerFlex Manager

PowerFlex Manager is the software component in the PowerFlex family that enables ITOM automation and LCM capabilities for PowerFlex systems. Starting with PowerFlex 4.0, the unified PowerFlex Manager brings together three separate components used in previous releases – PowerFlex Manager, the core PowerFlex UI, and the PowerFlex gateway. The new PowerFlex UI runs in Kubernetes and embraces a modern development framework.

PowerFlex file services

PowerFlex File Controllers, also known as File Nodes, are physical nodes that enable PowerFlex software defined File Services. They host the NAS Servers, which in turn host the tenant namespaces and file systems, mapping PowerFlex volumes to the file systems presented by the NAS Servers. All major protocols are supported, such as NFS, SMB/CIFS, FTP, and NDMP. PowerFlex file service is supported from PowerFlex 4.0.

PowerFlex with NVMe over TCP (NVMe/TCP)

NVMe is a family of open specifications that defines an industry standard for both a logical interface and a communication protocol for accessing block storage across an internal PCIe fabric or through external fabrics such as TCP/IP and FC. NVMe is designed to take advantage of the internal parallelism and low latency inherent in flash-based media SSDs.

NVMe/TCP is an extension of the NVMe base specification that defines the binding of the NVMe protocol to message-based fabrics using TCP. By binding the NVMe protocol to TCP, NVMe/TCP enables the efficient end-to-end transfer of commands and data between NVMe hosts and NVMe controller devices by any standard Ethernet-based TCP/IP networks.

From PowerFlex release 4.x onwards, support for NVMe/TCP has been added for on-premises deployment with the introduction of a new Storage Data Target (SDT) component installed in the storage layer. The NVMe initiator in the operating system or hypervisor communicates with the SDT, which then communicates with the Storage Data Server (SDS). The NVMe initiator is part of the kernel of the application host operating system or hypervisor.

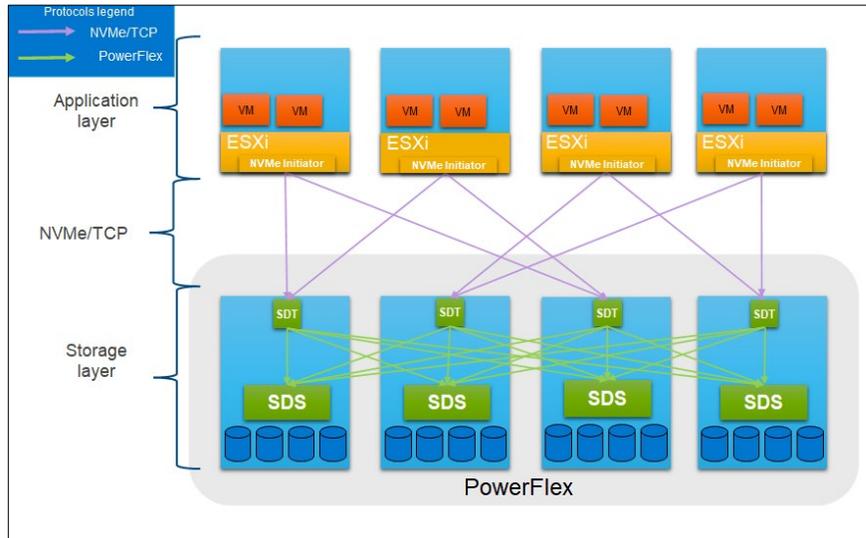


Figure 2. PowerFlex with NVMe/TCP

For more information, see [PowerFlex with NVMe over TCP](#).

PowerFlex deployment architectures

PowerFlex software-defined infrastructure excels in deployment flexibility. PowerFlex can be deployed in a two-layer (independent compute and storage layers), single-layer (Hyperconverged Infrastructure, or HCI), or a mixture of the two architectures (Mixed).

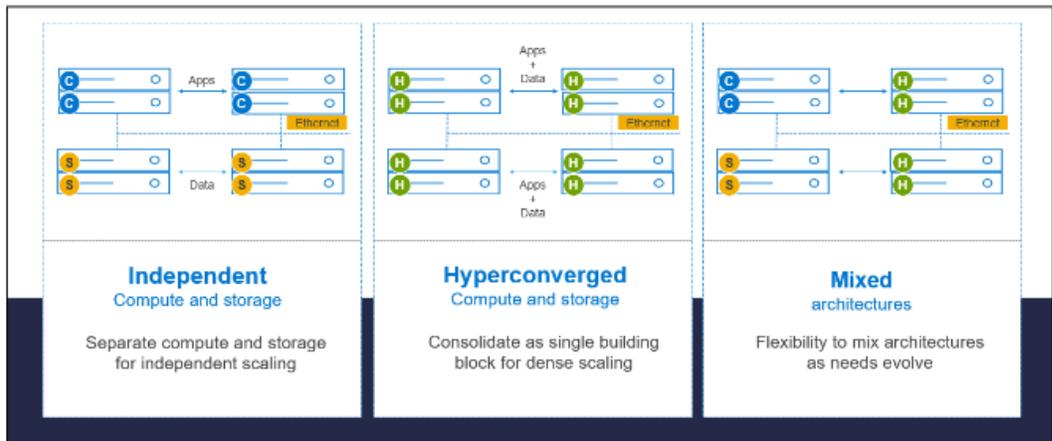


Figure 3. PowerFlex deployment architectures

Independent architecture

In an independent architecture or two-layer architecture, some nodes provide storage capacity for data in applications, while other separate and independent nodes provide compute resources for applications and workloads. Compute and storage resources can be scaled independently by adding nodes to the cluster while it remains active. This separation of compute and storage resources helps to minimize software licensing costs in certain situations. This architecture can be ideal for high-performance databases and application workloads.

Hyperconverged architecture

In an HCI architecture, each node in the cluster contributes storage and compute resources simultaneously to the applications and workloads. This architecture allows you to scale your infrastructure uniformly with building blocks that add both storage and compute resources. This architecture is appropriate for data center and workload consolidation.

Mixed architecture

A mixed architecture is a combination of both the HCI and Independent architectures. As shown in Figure 3 there are some storage-only nodes, compute-only nodes, and HCI nodes as a part of the same PowerFlex cluster. This is a desirable architecture when working with an existing compute infrastructure and adding high-performance software-defined infrastructure. This can also be a starting point for a two-layer deployment design as external workloads are migrated to PowerFlex.

PowerFlex consumption options

PowerFlex rack

PowerFlex rack is a software-defined infrastructure platform that delivers flexibility, elasticity, and simplicity with predictable performance and resiliency at scale by combining compute as well as high-performance storage resources in a managed unified network. This rack-based engineered system, with integrated networking, enables customers to achieve the scalability and management requirements of a modern data center.

PowerFlex appliance

PowerFlex appliance is a PowerEdge server which has been configured to be a node in a software-defined infrastructure deployment that runs PowerFlex software components. This offering allows customers the flexibility and savings to bring their own compatible networking.

PowerFlex custom nodes

PowerFlex Custom Nodes are validated server building-blocks configured for use with PowerFlex. They are available with thousands of configuration options and are available for customers who prefer to build their own environments.

Dell APEX Block Storage for Public Cloud

PowerFlex software is also available in the public cloud as Dell APEX Block Storage for Public Cloud. It offers the same on-premises benefits of high-performance, linear scalability, and high resilience as in cloud. It also adds cloud-specific benefits, such as large volume sizes, extreme performance based on NVMe drives, and predictable scalability. The Dell APEX Block Storage self-healing architecture simplifies and increases the availability of infrastructure. It is designed to remain running even when it encounters failures. With fast rebuilds, Dell APEX Block Storage recovers rapidly from failures and keeps operating, without degrading performance or compromising availability. You can also extend this availability and configure it to protect your application availability across multiple availability zones.

PowerFlex CSI/CSM

An important component outside of PowerFlex that enables a flexible consumption model for Kubernetes is the PowerFlex CSI driver, developed as a part of the Dell Kubernetes strategy. After loading the CSI driver for PowerFlex into Kubernetes, it can be used to

provision persistent volumes from the underlying PowerFlex storage resource. If the Kubernetes deployment is running low on PowerFlex storage resources, you can add PowerFlex storage nodes to increase the system capacity and performance.

The CSI driver connects the PowerFlex system and Kubernetes deployments. It is a storage broker which dynamically provisions volumes from PowerFlex through the PowerFlex API gateway to the Kubernetes cluster. Once the volume is available on PowerFlex, it is immediately mapped to the requesting pod. If a pod is destroyed or rescheduled, the CSI plug-in ensures that the volumes are remapped upon rescheduling of that pod.

- Customers running Kubernetes clusters on PowerFlex use the Dell Container Storage Modules (CSM), which extend the CSI driver capabilities. These modules:
- Provide enterprise storage capabilities to Kubernetes for cloud-native stateful applications.
- Reduce management complexity, so that developers can independently use enterprise storage with ease and automate daily operations.
- Extends storage functionality and capabilities beyond using the CSI driver alone.

These modules include snapshot, observability, authorization, application mobility, and resiliency.

PowerFlex supports multiple operating systems, different deployment options on-premises and public cloud deployment model (available in AWS and Azure). PowerFlex is validated with the leading Kubernetes distributions as shown in the following figure:

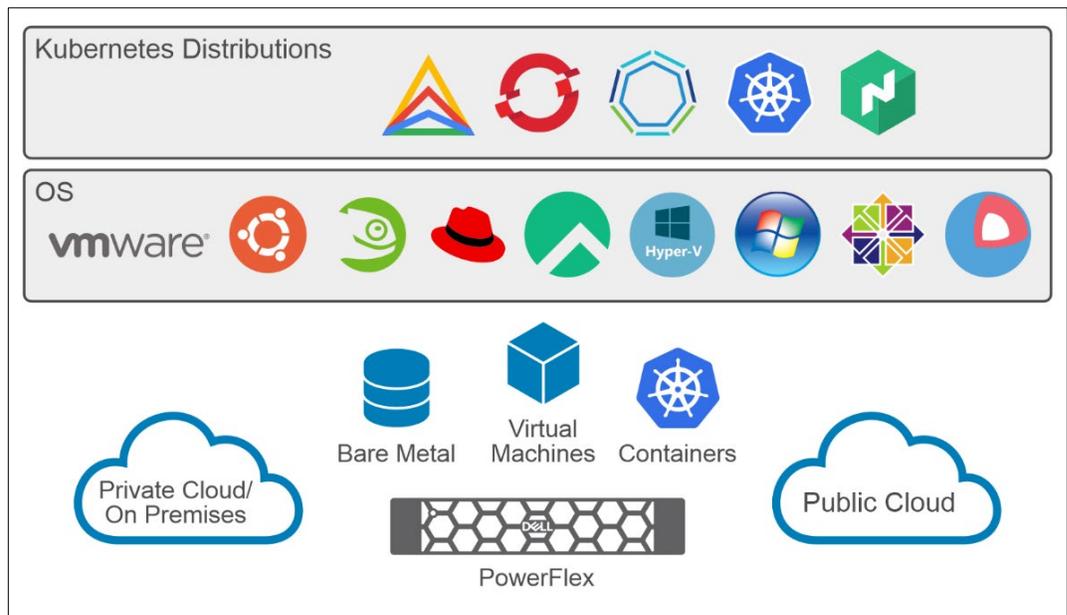


Figure 4. PowerFlex for different Kubernetes distributions

Apache CloudStack

Apache CloudStack is an open-source cloud software platform which is designed to deploy networks of scalable virtual machines in an Infrastructure-as-a-Service environment. As an open-source product, the platform can be molded to suit the needs of the customer. CloudStack can also be used by businesses that want to provide their own private cloud and hybrid cloud services on-premises.

CloudStack includes a compute function that allocates virtual machines (VMs) to individual servers, a network function uses software-defined virtual routers to create and manage logical networks, object and block storage systems, an image management function and a cloud computing management interface that supports all the software stack components.

Service providers and enterprises can use CloudStack to spin up elastic cloud computing services in an on-demand mode. Rather than managing virtual machines in the same way as physical machines, with CloudStack, enterprises can offer self-service virtual machines to the users without involving the IT department. Service providers can sell self-service virtual machine instances, storage volumes, and networking configurations over the Internet.

Users can manage their cloud deployments through a command line interface (CLI), a RESTful API, or Web interface. CloudStack application program interface (API) is compatible with Amazon EC2 and supports Terraform and Kubernetes Cluster API providers to facilitate hybrid cloud deployments.

For more information, see [Apache CloudStack documentation](#).

Solution overview

Overview

This section describes the architecture and design used to demonstrate CloudStack integration with PowerFlex.

In this solution, PowerFlex is deployed in a two-layer architecture, where the storage-only nodes are configured as the block storage system for the CloudStack environment.

The PowerFlex compute-only nodes, containing the SDC driver are configured as CloudStack KVM hypervisor nodes, which provide the compute platform for virtual machines (VMs).

Logical architecture

The following diagram shows the logical architecture of this solution:

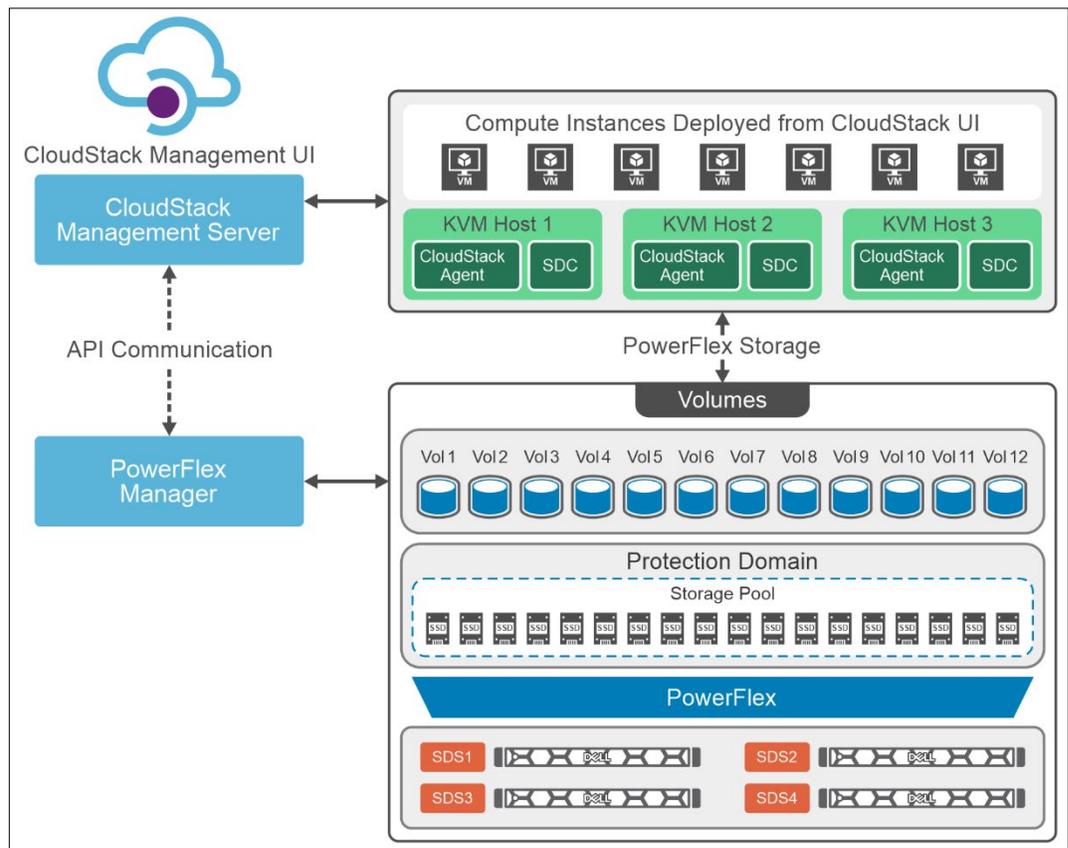


Figure 5. Logical design of CloudStack with PowerFlex

PowerFlex Manager is used to deploy the PowerFlex storage-only (SO) cluster on four bare metal servers followed by the deployment of PowerFlex compute-only (CO) nodes. PowerFlex storage-only cluster deployment is fully automated by the PowerFlex Manager, whereas the PowerFlex CO nodes are deployed manually by installing the KVM hypervisor followed by installing the SDC driver. In further sections, these CO nodes are referred to as KVM hosts.

A single PowerFlex protection domain with multiple storage pools is configured on PowerFlex. The PowerFlex storage pool topology consists of four storage-only nodes with 10 SSD disks in each of the four SO nodes in the PowerFlex system.

A CloudStack installation consists of two parts: the first is the Management Server and the second is the KVM hypervisor, as the managed cloud infrastructure. Once it is configured and managed as CloudStack cloud, resources such as hosts, storage devices, and IP addresses can be provisioned into the Management Server for consumption.

In this white paper, a single Management Server along with three KVM hosts are used for CloudStack setup with basic network configuration.

The PowerFlex SDC driver is installed on each of the KVM hypervisor hosts to enable PowerFlex to serve as the primary storage. The storage plugins feature in the CloudStack Management Server enables the use of PowerFlex storage pool as a managed primary storage for the KVM hypervisor, either as a zone-wide or cluster-wide pool. This pool can be added either from the UI or the API. Whenever VM or any storage related operations are performed, the respective volumes are created at the PowerFlex storage level.

Administrators and users can perform day-to-day operations such as VM provisioning, snapshots, migration, and so on, using the CloudStack UI.

For more information about the configuration of CloudStack and PowerFlex used in this solution, see [Configuration details](#).

CloudStack architecture

Apache CloudStack is a highly scalable IaaS cloud computing platform that provides an orchestration layer, automating the creation, provisioning, and configuration of IaaS components. CloudStack turns the existing virtual infrastructure into a cloud-based Infrastructure as a Service (IaaS) platform.

CloudStack works with multiple hypervisors, and a single deployment can contain multiple hypervisor implementations.

CloudStack deployment consist of the Management Server and the resources to be managed. The Management Server itself can be deployed in a multi-node installation where the servers are load balanced. MySQL can be configured to use replication to provide for failover in the event of database loss.

CloudStack offers an administrator web interface used for provisioning and managing the cloud and an end-user web interface that is used for running VMs and managing VM templates. CloudStack also provides a REST-like API interface for the operation, management, and use of the cloud.

Management Server

The Management Server orchestrates and allocates the resources. It runs in an embedded Jetty server container and requires a MySQL database to store data.

For this paper, a single Management Server is used, and the MySQL database is installed locally on that same VM.

Cloud infrastructure

The infrastructure consists of minimum one host running the hypervisor software.

Resources within the cloud are managed as follows:

- Regions: A collection of one or more geographically proximate zones.

- Zones: A zone is equivalent to a single data center. A zone consists of one or more pods and secondary storage.
- Pod: A pod is a rack of hardware that includes a layer-2 switch and one or more clusters.
- Cluster: A cluster consists of one or more hosts and primary storage.
- Host: A single compute node within a cluster.
- Primary storage is associated with a cluster and provides virtual disk volumes for all the VMs running on hosts in that cluster.
- Secondary storage is associated with a zone, and it stores templates, ISO images, and disk volume snapshots.

For more information about the terminologies, see [CloudStack Terminology](#)

The following figure shows a high-level logical architecture of CloudStack at zone level:

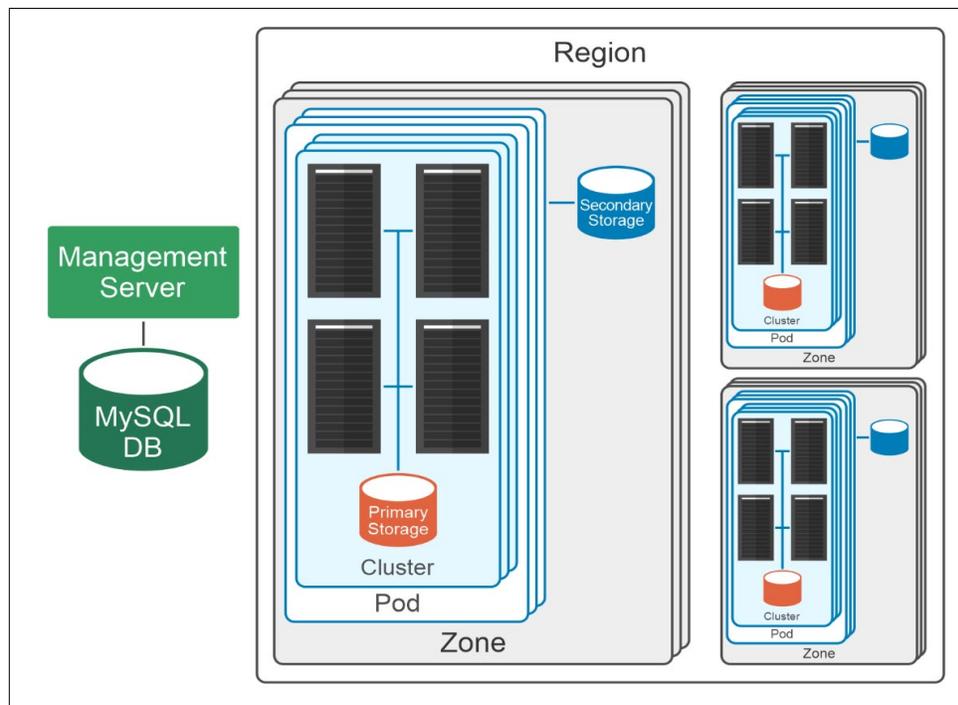


Figure 6. CloudStack Architecture

Note: For more information, see [CloudStack Architecture overview](#).

Network architecture

The following diagram shows the network architecture of this solution.

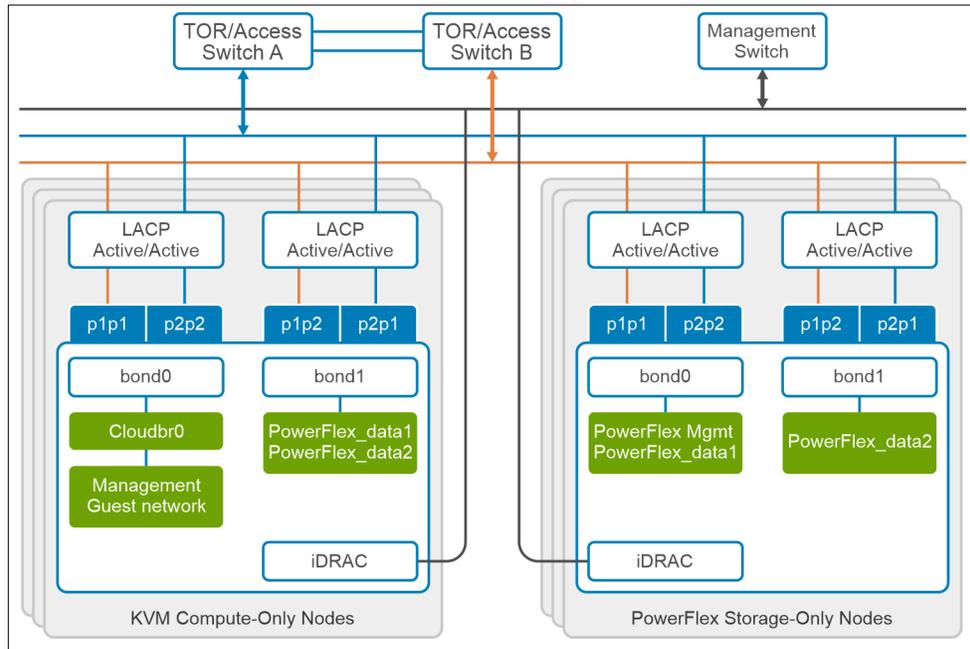


Figure 7. Network architecture

The following table describes the nodes in the architecture:

Node	Description
KVM compute-only node: bond0	This bond is configured with bridge networking for: Management and guest traffic
KVM compute-only node: bond1	This bond is for PowerFlex data networks. It is used for communication between PowerFlex compute and storage nodes.
PowerFlex storage-only node: bond0	This bond interface is for PowerFlex storage management and data network.
PowerFlex storage-only node: bond1	This bond interface is for the PowerFlex data network.

Table 2. Network description

At the physical layer, two supported Top of Rack (TOR) switches are used for redundancy and load-balancing purposes. A peer link is configured on both the TOR switches. VLANs are created to separate different traffic types on the network bonds.

CloudStack has two networking models, basic and advanced. This is decided based on the requirement. For this whitepaper, a basic networking model is used with the shared network concept.

In a Basic networking setup, the guest instances use the same IP range as the underlying CloudStack and hypervisor architecture.

For this whitepaper, CloudStack network bridges are used with KVM to connect the system VMs to the infrastructure, and guest instances to each other and the outside world.

The following figure shows an example of one of the KVM node network configurations:

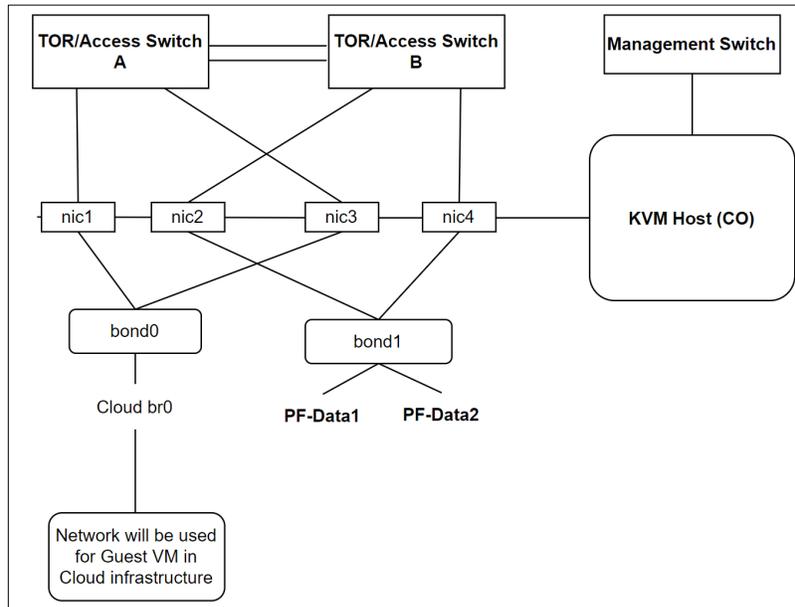


Figure 8. Example of a KVM Host network configuration

Note: It is important to maintain a consistent configuration across all the hypervisor hosts.

For more information about networking details, see [Beginner guide to CloudStack Networking](#) and [Network setup](#).

The following table shows the different networks that are configured for this solution:

Network type	VLAN ID
Management	105
PowerFlex Management	150
PowerFlex Data 1	152
PowerFlex Data 2	153
Guest network	Based on the requirement

Table 3. Network details

Note: OS Management network must be routable to the PowerFlex Management network to access the PowerFlex cluster.

Deployment of CloudStack

Overview

The deployment of CloudStack infrastructure involves the following steps:

- Installation and configuration of CloudStack Management Server.
- Setup of KVM hosts for infrastructure.

The following pre-requisites must be met before setting up the CloudStack infrastructure.

- DNS and NTP are required as a part of the infrastructure.
- PowerFlex storage cluster to be used as primary storage should be successfully deployed.

In this solution, a single Management Server is used along with three KVM hosts for CloudStack infrastructure with basic network configuration.

Management Server installation

This section describes the Management server installation process.

Prerequisites

The following requirements must be met before deploying:

Deploy a RHEL VM with the specified [Configuration details](#).

- VM should be configured with a fully qualified name (FQDN) and Internet access.
- Turn on and configure NTP for time synchronization.
- Ensure that the latest version of the Java package is installed.
- Set the SELinux to permissive.

Installation steps

The following steps describe the brief installation process necessary to deploy the Management Server. For more information, see [Management Server Installation](#).

1. Log in into the CloudStack management VM.
2. Run the following command to add the CloudStack repository (cloudstack.repo) under /etc/yum.repos.d

```
# vi /etc/yum.repos.d/cloudstack.repo
```

```
[cloudstack-4.18]
name=cloudstack
baseurl=http://packages.shapeblue.com/cloudstack/upstream/el
$releasever/4.18 [packages.shapeblue.com]
enabled=1
gpgcheck=1
gpgkey=http://packages.shapeblue.com/release.asc
[packages.shapeblue.com]
```

3. Run the following command to install the CloudStack management using dnf.

```
# dnf install cloudstack-management
```

4. The Management Server uses a MySQL database server to store its data. The following steps are used to install and configure the MySQL database:

- a. Run the following command to install the MySQL package.

```
# dnf install mysql-server
```

- b. Edit the MySQL configuration file `/etc/my.cnf.d/mysql-server.cnf` and insert the following lines in the `[mysqld]` section:

```
innodb_rollback_on_timeout=1
innodb_lock_wait_timeout=600
max_connections=350
log-bin=mysql-bin
binlog-format = 'ROW'
```

- c. Restart MySQL to put the new configuration into effect:

```
systemctl start mysqld
```

- d. Run the following command to secure the installation, answer **y** to all the questions:

```
# mysql_secure_installation
```

- e. `cloudstack-setup-databases` script is used for creating the CloudStack databases. The following command creates the **cloud** user in the database:

```
# cloudstack-setup-databases cloud:<DB
password>@localhost --deploy-as=root:<OS root_password>
Where: -
DBpassword: password to be assigned to the "cloud" user.
Set the password based on MySQL database password
policy.
```

`deploy-as:` specify the username and password of the user deploying the database. In the following command, it is assumed the root user is deploying the database and creating the "cloud" user.

Note: See [Install the database on separate nodes](#) to install the MySQL database on a standalone machine which is separate from the Management Server when using a deployment that includes several Management Server nodes.

5. Finish configuring the OS for the Management Server by running the following command:

```
# cloudstack-setup-management
```

6. The following message is displayed on completion: **CloudStack Management Server setup is done.**

Setting up secondary storage

CloudStack is designed to work with any scalable secondary storage system. Secondary storage uses the Network File System (NFS) protocol, as this ensures that the files can be

accessed by any host in the zone. Setup NFS share as secondary storage on the Management Server.

For more information, see [Secondary Storage](#).

Steps to setup secondary storage:

1. Run the following command to install the NFS utility package:

```
# dnf install nfs-utils
```

2. Run the following command to create a directory that will be used for secondary storage:

```
# mkdir -p /export/secondary
```

3. Run the following command to configure the new directories as NFS exports:

```
# vi /etc/exports
```

4. Insert the following line:

```
/export *(rw,async,no_root_squash,no_subtree_check)
```

5. Export the /export directory: `exportfs -a`.

6. Edit the `/etc/sysconfig/nfs` file. Uncomment the following lines to enable the ports used by CloudStack:

```
LOCKD_TCPPORT=32803
LOCKD_UDPPORT=32769
MOUNTD_PORT=892
RQUOTAD_PORT=875
STATD_PORT=662
STATD_OUTGOING_PORT=2020
```

7. Edit the `/etc/sysconfig/iptables` file if the firewall is not disabled. Add the following lines at the beginning of the INPUT chain, where `<NETWORK>` is the network that is used:

```
# -A INPUT -s <NETWORK> -m state --state NEW -p udp --dport
111 -j ACCEPT
# -A INPUT -s <NETWORK> -m state --state NEW -p tcp --dport
111 -j ACCEPT
# -A INPUT -s <NETWORK> -m state --state NEW -p tcp --dport
2049 -j ACCEPT
# -A INPUT -s <NETWORK> -m state --state NEW -p tcp --dport
32803 -j ACCEPT
# -A INPUT -s <NETWORK> -m state --state NEW -p udp --dport
32769 -j ACCEPT
# -A INPUT -s <NETWORK> -m state --state NEW -p tcp --dport
892 -j ACCEPT
# -A INPUT -s <NETWORK> -m state --state NEW -p udp --dport
892 -j ACCEPT
# -A INPUT -s <NETWORK> -m state --state NEW -p tcp --dport
875 -j ACCEPT
```

```
# -A INPUT -s <NETWORK> -m state --state NEW -p udp --dport
875 -j ACCEPT
# -A INPUT -s <NETWORK> -m state --state NEW -p tcp --dport
662 -j ACCEPT
# -A INPUT -s <NETWORK> -m state --state NEW -p udp --dport
662 -j ACCEPT
```

8. Restart iptables service.

```
# service iptables restart
# service iptables save
```

9. Reboot the Management Server host. Now the NFS secondary storage has been set.

System VM template

CloudStack is bundled with the system VM template and automatically seeded by Management Server to the KVM hosts when deploying a zone. Make sure that Management Server can Mount NFS secondary storage for this.

For more information about the template, see [System VM template](#).

CloudStack Infrastructure: KVM Host setup

This section describes the CloudStack infrastructure KVM Host setup.

Prerequisites

The following requirements must be met on all the hosts before the host setup:

- Based on the current [compatibility matrix](#), only KVM is supported with PowerFlex. In this solution, RHEL is used with KVM.
- Install hypervisor on all the hosts. Configure networking as required. For more information, see the [Network architecture](#) section.
- Update the fully qualified domain name, the host should have Internet access and turn on the NTP for time synchronization.
- Set SELinux to permissive.
- Ensure that the latest version of the Java package is installed.
- Each hypervisor could communicate with the other hypervisors and the Management Server. Verify the connectivity using the following ping command:

```
ping <hypervisor or management server IP>
```
- For more information about the requirement, see [System Requirement for KVM Hypervisor Host](#).

PowerFlex SDC installation

Steps to install the SDC driver which is configured on all the KVM hypervisor hosts.

1. Before proceeding with the SDC installation, ensure that the PowerFlex data networks (PowerFlex data1 and data2) have been configured on the KVM hosts.
2. Run the following command to copy the PowerFlex SDC packages onto KVM hosts and install them:

```
# MDM_IP=<LIST_VIP_MDM_IPS> rpm -i <SDC_PATH>.rpm
```

where

<LIST_VIP_MDM_IPS> is a comma-separated list of all the MDM IP addresses or the virtual IP address of the MDM. Do not include tiebreaker addresses here.

<SDC_PATH> is the path where the SDC installation package is located.

3. Reboot the host.
4. Verify that the driver is loaded by running the following command:


```
# service scini status
```
5. Repeat steps 1 to 4 for all the KVM hosts.
6. Log in to the PowerFlex Manager GUI. Navigate to **Block**→**Hosts** and verify that the KVM hosts are listed as shown in the following figure:

<input type="checkbox"/>	Name	Host IP	Protocol Type
<input type="checkbox"/>	SDC:9923c1e500000000	192.168.151.41 +1	SDC
<input type="checkbox"/>	SDC:9923e8f500000001	192.168.151.42 +1	SDC
<input type="checkbox"/>	SDC:9923e8f600000002	192.168.151.44 +1	SDC

Figure 9. SDC Host lists

Verify NFS share accessibility

Each KVM host should have access to the NFS share which is created in the [Setting up Secondary storage](#) section.

1. Run the following commands to ensure that the NFS and rpcbind services are up:

```
# service rpcbind start
# servicenfs start
# chkconfig nfs on
# chkconfig rpcbind on
# reboot
```

2. Mount the /export directories as shown below:

```
# mkdir /secondary
# mount -t nfs <management-server-name>:/export/secondary
# umount /secondary
```

Install and configure the CloudStack agent

The CloudStack agent communicates with the Management Server and controls the instances on the KVM hosts.

Run the following command to install the agent:

```
# dnf install -y epel-release
# dnf install cloudstack-agent
```

Note: if a non-root user is used to add KVM hosts UI, add the user to the sudoers file:

```
cloudstack ALL=NOPASSWD: /usr/bin/cloudstack-setup-agent
Defaults:cloudstack !requiretty
```

Install and Configure libvirt

Libvirt is required by the CloudStack-agent and should already be installed. It is vital that libvirt be configured properly as CloudStack uses libvirt for managing the virtual machines.

1. To avoid any potential security risk to the VMs, turn off the libvirt unsecure TCP port. While adding KVM hosts, CloudStack management uses its CA framework to install TLS certificates that are also used to secure agent-mgmt server, libvirtd and qemu/vnc to enable mutual-TLS secured connections.

Set the following parameters in `/etc/libvirt/libvirtd.conf`:

```
listen_tls = 0
listen_tcp = 0
tls_port = "16514"
tcp_port = "16509"
auth_tcp = "none"
mdns_adv = 0
```

2. Uncomment the following line in `/etc/sysconfig/libvirtd`:

```
LIBVIRTD_ARGS="--listen"
```

3. Update the changes by running the following command:

```
systemctl mask libvirtd.socket libvirtd-ro.socket libvirtd-
admin.socket libvirtd-tls.socket libvirtd-tcp.socket
```

4. Restart libvirt: `systemctl restart libvirtd`

Configuring the firewall

If firewalld is not disabled in the KVM hosts, then ensure that you open the following TCP ports for proper communication between KVM hosts and Management Server. Also update the settings in iptables and keep it persistent across reboots. For more information, see [Configuring the firewall](#).

- 22 (SSH)
- 1798
- 16514 (libvirt)
- 5900 - 6100 (VNC consoles)
- 49152 - 49216 (libvirt live migration)

Now the host is ready to be added to the cluster.

Accessing the CloudStack UI

For more information about the configuration, see [KVM Host Installation](#).

CloudStack provides a web-based UI that can be used by both administrators and end users. The UI uses API auto-discovery to discover APIs allowed for a logged-in user. The CloudStack UI helps users of cloud infrastructure to view and use their cloud resources, including instances, templates and ISOs, data volumes and snapshots, guest networks, and IP addresses.

Access the URL using Management Server IP: <http://<management-server-ip-address>:8080/client>

In the login screen as shown in [Figure 10](#), specify the following to go to the UI dashboard:

- Username is the User ID of an individual account. The default username is **admin**.
- The password is associated with the User ID. The password for the default username is **password**.
- Domain: For admin leave this field blank. For users in the sub-domains, enter the full path to the domain.



Figure 10. CloudStack UI

CloudStack configuration

This section describes how to add regions, zones, pods, clusters, hosts, storage, and networks to the CloudStack environment for day-to-day operations.

Note: For more information about the CloudStack infrastructure, see [CloudStack Architecture](#). For detailed information about CloudStack configuration, see [Configuring CloudStack Installation](#).

1. Log in to CloudStack UI. Navigate to the left and choose **Infrastructure**.
2. On the Zones section, click **Add Zone**. The zone creation wizard appears.
3. Choose the Zone type as **Core** and click **Next**.

4. Select **Basic** networking type, as all the host are configured to use basic network configuration. See the [Network Architecture](#) section for networking information. Click **Next**.
5. Enter the following **zone details** and click **Next**:
 - Name: A name for the zone.
 - DNS: This is the DNS servers for use by guest VMs in the zone.
 - Internal DNS: This is DNS servers for use by system VMs in the zone. These DNS servers are accessed by the management traffic network interface of the System VMs.
 - Hypervisor: Choose the hypervisor for the first cluster in the zone.
 - Network Offering: Select offering for shared network.

1 2 3 4 5 6
 Zone type Core zone type **Zone details** Network Add resources Launch

A zone is the largest organizational unit in CloudStack, and it typically corresponds to a single datacenter. Zones provide physical isolation and redundancy. A zone consists of one or more pods (each of which contains hosts and primary storage servers) and a secondary storage server which is shared by all pods in the zone.

* Name:

* IPv4 DNS1:

IPv4 DNS2:

* Internal DNS 1:

Internal DNS 2:

* Hypervisor: KVM ▼ ✓

Network offering: Offering for Shared Security group enabled networks ▼

Network domain:

Dedicated:

Enable local storage for user VMs:

Enable local storage for system VMs:

Figure 11. Zone details

6. Update the network details.
 - a. Choose the **traffic types** which will be carried by the physical network such as Guest and Management. Click the **Edit** button under the traffic type icon and **update traffic label** to **cloudbr0**. These labels must match the labels which are already defined on the hypervisor host which carries the respective traffic. Click **Next**.

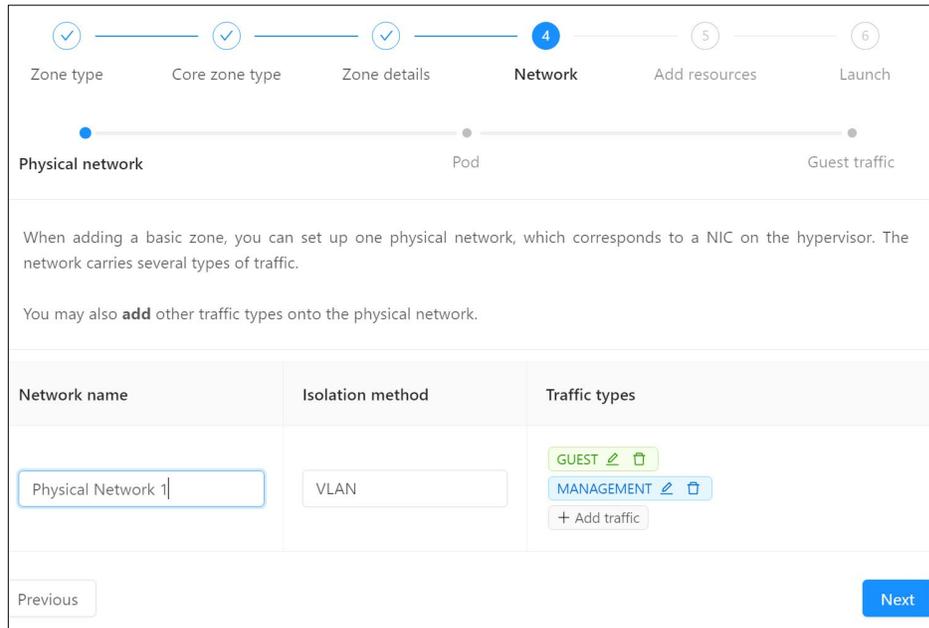


Figure 12. Network update

- b. In a new zone, CloudStack adds the first pod and multiple pods can be added later. Enter the following details for the **Pod** and click **Next**.
 - Pod Name: A name for the pod.
 - Reserved system gateway: The gateway for the hosts in that pod.
 - Reserved system netmask: The network prefix that defines the pod’s subnet.
 - Start/End Reserved System IP: The IP range in the management network that CloudStack uses to manage various system VMs, such as Secondary Storage VMs, Console Proxy VMs, and DHCP.

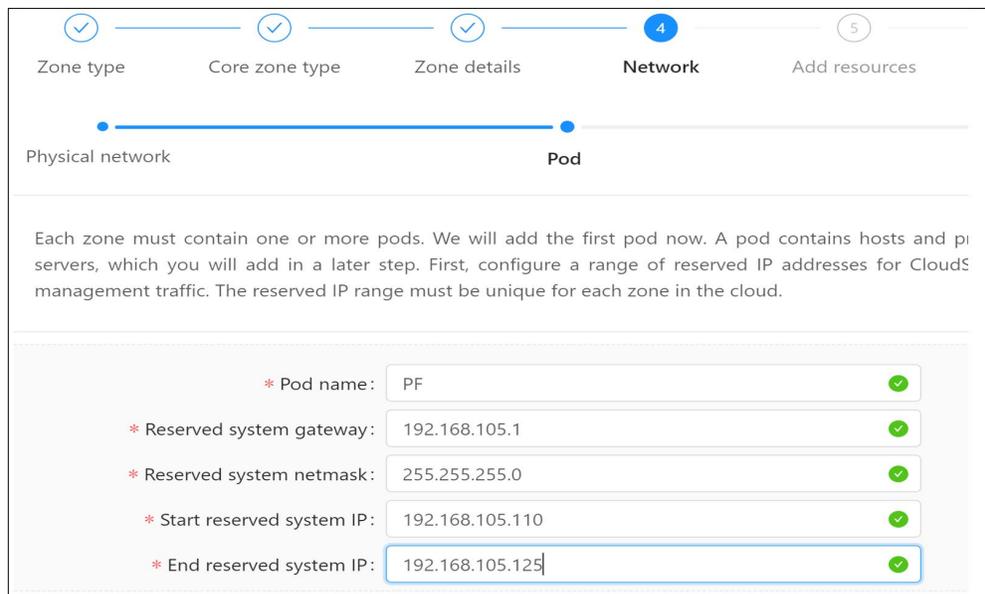


Figure 13. Pod details

7. Configure the network for **guest traffic**. Provide the following details and then click **Next**:
 - **Guest gateway**: The gateway that the guests should use.
 - **Guest netmask**: The netmask in use on the subnet that the guests use.
 - **Guest starts IP/End IP**: Enter the first and last IP addresses that define a range that CloudStack can assign to guests.

Note: Since the basic network type is selected with a shared network model, the same gateway has been used for both management and guest network. Provide different IP ranges for management and guest networks to ensure that there is no network conflict.

Figure 14. Guest network details

8. Provide a **Cluster name** and Click **Next**.
9. To configure the first **host**, enter the host IP, its credentials, and tag if required. Click **Next**.

Note: When adding a hypervisor host to CloudStack, the host must not have any VMs already running.

10. In a new cluster, CloudStack adds the first primary storage server. One can always add more servers later. For adding PowerFlex as **primary storage** provide the following details and click **Next**:
 - **Name**: Enter a **name** for storage.
 - **Scope**: Define the scope for the storage, whether it serves the cluster or **zone** level.
 - **Protocol**: For PowerFlex storage, select **custom**.
 - **Provider**: Select **PowerFlex** as provider.
 - **Gateway**: Provide the **PowerFlex Manager IP**.

- Gateway username and password: Provide the PowerFlex Manager UI credentials.
- Storage pool: Provide PowerFlex storage pool name.

Note: Make sure storage pool is already defined in the PowerFlex storage and the KVM host can access the storage pool.

Add zone

Zone type Core zone type Zone details Network **Add resources** Launch

Cluster IP address Primary storage Secondary storage

Each cluster must contain one or more primary storage servers. We will add the first one now. Primary storage contains the disk volumes for all the VMs running on hosts in the cluster. Use any standards-compliant protocol that is supported by the underlying hypervisor.

* Name: Powerflex ✓

Scope: Zone ✓

* Protocol: custom ✓

* Provider: PowerFlex ✓

* Gateway: 192.168.105.15 ✓

* Gateway username: admin ✓

* Gateway password: ✓

* Storage pool: SP-SSD-1 ✓

Figure 15. Primary storage details

11. Adding secondary storage. Since the secondary storage/NFS is already defined in the CloudStack Management Server, update the details here and click **Next**.

Note: For secondary storage, see [Setting up Secondary storage](#).

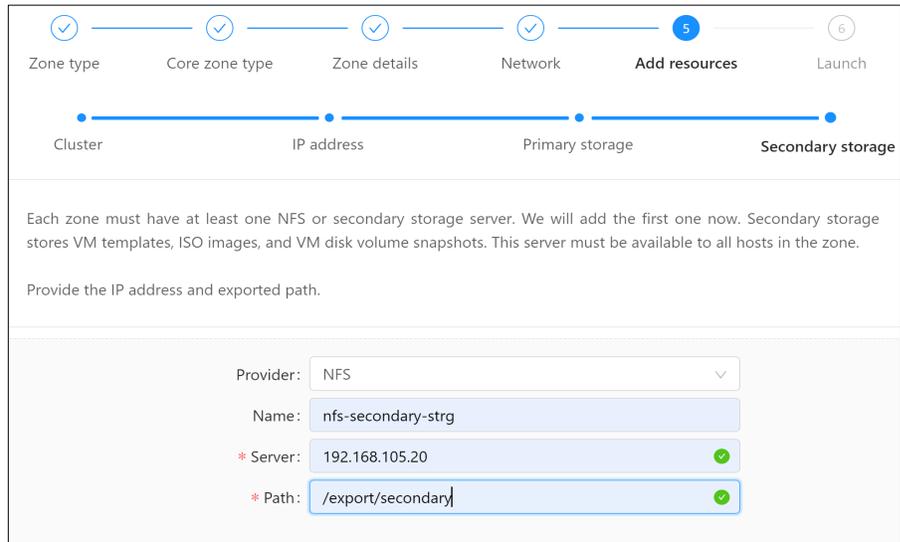


Figure 16. Secondary storage details

12. After updating everything correctly, a **Zone is ready to launch** message appears. Click on the **Launch zone** button.
13. The configuration takes time to complete. Once it is done, all the configurations can be viewed by navigating to **Infrastructure → Summary**.

Zones 1	Pods 1	Clusters 1	Hosts 3
Primary storage 3	Secondary storage 1	System VMs 2	Virtual routers 1

Figure 17. Summary of the infrastructure

Follow these steps if the VMs are not running or failed to run on the PowerFlex storage for the first time:

1. Navigate to **Infrastructure → Zone** → click **Disable zone**.

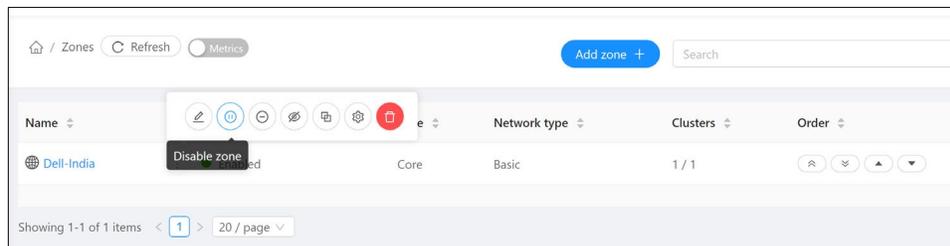


Figure 18. Disabling the zone

2. Disable the primary storage pool by selecting the storage pool, **Infrastructure → Primary storage** → select storage pool → **Disable storage pool**.
3. Next, navigate to **Infrastructure → System VMs** → select VM (secondary storage VM and console proxy VM) → **Destroy the System VMs**.

4. Enable the Zone and CloudStack automatically creates a new System VM. Since primary storage is disabled, now VMs can use the default secondary storage.
5. Once the VMs come up, navigate to **Infrastructure** → **Primary storage** → select storage pool → **Enable storage pool**. The System VMs are automatically migrated to use Primary storage.

Note: When creating a new zone, CloudStack provides an option to create other configurations as well such as pod, cluster, adding host, storage pool and so on. However, pods/cluster/host can be added anytime later by directly navigating to the respective section in the UI.

Adding more hosts

Make sure that the KVM is already installed on the host and configured with networking. Networking should be common across all the nodes inside a single zone.

Steps to add multiple KVM hosts in the existing cluster.

1. Log in to the CloudStack UI.
2. In the left navigation, select **Infrastructure**. In Hosts, click **Add host**.
3. Enter the following details and click **OK**:
 - **Zone, Pod, and cluster name:** Select the accordance in which host needs to be added.
 - **Host name:** DNS name or IP address of the host in URL format as shown in the following figure.
 - **Host username and Password:** Provide the credential of the host.
 - **Host tag:** It is optional, any labels that are used to categorize hosts for ease of maintenance.

Note: For a more secured way, a host can also be added using CloudStack's SSH key without having to provide a host password. Before adding the host in CloudStack do the following. Copy the SSH public key from `/var/lib/cloudstack/management/.ssh/id_rsa.pub` on the management server. Add the copied key to the `/root/.ssh/authorized_keys` file on the KVM host's.

Figure 19. Adding a host

- Repeat for additional hosts. Once done, verify that the hosts are added by navigating to the Host section as shown in the following figure:

Name	State	Resource state	IP Address	Hypervisor
cs-co-node4.pflex.com	: ● Up	● Enabled	192.168.105.44	KVM
cs-co-node2.pflex.com	: ● Up	● Enabled	192.168.105.42	KVM
cs-co-node1.pflex.com	: ● Up	● Enabled	192.168.105.41	KVM

Figure 20. KVM host list

Adding primary storage pool

When creating a new zone, CloudStack provides an option to create a primary storage. However, primary storage can be added at any time, such as when adding a new cluster or adding more servers to an existing cluster.

Note: When using pre-allocated storage for primary storage, make sure that there is no data residing in the storage pool.

Note: Make sure the storage pool is already configured in the PowerFlex storage and the KVM host can access the storage pool. Adding the storage to CloudStack destroys any existing data.

- Log in to the CloudStack UI.
- In the left navigation, choose Infrastructure. In Primary storage, click **add primary storage**.

3. Provide the following information and click **OK**.
 - **Scope:** Indicate whether the storage is available to all hosts in the zone or only to hosts in a single cluster.
 - **Pod** (available only if Cluster selected in the Scope field): The pod for the storage device.
 - **Cluster** (available only if Cluster is selected in the Scope field): The cluster for the storage device.
 - **Name:** Name for the storage.
 - **Protocol:** For PowerFlex storage. Select **custom**.
 - **Provider:** Select **PowerFlex** as provider.
 - **Gateway:** Provide **PowerFlex Manager IP**.
 - **Gateway username and password:** Provide PowerFlex Manager UI credentials.
 - **Storage pool:** Provide a PowerFlex storage pool name, which is already defined on the PowerFlex storage.
 - **Storage tags:** It is optional, it is a label that you use to categorize for ease of maintenance. If it is defined, then the tag set on primary storage across clusters in a zone must be identical.

The screenshot shows a dialog box titled "Add primary storage" with a close button (X) in the top right corner. The dialog contains the following fields and values:

- Scope:** Zone
- Hypervisor:** KVM
- * Zone:** Dell-India
- * Name:** PF-strg
- * Protocol:** custom
- * Provider:** PowerFlex
- * Gateway:** 192.168.105.15
- * Gateway username:** admin
- * Gateway password:** (masked with dots)
- * Storage pool:** SP1
- Storage tags:** the tags for the storage pool

At the bottom right of the dialog, there are two buttons: "Cancel" and "OK".

Figure 21. Adding a primary storage pool

Note: The tag set on primary storage across clusters in a Zone must be identical.

4. Verify the storage capacity of the added storage pool with PowerFlex side storage pool.
 - Log in to CloudStack UI and navigate to **Infrastructure**→ **Primary storage** and click on the respective **storage pool** to see the size as shown below.

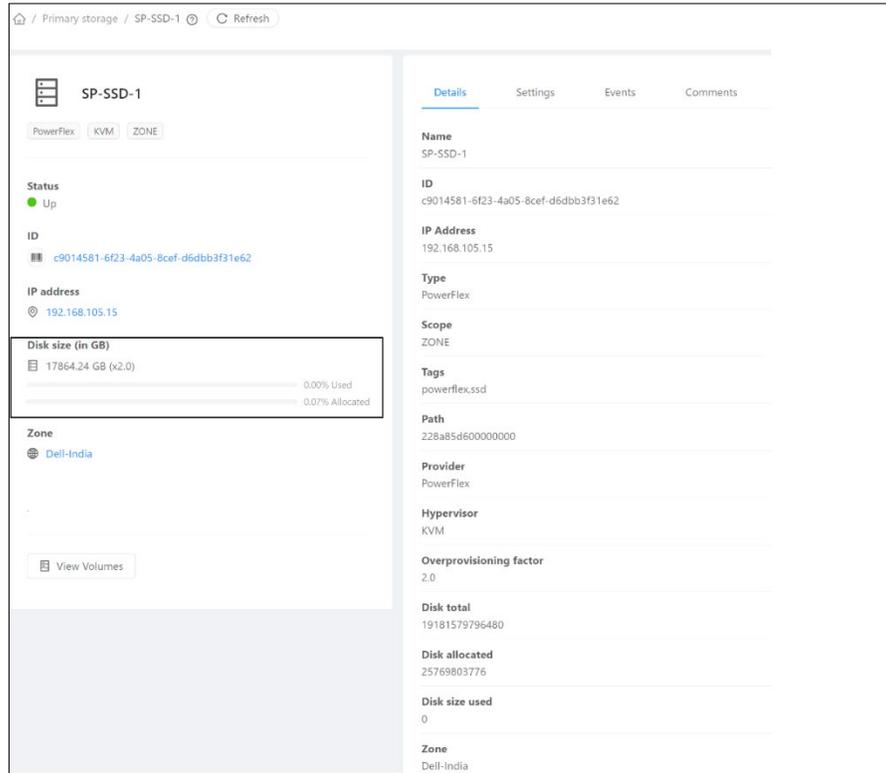


Figure 22. Capacity details of PowerFlex storage in CloudStack

- Log in to PowerFlex Manager, navigate to **Block**→ **Storage Pool** and select the respective storage pool. And then click on the **Capacity** tab to verify the capacity.

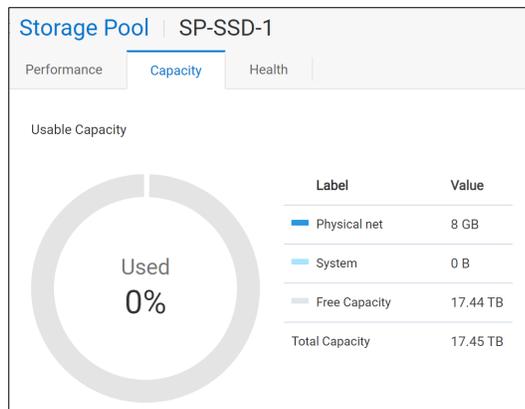


Figure 23. Storage pool capacity from the PowerFlex Manager

The user can add multiple storage pools and define the scope to use as either cluster or zone wide.

Name	State	IP Address	Scope	Type
SP2	● Up	192.168.105.15	CLUSTER	PowerFlex
SP1	● Up	192.168.105.15	ZONE	PowerFlex
SP-SSD-1	● Up	192.168.105.15	ZONE	PowerFlex

Figure 24. Storage pool list with different scopes

Other administrator and user configuration options

The CloudStack UI allows users of cloud infrastructure to view and use their cloud resources, including instances, templates and ISOs, data volumes and snapshots, guest networks and IP addresses.

A few major resource configurations are:

Service Offerings

A user creating a new Instance can make various choices about its characteristics and capabilities. CloudStack provides several ways to present users with choices when creating a new instance such as compute, disk, and network offerings. For more information about the offerings and its configuration, see [Service Offerings](#).

Templates

CloudStack comes with a default template for the CentOS operating system. There are many ways to add more templates, administrators, and end users can add multiple templates.

The default template is downloaded by the secondary storage VM after the primary and secondary storage are configured. Users can use this template for production deployment, or they can also use custom template by importing a new template/IOS.

For more information about the template and about creating or importing a new template, see [Working with Templates](#).

Instances

CloudStack provides several guest management operations for end users and administrators. Instances may be stopped, started, rebooted, and destroyed. The administrator has complete control over the lifecycle of all guest instances in the cloud. Instances are created from a template.

For more information, see [Working with Instances](#).

CloudStack and PowerFlex integration

The PowerFlex Gateway details provided when adding PowerFlex storage pool as the primary storage enables the CloudStack communication with the PowerFlex REST APIs for various storage operations.

When a PowerFlex storage pool is configured and used with Compute or Disk Offerings, an administrator can build an environment in which a root or data disk that a user creates leads to the dynamic creation of a PowerFlex volume.

The lifecycle operations of CloudStack resources, templates, volumes, and snapshots in a PowerFlex storage pool can be managed through the CloudStack UI. The following operations are supported for a PowerFlex storage pool:

- Volumes are created from the PowerFlex Storage in the form of root disk and data disk whenever the instance gets deployed.
- Instance snapshot (disk-only, snapshot with memory is not supported).
- Migrating instances from one host to another (within and across clusters, for zone-wide primary storage).
- Volume lifecycle and operations:
 - Create ROOT disks using the selected template.
 - List, Detach, Resize ROOT volumes.
 - Create, List, Attach, Detach, Resize, Delete DATA volumes.
 - Create, List, Delete snapshots of volumes (with backup in Primary, no backup to secondary storage).
 - Create a template from PowerFlex volume or snapshot.
 - Live migration of volume from one PowerFlex storage pool to another.
 - Migration of instance with storage: Offline.

Once instances and their volumes are created, the respective volume can be identified on the PowerFlex storage using the path of the volume as shown below.

- Log in to the CloudStack UI. Navigate to **Storage**→ **Volumes** and select the volumes and identify the path.

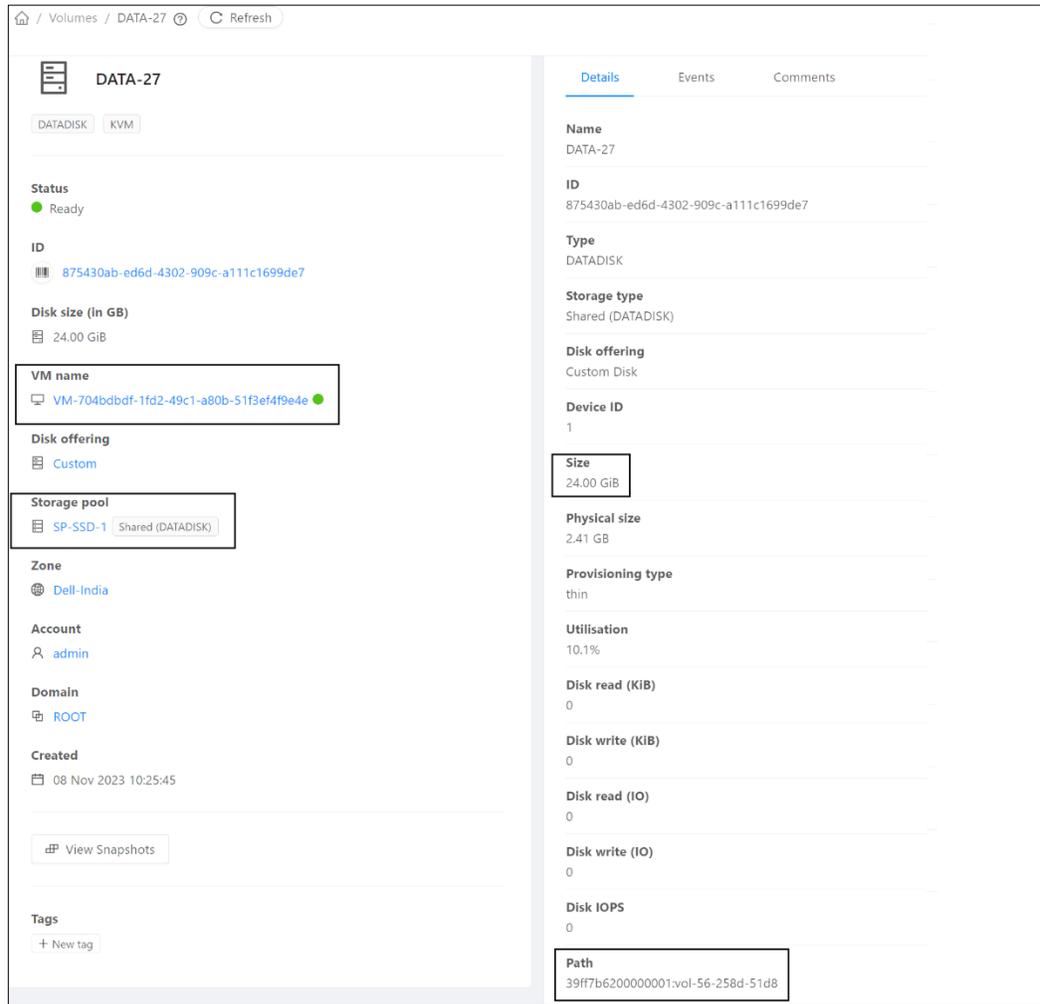


Figure 25. Identifying the volume path at CloudStack UI

- Log in to the PowerFlex Manager UI, navigate to **Block** → **Volumes**, cross verify the volume with the identified path name.

Name	Volume ID	Type	Size	Mapped	# Hosts	Hosts Protocol	Creation Time	Read Only	Storage Pool
vol-60-4581-51d8	39ff7b8500000004	Thin	8 GB	Yes	1	SDC	Nov 9, 2023, 11:29:43 PM	No	SP-SSD-1
vol-59-4581-51d8	39ff7b830000000b	Thin	24 GB	Yes	1	SDC	Nov 9, 2023, 2:13:35 AM	No	SP-SSD-1
vol-58-4581-51d8	39ff7b8200000006	Thin	8 GB	Yes	1	SDC	Nov 9, 2023, 2:12:21 AM	No	SP-SSD-1
vol-57-258d-51d8	39ff7b8100000005	Thin	24 GB	Yes	1	SDC	Nov 8, 2023, 3:43:17 AM	No	SP2
vol-56-4581-51d8	39ff7b6200000001	Thin	24 GB	Yes	1	SDC	Nov 7, 2023, 2:52:33 AM	No	SP-SSD-1
vol-51-4581-51d8	39ff7b7a00000002	Thin	8 GB	Yes	1	SDC	Nov 8, 2023, 1:45:31 AM	No	SP-SSD-1

Figure 26. Identified volume in the PowerFlex Manager UI

- The following image shows the list of volumes which are created on PowerFlex when the volumes, instances, snapshot, and templates are created from CloudStack using PowerFlex as the primary storage.

Volumes											+ Create Vol
Mapping ▾		Modify ▾		More Actions ▾						16 Volumes	
<input type="checkbox"/>	Name	Volume ID	Type	Size	Mapped	# Hosts	Hosts Protocol	Creation Time	Read Only	Storage Pool	
<input type="checkbox"/>	vol-23-258d-51d8	39ff7b3800000007	Thin	8 GB	Yes	1	SDC	Nov 6, 2023, 4:44:32 AM	No	SP2	
<input type="checkbox"/>	tmpl4-258d-51d8	39ff7b3b00000003	Thin	8 GB	No	0	NONE	Nov 6, 2023, 5:29:25 AM	No	SP2	
<input type="checkbox"/>	snap-10-258d-51d8	39ff7b7e00000009	Snapshot	24 GB	No	0	NONE	Nov 8, 2023, 1:47:29 AM	No	SP2	
<input type="checkbox"/>	vol-57-258d-51d8	39ff7b8100000005	Thin	24 GB	Yes	1	SDC	Nov 8, 2023, 3:43:17 AM	No	SP2	
<input type="checkbox"/>	vol-58-4581-51d8	39ff7b8200000006	Thin	8 GB	Yes	1	SDC	Nov 9, 2023, 2:12:21 AM	No	SP-SSD-1	
<input type="checkbox"/>	vol-59-4581-51d8	39ff7b830000000b	Thin	24 GB	Yes	1	SDC	Nov 9, 2023, 2:13:35 AM	No	SP-SSD-1	
<input type="checkbox"/>	vol-60-4581-51d8	39ff7b8500000004	Thin	8 GB	Yes	1	SDC	Nov 9, 2023, 11:29:43 PM	No	SP-SSD-1	
<input type="checkbox"/>	vmsnap-15-59-4581-51d8	39ff7b870000000d	Snapshot	24 GB	No	0	NONE	Dec 11, 2023, 3:37:41 AM	No	SP-SSD-1	
<input type="checkbox"/>	snap-12-4581-51d8	39ff7b890000000f	Snapshot	8 GB	No	0	NONE	Dec 11, 2023, 3:39:17 AM	No	SP-SSD-1	
<input type="checkbox"/>	vol-24-258d-51d8	39ff7b3900000008	Thin	8 GB	Yes	1	SDC	Nov 6, 2023, 4:45:14 AM	No	SP2	
<input type="checkbox"/>	vol-25-258d-51d8	39ff7b3a0000000a	Thin	8 GB	Yes	1	SDC	Nov 6, 2023, 4:45:54 AM	No	SP2	
<input type="checkbox"/>	vol-42-258d-51d8	39ff7b6100000000	Thin	8 GB	Yes	1	SDC	Nov 7, 2023, 2:52:21 AM	No	SP2	
<input type="checkbox"/>	vol-51-4581-51d8	39ff7b7a00000002	Thin	8 GB	Yes	1	SDC	Nov 8, 2023, 1:45:31 AM	No	SP-SSD-1	
<input type="checkbox"/>	vmsnap-15-58-4581-51d8	39ff7b860000000c	Snapshot	8 GB	No	0	NONE	Dec 11, 2023, 3:37:41 AM	No	SP-SSD-1	
<input type="checkbox"/>	vol-56-4581-51d8	39ff7b6200000001	Thin	24 GB	Yes	1	SDC	Nov 7, 2023, 2:52:33 AM	No	SP-SSD-1	
<input type="checkbox"/>	vmsnap-15-60-4581-51d8	39ff7b880000000e	Snapshot	8 GB	No	0	NONE	Dec 11, 2023, 3:37:41 AM	No	SP-SSD-1	

Figure 27. Volume list in the PowerFlex Manager UI

Conclusion

This whitepaper documents the steps to deploy the Apache CloudStack Infrastructure-as-a-Service platform and configure Dell PowerFlex as the primary storage layer. The integration of Dell PowerFlex software-defined infrastructure with Apache CloudStack presents a powerful solution for effectively navigating market trends such as multi-cloud adoption, flexibility, scalability, automation/orchestration and so on. This integration is key to seamlessly managing and scaling the cloud infrastructure in a rapidly evolving digital landscape.

References

Dell Technologies documentation

The following Dell Technologies documentation provides additional and relevant information. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [Dell PowerFlex product overview](#)
- [PowerFlex specification sheet](#)
- [PowerFlex Solutions Document – Info Hub](#)
- [Dell PowerFlex Software - Technical Documentation](#)
- [Dell PowerFlex Software Defined Infrastructure](#)

Apache CloudStack documentation

The following CloudStack documentation provides additional and relevant information:

- [CloudStack Documentation](#)
- [CloudStack Network Consideration](#)
- [Compatibility Matrix](#)
- [Solution Brief: IaaS Transformation: Apache CloudStack and Dell PowerFlex](#)

Appendix A

Configuration details

The following tables show the hardware and software configuration of storage and compute nodes that were used for Dell PowerFlex deployment as well as the CloudStack Management Server:

Table 4. CloudStack Management Server (Virtual Machine) configuration

Hardware	Configuration
vCPU	8
Memory (GiB)	16 GB
vNIC	VMXNET 3
Storage	Operating System disk: 100 GB For Secondary storage: 500 GB
Operating System	Red Hat Enterprise Linux 8.6
CloudStack Software version	4.18.1.0

Table 5. Storage node configuration

Hardware	Configuration
Server model	PowerFlex R650
CPU	2 x Intel(R) Xeon(R) Gold 6336Y CPU @ 2.40 GHz
Memory (GiB)	544 GB (16 x 32 GB + 2 x 16 GB NVDIMM)
NIC	4 ports 25 GB Mellanox Connect-X 5
Storage	10 x 1.4 TB SAS SSDs
Software	Configuration
Operating System	Embedded OS
Dell PowerFlex software release	4.5

Table 6. Compute node configuration

Hardware	Configuration
Server model	PowerFlex R740
CPU	2 x Intel(R) Xeon(R) Gold 6150 CPU @ 2.70 GHz
Memory (GiB)	384 GB (6 x 64 GB)
NIC	4 ports 25 GB Mellanox ConnectX-4 LX SFP
Storage	BOSS S1 Controller 2 x 224 GB SATA SSD
Software	Configuration
Operating System	RHEL 8.6 with KVM
Dell PowerFlex software release	4.5
CloudStack Software version	4.18.1.0