Intel® Solutions
Reference Architecture

VM Image Hosting Using the Fujitsu* Eternus CD10000 System with Ceph* Storage Software

Intel® Xeon® Processor E5-2600 v3 Product Family

SRA Section: Audience and Purpose

Fujitsu is one of the first companies to offer a turnkey solution for software-defined storage implementing an open source application stack known as Ceph*. The hardware is based on the Intel® Xeon® E5-2600 v2 family of high performance processors, Intel® Solid-State Drive Data Center Family, and 1 Gb and 10 Gb Intel® converged network adapters. Using those ingredients, Fujitsu has produced a fully validated, reliable, cost effective storage solutions, with the backing of a major corporation in the storage segment. This solution minimizes risk and provides an accelerated ramp for onsite IT at cloud service providers, and medium to large enterprises, especially in addressing the need to provide reliable block storage services for VM image hosting. Intel based solutions provide powerful, scalable and reliable platforms that are fully supported by Intel with an extensive roadmap into the future. A customer deciding to use a Fujitsu product with Intel Xeon® processors can be assured of improvements in speed, power and reliability year over year. Fujitsu's design effort and customer support provide the backing required in a solution of this complexity and size.

January 2015

Table of Contents

Executive Summary

This White Paper addresses Ceph based VM cloud hosting, as well as the general functionality and capabilities of the FUJITSU* Storage ETERNUS CD10000 hyper-scale system. Based on the Intel Xeon E5-2600 v2 family of processors, Intel Solid-State Drive Data Center Family, and 1 Gb and 10 Gb Intel Converged Network Adapters, the ETERNUS CD10000 provides virtually unlimited, modular scalability of storage capacity and performance with a zero downtime target for online access to data. Integrating open-source Ceph software into a storage system pre-configured and delivered, with end-to-end support from Fujitsu, enables IT organizations to fully benefit from open source software while minimizing implementation and operational risks. Providing hyper-scalable object, block, or file storage scalable to more than 50 Petabytes, the ETERNUS CD10000 is a viable solution for service providers, cloud, IT and telecommunications as well as media-broadcasting companies, financial and public institutions with ever-growing document repositories, large scale business analytics/big data applications as well as organizations with comprehensive multimedia data. ETERNUS CD10000 can serve all these segments.

Distributed Scale-out Storage

Scale-out storage is rapidly emerging as a viable alternative to address a wide variety of enterprise use cases, it allows the enterprise to add storage on a “pay-as-you-grow” basis for incremental capacity minimizing ongoing costs and better matching performance to requirements. Due to the explosion of unstructured data, the need to provide differentiated services, and the need for vendor provided professional support, demand will increase for scale-out object, block or file based storage.

Scalability

- Practically unlimited scalability in terms of performance & capacity
- No performance bottlenecks
- No hot spots regarding data access
- Zero planned or unplanned downtime

Reliability

- Full redundancy
- Self-healing capability
- Available geographical dispersion
- Fast rebuild of failed disks and nodes

Manageability

- Central management of storage
- Unified multi-protocol access (block, file and object)
- Seamless introduction of additional storage capacity

While traditional scale-up storage systems distribute volumes across sub-sets of spindles, scale-out systems use mathematical algorithms to distribute volumes across all or many spindles to provide maximum utilization of system resources. Comparing the distributed scale-out storage solution versus the traditional approach, the distributed scale-out architecture shows clear benefits.

The data allocation scheme distributes data across multiple disks. The separation in figure 1 shows the different storages nodes, which are also part of the data distribution. The objects, files or block data are allocated homogeneously. Disks or node failures can occur without losing data or affecting performance. The system thus has very high fault tolerance, elasticity, and scalability of capacity and data migration capabilities with zero downtime.

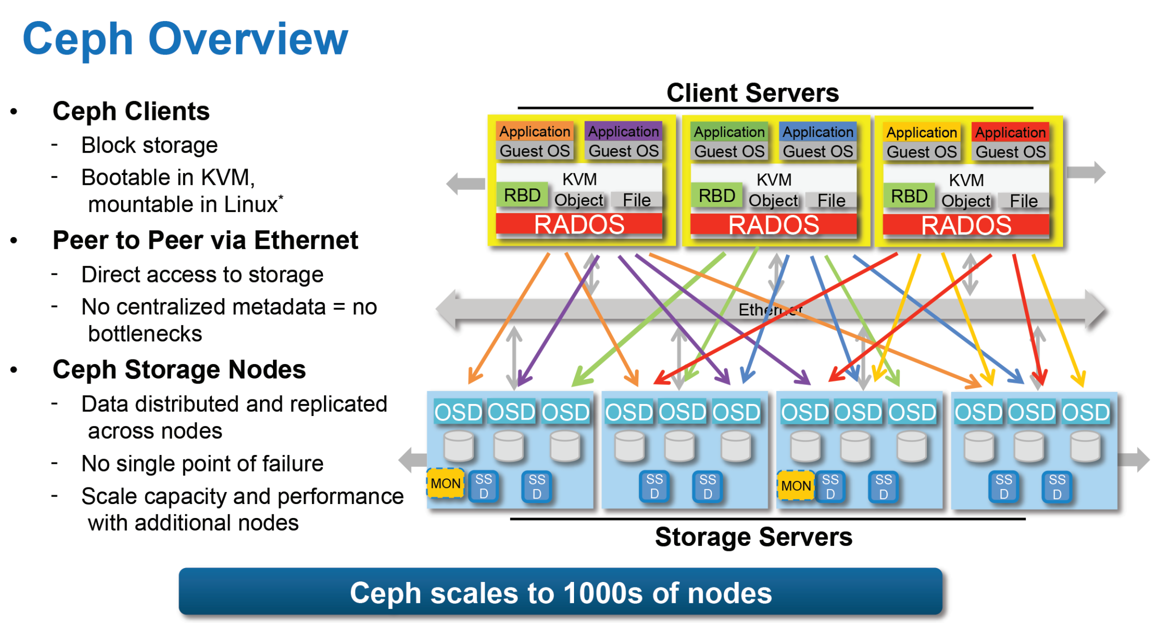

Figure 1 – Ceph Overview

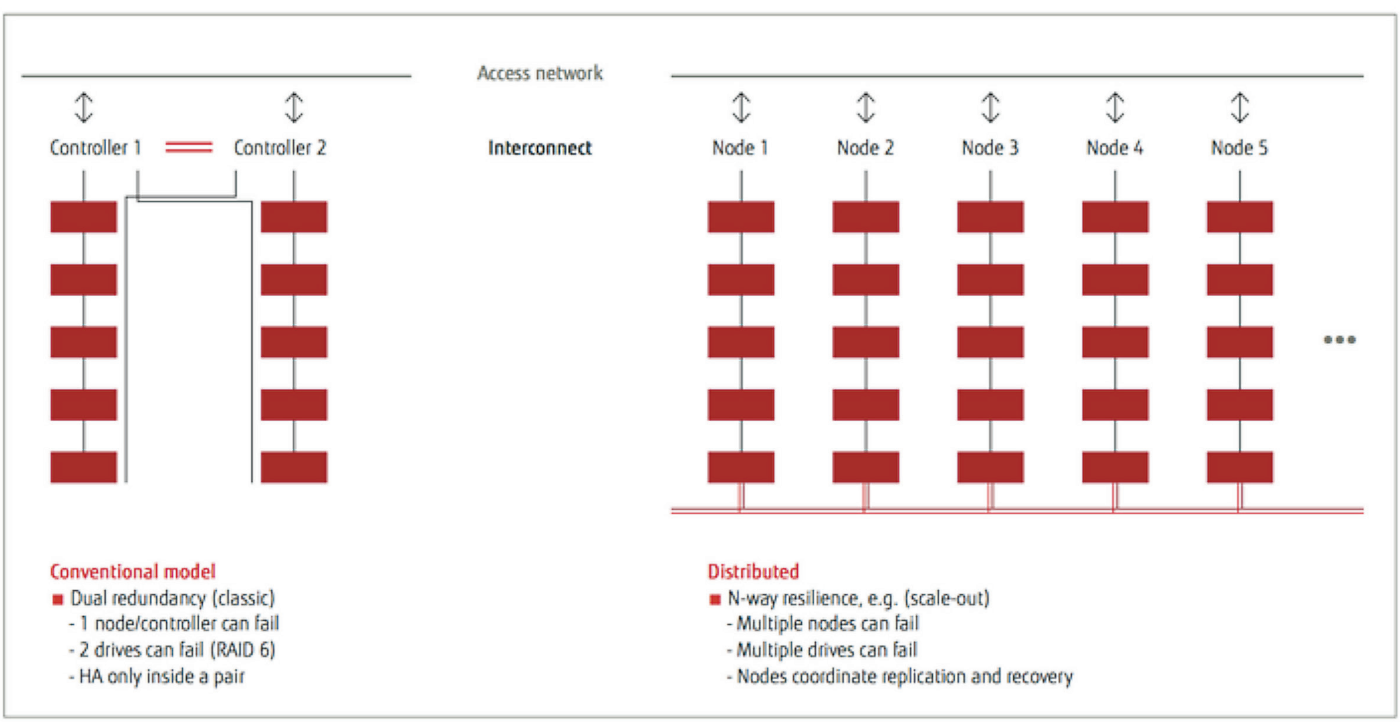

Figure 2 – Availability and Scalability Through Distributed Scale-Out Model

High Availability

In the distributed scale-out architecture of ETERNUS CD10000, multiple storage nodes and drives can fail because of the inherent fault tolerance of the architecture without impacting business continuity. All nodes can automatically coordinate the recovery in a short, automatic, self-healing process.

The system is fault tolerant by design using multiple compute platforms based on Intel Xeon® E5 processors. Planned downtime for maintenance or for technology upgrades is no longer required or necessary. As storage nodes of a particular generation are being updated with nodes of the next generation, during operation the lifecycle of the solution can be extended and migration efforts can be reduced.

Architecture

Figure 3 shows the architecture building blocks of the ETERNUS CD10000 system. The system hardware consists of Intel Xeon® E5-2640v2 processors in the storage nodes, a fast, low latency Infiniband backend network and a 10 Gb Ethernet front-end network which is crucial for consistent data access and performance.

The internal system management based on Ceph open-source software plus VSM and related enhancements from Fujitsu to improve and enhance ease of use. The management system is responsible for operating the complete system platform with all its nodes from a single instance. The software design enables unified object, block and file access. Fujitsu also supports user specific interfaces for cloud storage, synch & share, archiving and file service as options. Here's an overview of key solution characteristics:

- Up to 200 Dual Intel Xeon® E5-2640 v2 processor based storage nodes, able to deliver over 50 Petabytes raw capacity.
- Host interface: 10 Gb Ethernet – front-end access network
- Intra-node interconnect via Infiniband 40 GB/s
- Ceph Storage Software
  - » VSM software enhancements for ease-of-management
- Unified access for object, block and file storage
  - » Central management system based on the Intel Xeon® ES-2640 v2.

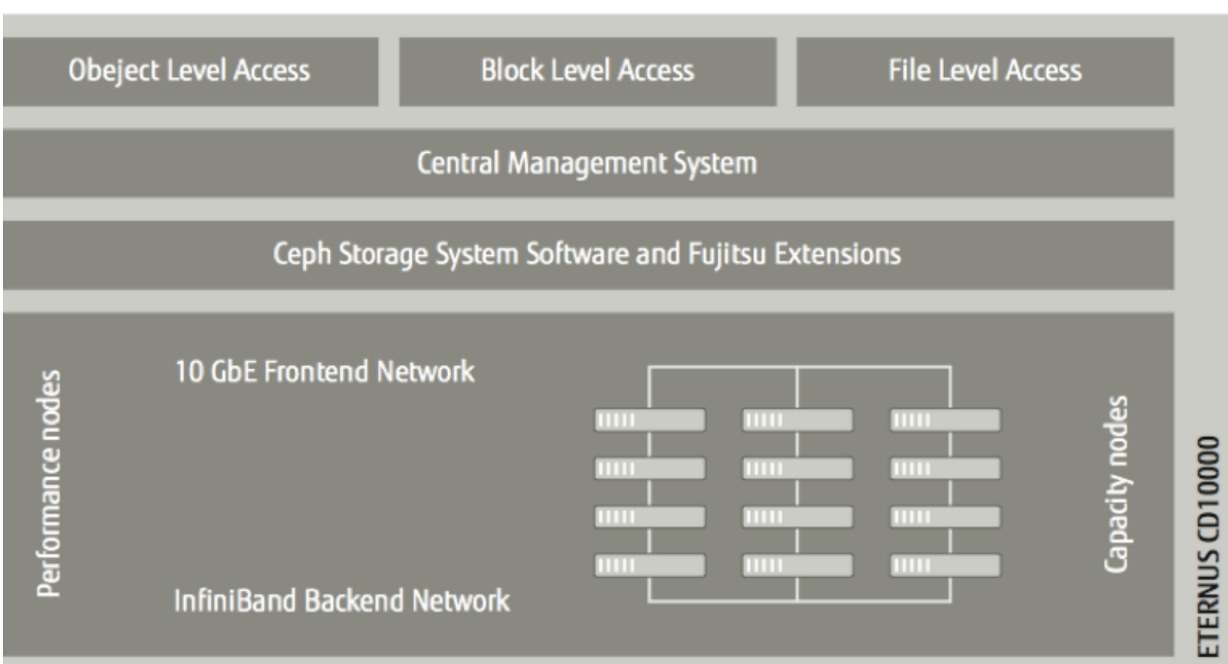

Figure 3 – Eternus CD10000 Architecture

Product Overview

ETERNUS CD10000 Hardware Architecture

Configuration schemes of ETERNUS CD10000 nodes:

This section describes the architecture of the ETERNUS CD10000.

Management Node

Using one Intel Xeon E5-2640 v2 processor, the management node is responsible for collecting all the data logs of the running system. The management node acts as a "housekeeper" and stores all the upcoming event logs from the system, but has no further action in system operation. For redundancy purposes, installing two management nodes per system is recommended.

Technical specs:

- 1x Intel® Xeon® E5-2640 v2 processor
- 64GB RAM
- Front-End Interface: 2x 1GbE onboard for administration, 4x 1GbE for administration and management
- 4x 2.5" 900GB SAS 10k HDDs (2x for OS, 2x for user data) The total raw capacity of the node is 3.6 TB SAS HDD.
- Delivers 3.6TB using 2.5" SAS disks (10k rpm) Management Node

Basic Storage Node

Using two Intel Xeon E5-2640 v2 processors, the storage node is responsible for storing and retrieving data on the system drive array. In scaling out the system, one would add basic storage nodes to achieve the required storage capability.

Technical specs:

- Dual Intel® Xeon® E5-2640 v2 processors
- 128GB RAM
- Node-Interconnect: Infiniband 2x 40GbIB (Mellanox IBHCR 40 GB dual channel QDR)
- Front-end interface: 10GbE (2x 10GbE PCIe x8 D2755 SFP+)
- Intel® P3700 800GB PCIe SSD for journal, metadata and caching
- 16x 2.5" 900GB SAS 10K RPM HDDs (2x for OS, 14x for user data)

The total raw capacity of each node is 12.6TB.

NOTE: The usable capacity depends on the number of the replicas.

Typically usable capacity = raw capacity / number of replicas, e.g. 2 or 3

Figure 4 – Basic Storage Node

Performance Node

Designed to deliver fast access to data, the performance node uses fast rotating SAS disks and PCIe connected SSDs.

- Intel® P3700 800GB PCIe SSD with 800GB capacity - also used as storage tier
- ETERNUS JX40 JBOD adds 21.6TB of 2.5" SAS disks (10K rpm)

The total raw capacity is 34.2 TB.

Typically usable capacity = raw capacity / number of replicas, e.g. 2 or 3

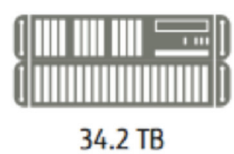

Figure 5 – Performance Node

Capacity Node

Designed to deliver high data density through a very compact housing for disks, the capacity nodes can host high capacity disks to store high data volumes in a cost effective manner.

- 14x SAS 900GB HDD plus 60x 3.5" NL-SAS 4TB HDD total

The total raw capacity is 252.6TB.

Typically usable capacity = raw capacity / number of replicas, e.g. 2 or 3

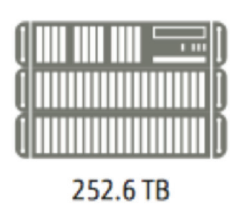

Figure 6 – Capacity Node

Note: The basic configuration of the ETERNUS CD10000 system is 4+1, which contains 1 management node and 4 basic nodes. The maximum amount of storage nodes is 200. The intra-node interconnect is based on an Infiniband 2x 40 GB 2 channel QDR. The various configuration details are shown in Figure 6.

Storage Node Architecture

The design of the storage nodes is optimized to provide consistent high performance. Even in failure situations no service degradation should occur.

This is achieved through:

- High performance Dual Intel Xeon E5-2640 v2 processors with a balanced I/O architecture
- 1 GB/s throughput including all redundant data copies (replicas)
- Very fast Infiniband dual channel back-end network for fast, low latency distribution of data between nodes and quick rebuild of data redundancy/redistributing data after hardware failure.
- PCIe Intel P3700 800GB SSDs for fast access of journal and metadata and also acting as a data cache. This helps to reduce or avoid latency in distributed storage architecture.

Figure 7 shows the data throughput in the ETERNUS CD10000 server nodes. This architecture ensures a balanced I/O performance with approximately 1 Gigabyte/s sustained speed including all redundant data copies per performance node. Every component runs at the required speed with a dual redundant network Infiniband switch in the backend.

ETERNUS CD10000 Software Architecture

Ceph Storage Software

Ceph is an open source, software defined storage platform that is designed to present object, block and file storage from a distributed Intel Architecture-based compute cluster. Ceph supports scalable clusters up to the Exabyte level. The data is replicated over disks and nodes and supports system fault tolerance; Ceph is designed to be both self-healing and self-managing. Ceph helps to reduce investment and operational costs. Ceph software is part of the OpenStack* project and the core storage element. OpenStack enables IT organizations to build private and public cloud environments with open source software.

One of the underlying Ceph technologies is the CRUSH algorithm that manages the homogeneous allocation of the data placement. The CRUSH algorithm (Controlled Replication Under Scalable Hashing) describes a method

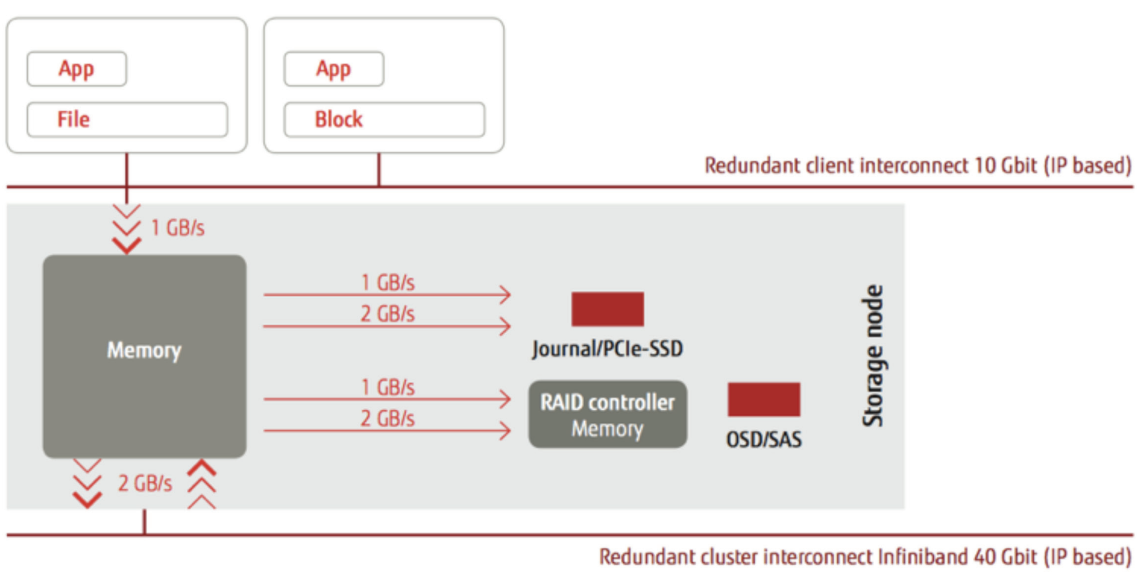

Figure 7 – Data Throughput

which determines how data is stored and retrieved by computing data storage locations. It works with a pseudo-random and uniform (weighted) distribution to establish a homogeneous allocation of the data between all available disks and nodes. Adding new devices (HDDs or nodes) will thus not have a negative impact on data mapping. The new devices become part of the system without any bottleneck or hot-spot effects. Due to the infrastructure algorithm powered by CRUSH, the placement based on physical infrastructure (e.g. devices, servers, cabinets, rows, DC's) becomes very easy. This enables high performance and provides the system with high resilience.

Figure 8 shows how the data is stored on object storage devices and distributed on various HDDs and server nodes. In order to simplify the visualization of the data distribution process, we use a sample object with 4x 4MB blocks. The 4MB blocks are stored in each object store container (HDDs) of the various server nodes. In this example a total of 3 copies (replicas) of the objects are stored and distributed on different containers providing data high availability trough redundancy. If the client asks for this 16MB file, the CRUSH algorithm collects the respective data blocks from the storage nodes and rebuilds it as an object file. Block sizes are not fixed and can be individually defined.

In the event of an unexpected hardware failure (disks or nodes), the system is able to seamlessly proceed without downtime. Since data copies are available on other locations within the system, data is redistributed over the remaining operational HDDs, with no significant negative influence on reliability or performance. The system automatically recreates new data copies (replicas) for those lost by the failure of storage components. Through the distributed nature of storing data, this happens in only a fraction of time when compared with the rebuild of data on a spare disk within a RAID array. Ceph storage software overcomes the central access bottleneck with an algorithmically determined method of storing and retrieving data, and therefore avoids a single point of failure. It also avoids performance bottlenecks through extensive parallelism. Theoretically it scales to the Exabyte level. CRUSH uses a map of the storage node cluster to pseudo-randomly store and retrieve data in OSDs across the complete system. It also uses intelligent data replication to ensure resiliency.

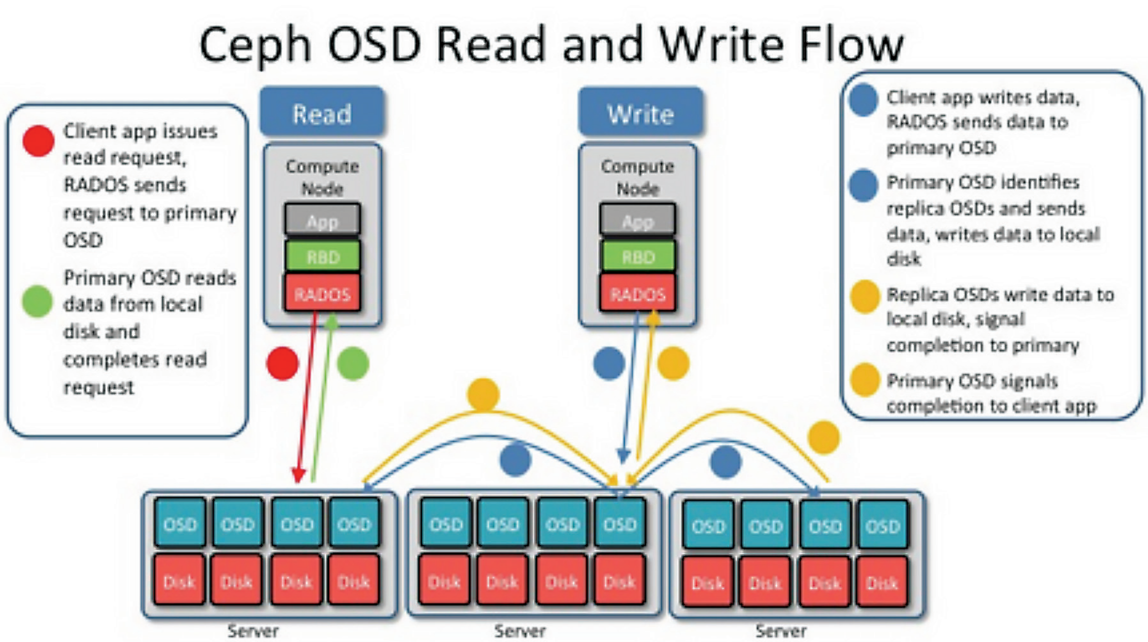

Figure 8 – How Ceph Stores Dataput

ETERNUS CD10000 Management System

Figure 9 shows how the ETERNUS CD10000 software architecture and the clients access the object store devices (OSDs). Monitoring Clusters (MONs), Object Storage Devices (OSDs) and the Metadata Servers (MDSs) are part of the software storage solution of Ceph. The Cluster Monitors (MONs) are responsible for the authentication of the clients. MONs also controls the cluster membership, the cluster state and the cluster map of accessed data. The Meta-Data Server (MDS) controls the block file objects of the POSIX compliant metadata and support rules for the namespace management. The Object Storage Device (OSD) stores and organizes all the data and metadata on the object stores in flexibly sized containers; up to 10,000 containers may be enabled.

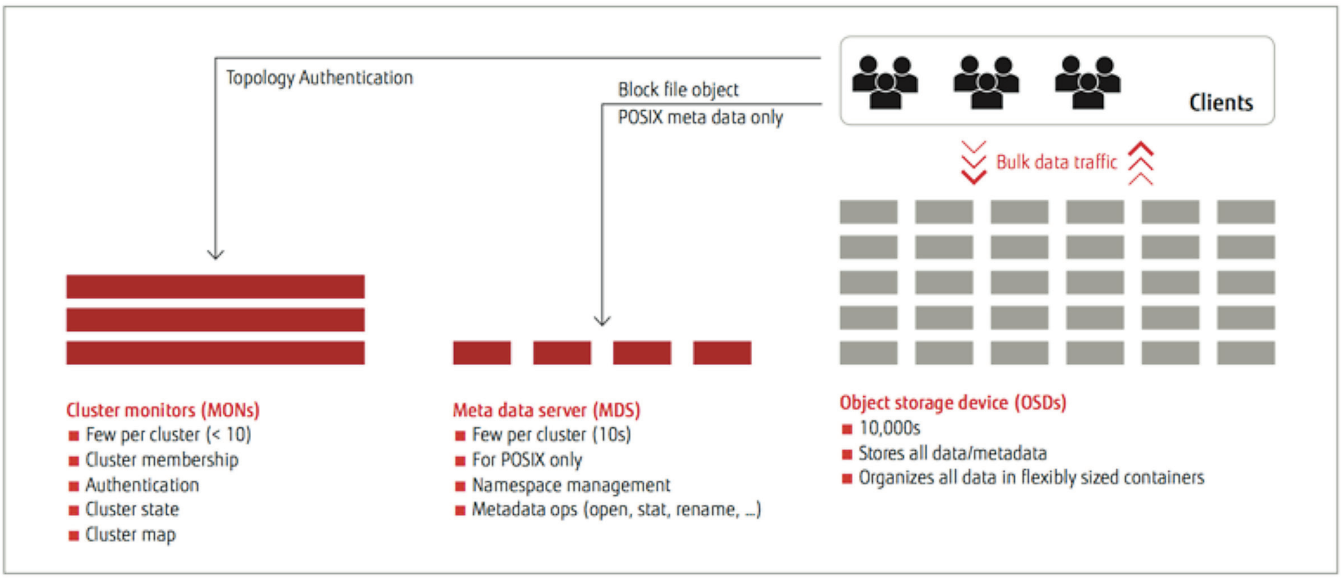

Figure 9 – ETERNUS CD10000 Software Architecture

Virtual Storage Manager for Ceph

The basic system management of Ceph offers a generic command line interface that requires suitable skills in batch files and Linux support. Fujitsu has complemented the basic management functions of Ceph with a GUI based deployment of the Intel developed, open source Virtual Storage Manager (VSM), plus additional Fujitsu enhancements for supporting operational speed and efficiency.

VSM is a browser-based management application for Ceph storage systems. It supports Ceph clusters comprising a mix of hard disk drives (HDDs), solid state drives (SSDs), and SSD-cached HDDs.

VSM provides a unified, consistent view of a Ceph storage system. It provides a framework for classifying and organizing storage system assets, and simplifies the initialization, configuration, monitoring, and management of the Ceph storage system.

The following diagram provides a high-level logical view of a VSM-managed Ceph cluster in a cloud datacenter application.

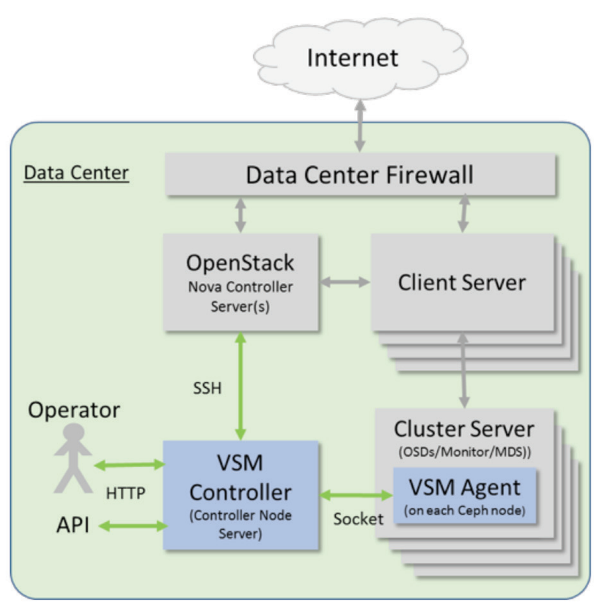

Figure 10 – VSM High Level Architecture

The VSM Storage Controller typically runs on a dedicated server or server instance. The controller maintains the system state and manages storage system configuration, including configuration of Ceph data placement rules. The controller also provides a browser-based graphical interface that supports operations management. The controller interfaces to OpenStack for managing Ceph storage pools under OpenStack Cinder*, system monitoring of cluster health, and capacity utilization.

VSM agents run on each storage node. The VSM agent retrieves configuration information from the server, and monitors performance metrics

The key elements of a typical VSM-managed Ceph cluster in a cloud datacenter application are shown in Figure 11.

This in the CD1000, VSM delivers a single pane of:

- Central software deployment
- Central network management
- Central log file management
- Central cluster management
- Central configuration, administration and maintenance
- SNMP integration of all modes and network components. The GUI dashboard enables easy hassle-free administration of the system with simultaneous time and cost savings for complex storage administration and management procedures.

The following diagrams show you a few of the primary VSM displays.

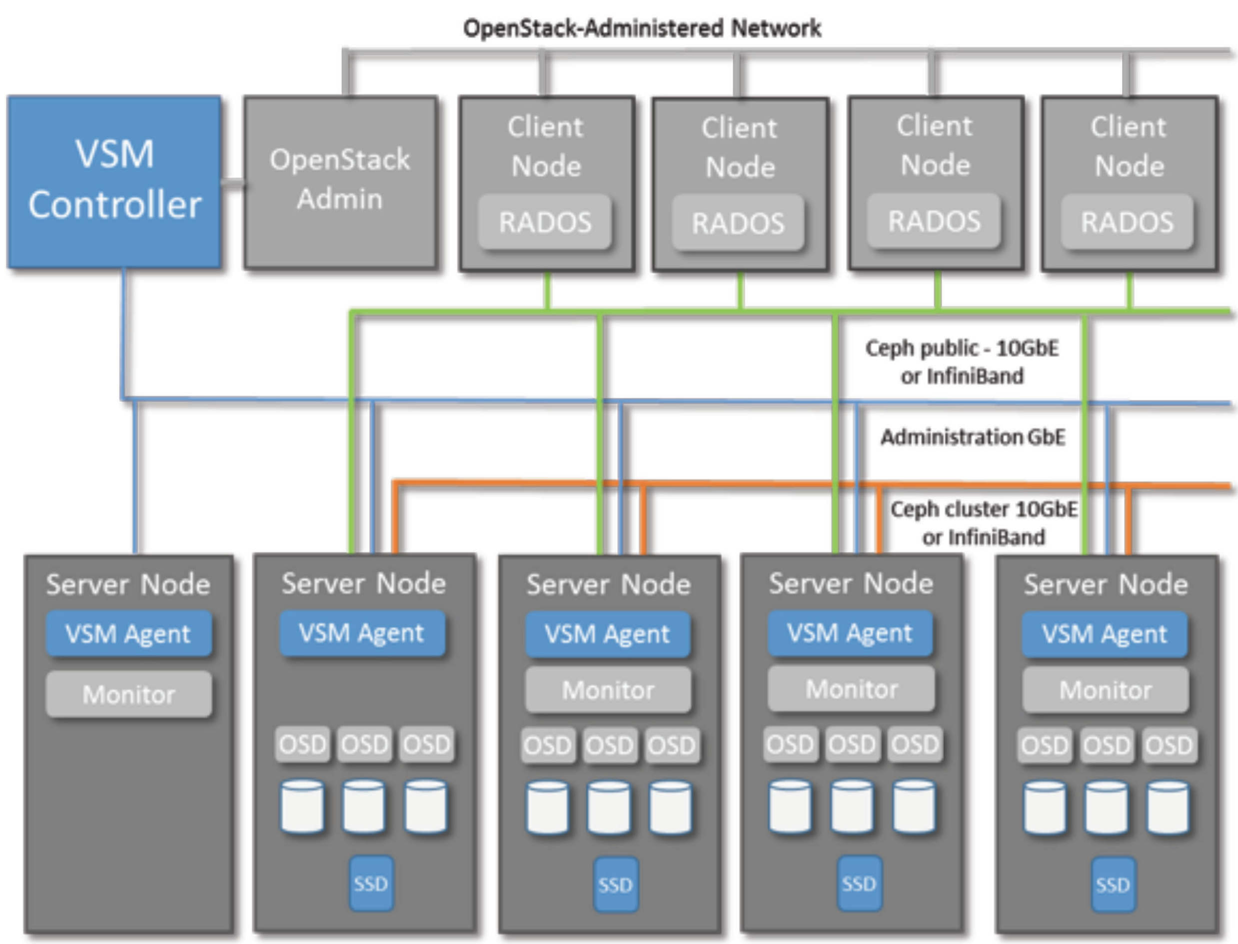

Figure 11 – How VSM Manages Ceph Storage

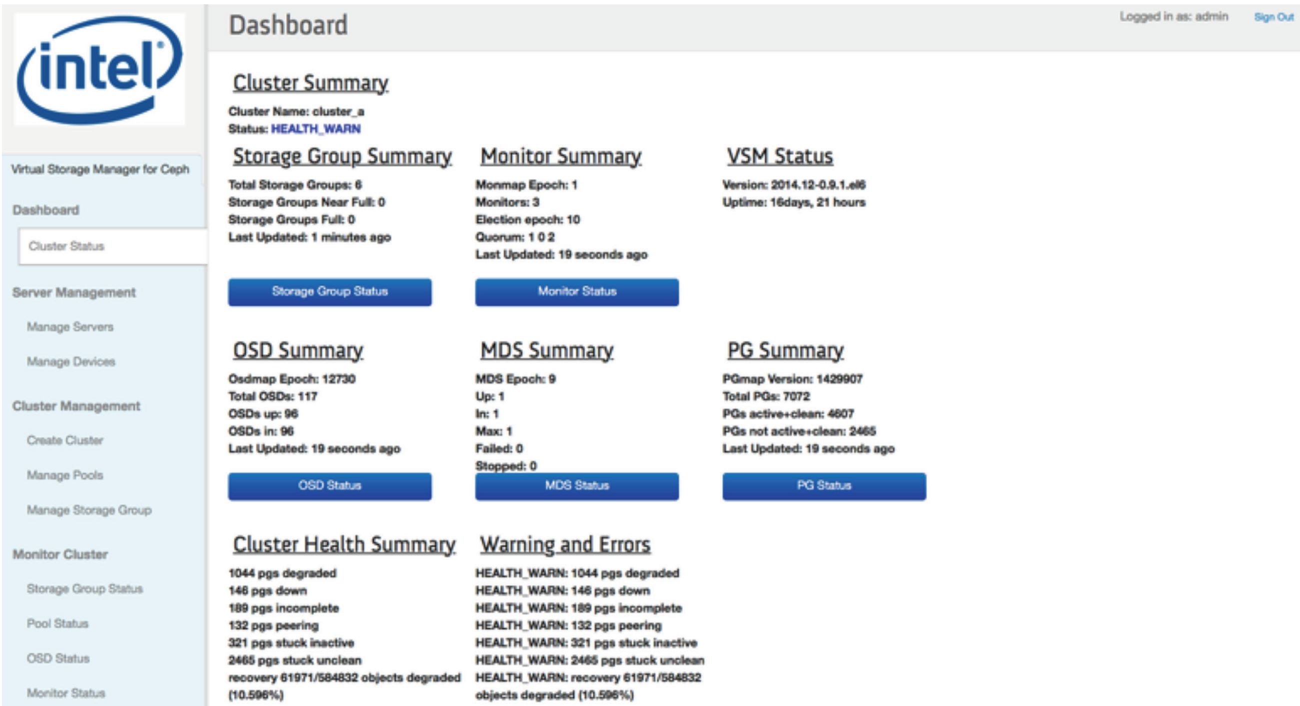

Figure 12 – Main VSM Dashboard

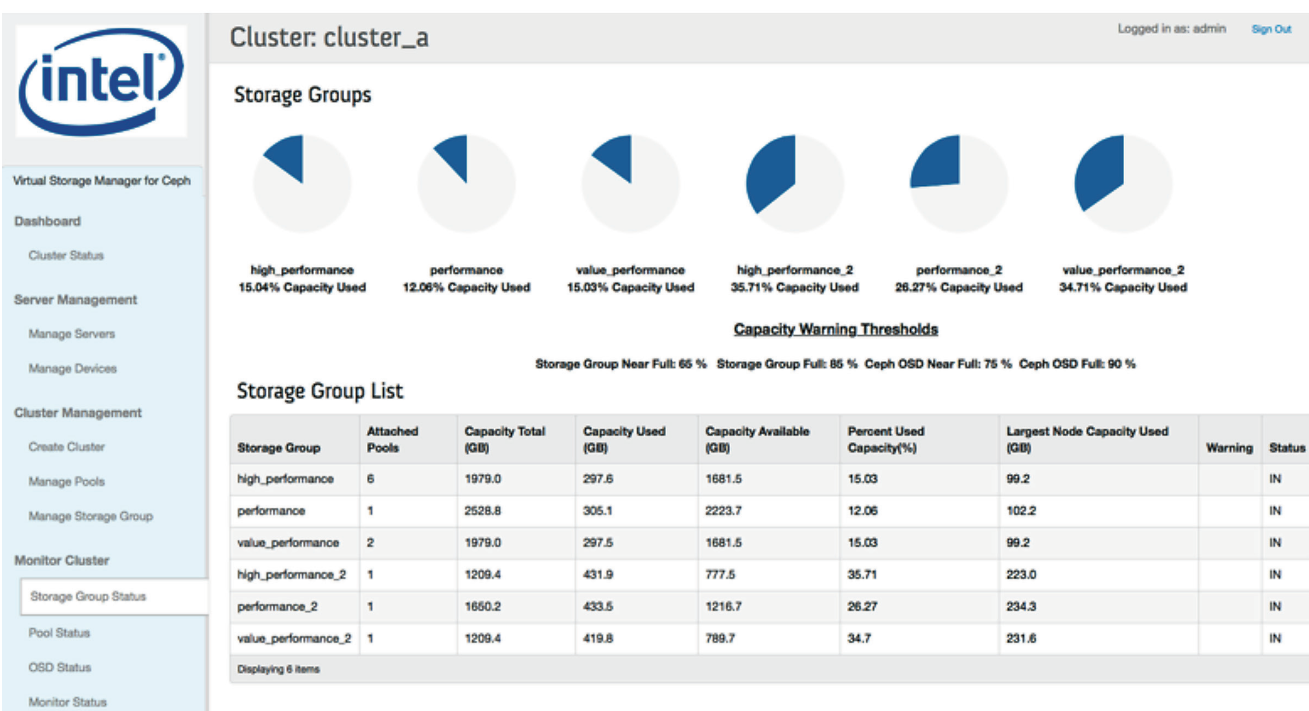

Storage Group	Attached Pools	Capacity Total (GB)	Capacity Used (GB)	Capacity Available (GB)	Percent Used Capacity(%)	Largest Node Capacity Used (GB)	Warning	Status
high_performance	6	1979.0	297.6	1681.5	15.03	99.2		IN
performance	1	2528.8	305.1	2223.7	12.06	102.2		IN
value_performance	2	1979.0	297.5	1681.5	15.03	99.2		IN
high_performance_2	1	1209.4	431.9	777.5	35.71	223.0		IN
performance_2	1	1650.2	433.5	1216.7	26.27	234.3		IN
value_performance_2	1	1209.4	419.8	789.7	34.7	231.6		IN

Figure 13 – VSM Storage Pools Tracking

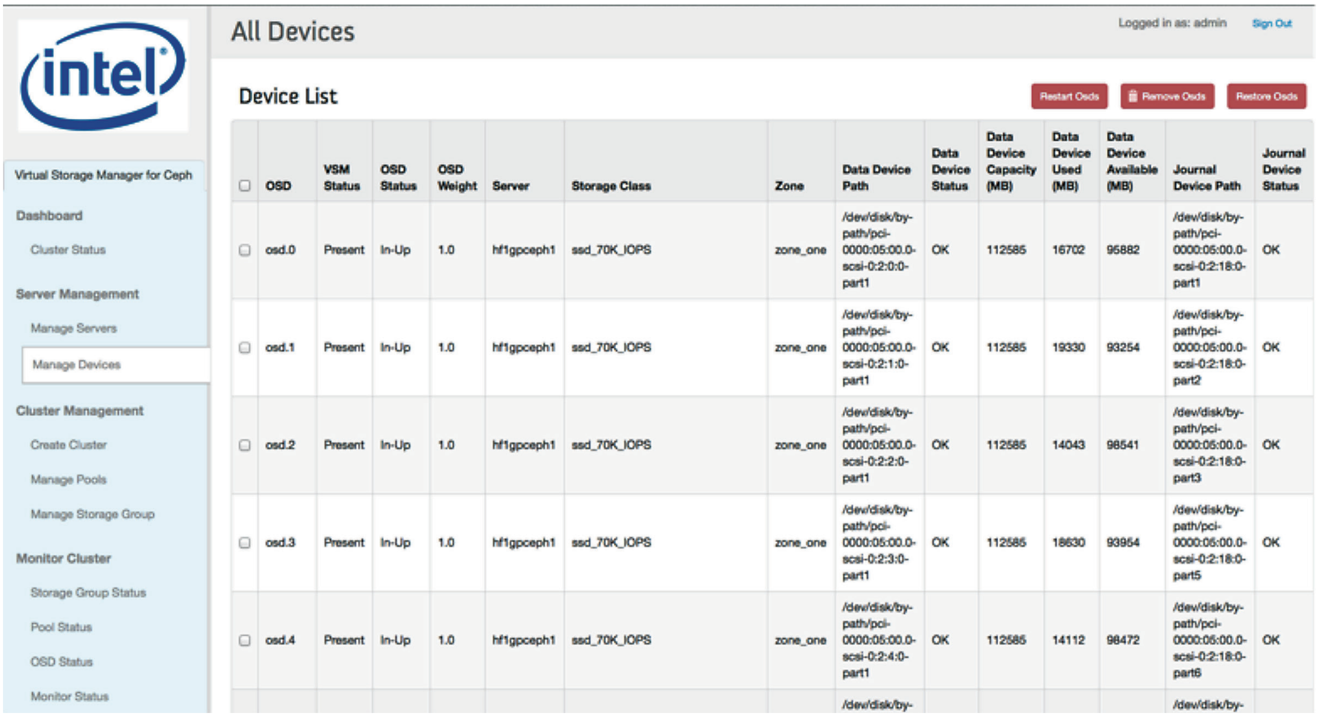

OSD	VSM Status	OSD Status	OSD Weight	Server	Storage Class	Zone	Data Device Path	Data Device Status	Data Device Capacity (MB)	Data Device Used (MB)	Data Device Available (MB)	Journal Device Path	Journal Device Status
osd.0	Present	In-Up	1.0	hf1gpceph1	ssd_70K_IOPS	zone_one	/dev/disk/by-path/pci-0000:05:00.0-scsi-0:2:0:0-part1	OK	112585	16702	95882	/dev/disk/by-path/pci-0000:05:00.0-scsi-0:2:18:0-part1	OK
osd.1	Present	In-Up	1.0	hf1gpceph1	ssd_70K_IOPS	zone_one	/dev/disk/by-path/pci-0000:05:00.0-scsi-0:2:1:0-part1	OK	112585	19330	93254	/dev/disk/by-path/pci-0000:05:00.0-scsi-0:2:18:0-part2	OK
osd.2	Present	In-Up	1.0	hf1gpceph1	ssd_70K_IOPS	zone_one	/dev/disk/by-path/pci-0000:05:00.0-scsi-0:2:2:0-part1	OK	112585	14043	98541	/dev/disk/by-path/pci-0000:05:00.0-scsi-0:2:18:0-part3	OK
osd.3	Present	In-Up	1.0	hf1gpceph1	ssd_70K_IOPS	zone_one	/dev/disk/by-path/pci-0000:05:00.0-scsi-0:2:3:0-part1	OK	112585	18630	93954	/dev/disk/by-path/pci-0000:05:00.0-scsi-0:2:18:0-part5	OK
osd.4	Present	In-Up	1.0	hf1gpceph1	ssd_70K_IOPS	zone_one	/dev/disk/by-path/pci-0000:05:00.0-scsi-0:2:4:0-part1	OK	112585	14112	98472	/dev/disk/by-path/pci-0000:05:00.0-scsi-0:2:18:0-part6	OK
							/dev/disk/by-					/dev/disk/by-	

Figure 14 – Storage Devices Connected to OSDs and Their Status

ETERNUS CD10000 Unified Storage

Ceph Interface Architecture

Figure 15 shows the interface for the unified object, block and file level access. The software core is the Ceph Object Storage (RADOS). RADOS stands for "Reliable Autonomic and Distributed Object Store". RADOS contains the interfaces to access Ceph cluster hosted data for: Librados, Ceph Object Gateway, Ceph Block Device and Ceph File System services.

Librados: The software libraries provide client applications with direct access to the (RADOS) object-based storage system, and also provide a foundation for some of Ceph's features. The Librados interface libraries provide for apps directly access to RADOS with support for C, C++, Java, Python, Ruby and PHP.

Ceph Object Gateway (RGW): The Ceph object gateway offers direct object storage access via the well known Amazon* S3 (Simple Storage Service) or OpenStack Swift API's.

Ceph Block Device (RBD): The Ceph block device is offering block storage for e.g. QEMU/KVM virtual machines, virtual machine volumes as well as virtual machines snapshots and images if configured to do so. Block Devices can be mounted on all OS's as they mounted via the hypervisor with no client drivers required.

By design, Ceph automatically replicates data across the cluster and reacts automatically to node or drive failures and is in fact self-healing.

Ceph File System (CephFS): The Ceph file system allows direct file and directory access for applications. Clients mount the POSIX-compatible file systems e.g. using a Linux kernel client.

Customer values – Fujitsu blueprint scenarios: In addition to the generic interfaces, Fujitsu offers service concepts to use the system on standardized scenarios, e.g. OpenStack, Drop Box*-like file-sharing, and Backup/Archiving. These types of customized blueprint scenarios are not part of the system, but regarding the various customer capabilities, this can be individually considered and implemented.

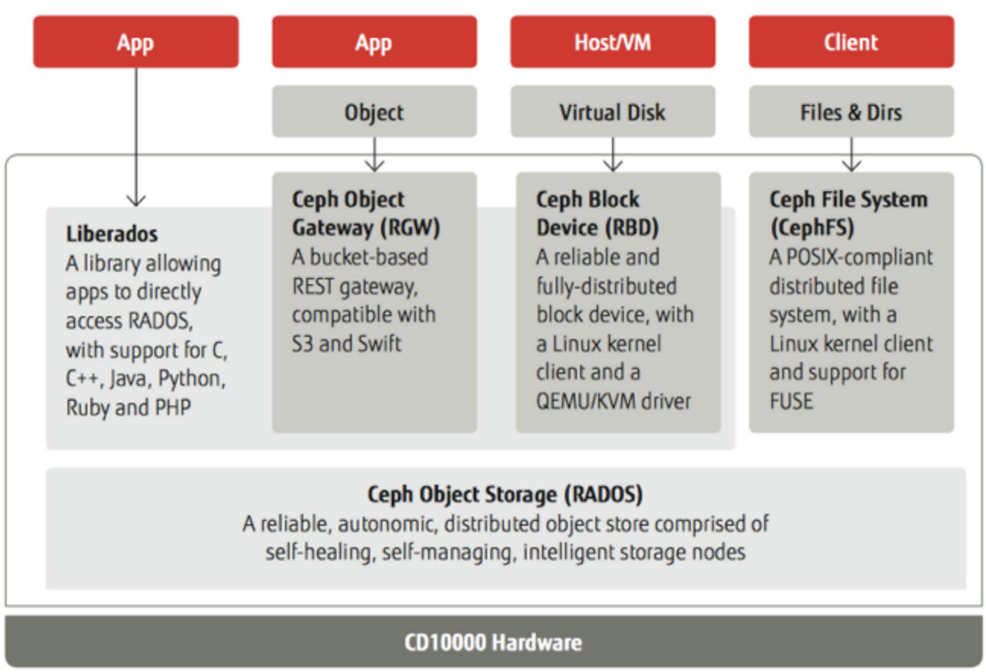

Figure 15 – Interfaces for Unified Object, Block and File Level Access

One Typical Use Case for Ceph – Hosting VM Images

The OpenStack documentation notes that Ceph is designed to expose different types of storage interfaces to the end user: it supports object storage, block storage, and file-system interfaces. Ceph supports the same API as Swift for object storage, can be used as a backend for Cinder block storage as well as backend storage for Glance* images. Ceph supports "thin provisioning", implemented using copy-on-write methodology. When Openstack Glance, Cinder, Nova* and Libvirt/KVM* are all integrated together to use Ceph storage, live instant volume snapshots, live migration of virtual machines between hosts as well as more advanced provisioning use cases can be implemented.

This can be useful when booting from volume because a new volume can be provisioned very quickly. Ceph also supports Keystone*-based authentication and integration, so native authentication and authorization of both actual users and service users can be offloaded to this server with minimal configuration changes in Ceph.

If you want to manage your object and block storage within a single system, or if you want to support fast boot-from-volume, Ceph is a good choice. Figure 16 shows the key components OpenStack. The arrows represent a logical flow and not data movement.

Glance is an OpenStack image repository for virtual machine images. In a large-scale implementation, you need a place to maintain all virtual machine templates for new VMs as well as VM backups. In OpenStack this is done by Glance.

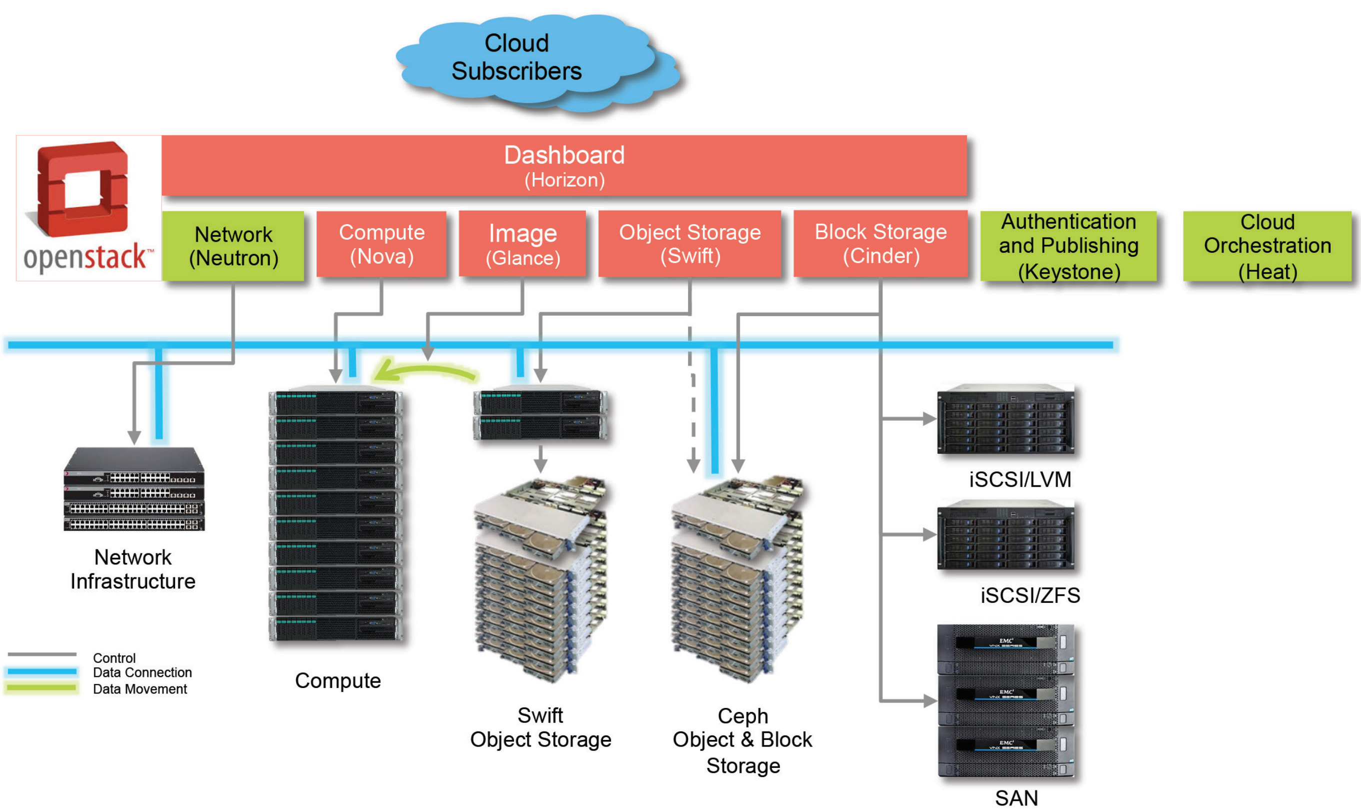

Figure 16 – OpenStack and its Primary Components

Adding a VM image to Glance is straightforward. Here is an example of adding a RHEL image:

```
# name=RHEL
# image=RHEL _ 6.2-disk.img
# glance image-create --name=$name --is-
public=true --container-format=bare
--disk-format=qcow2 < $image
```

Note that images can also be uploaded via the Openstack UI at either per tenancy or fully public available images for all users of the platform. Images can also be downloaded via the UI if needed.

Nova manages the OpenStack compute resources for virtual machines. You can create, start/stop/delete virtual machines. Similar to Cinder, Nova uses existing virtualization technologies (e.g. KVM, QEMU*, Xen*, VMware*) to do this. We can treat this as virtual machine management framework. With Nova, you can start a virtual machine using a Glance image template. Glance (similar to Cinder) provides a variety of backends for image.

Using Nova, you can list the image we just added in the previous step with glance:

```
$ nova image-list
+--------------------------------------+-----------+--------+--------+
|                  ID                  |    Name   | Status | Server |
+--------------------------------------+-----------+--------+--------+
| 17a34b8e-c573-48d6-920c-b4b450172b41 | RHEL 6.2  | ACTIVE |        |
+--------------------------------------+-----------+--------+--------+
```

And then using Nova we can boot a VM using this image:

```
# nova boot --flavor 2 --key _ name oskey --image 17a34b8e-c573-48d6-920c-b4b450172b41 rhel
```

The Nova list command allows you to watch the image startup and then connect to your new image

```
 $ nova list
+--------------------------------------+--------+--------+------------------+
|                  ID                  |  Name  | Status |     Networks     |
+--------------------------------------+--------+--------+------------------+
| 0e4011a4-3128-4674-ab16-dd1b7ecc126e | rhel   | BUILD  | demonet=10.0.0.2 |
+--------------------------------------+--------+--------+------------------+
$ nova list
+--------------------------------------+--------+--------+------------------+
|                  ID                  |  Name  | Status |     Networks     |
+--------------------------------------+--------+--------+------------------+
| 0e4011a4-3128-4674-ab16-dd1b7ecc126e | rhel   | ACTIVE | demonet=10.0.0.2 |
+--------------------------------------+--------+--------+------------------+
$ ssh -i oskey.priv root@10.0.0.2
```

Or you can also use the Openstack management UI to launch/remote console access/start/stop/destroy/add or delete volumes/ add or delete floating public IP addresses and rules as needed.

OpenStack Ceph is the storage for this new VM image. Ceph is well suited for this due to its support of block storage, and interface with Cinder and OpenStack. This happens because we previously configured Glance to use Ceph to store its images in a Ceph storage pool, and we enabled copy-on-write cloning of images. Additionally, we told Cinder (the OpenStack block device module) what Ceph pool to use, and the location of the Ceph configuration file, added secrets to Libvirt/KVM – and setup authentication with Keystone as well as the Glance/Nova/Cinder/Libvirt service clients.

Figure 17 shows the process of creating a new VM for a user, using OpenStack and Ceph.

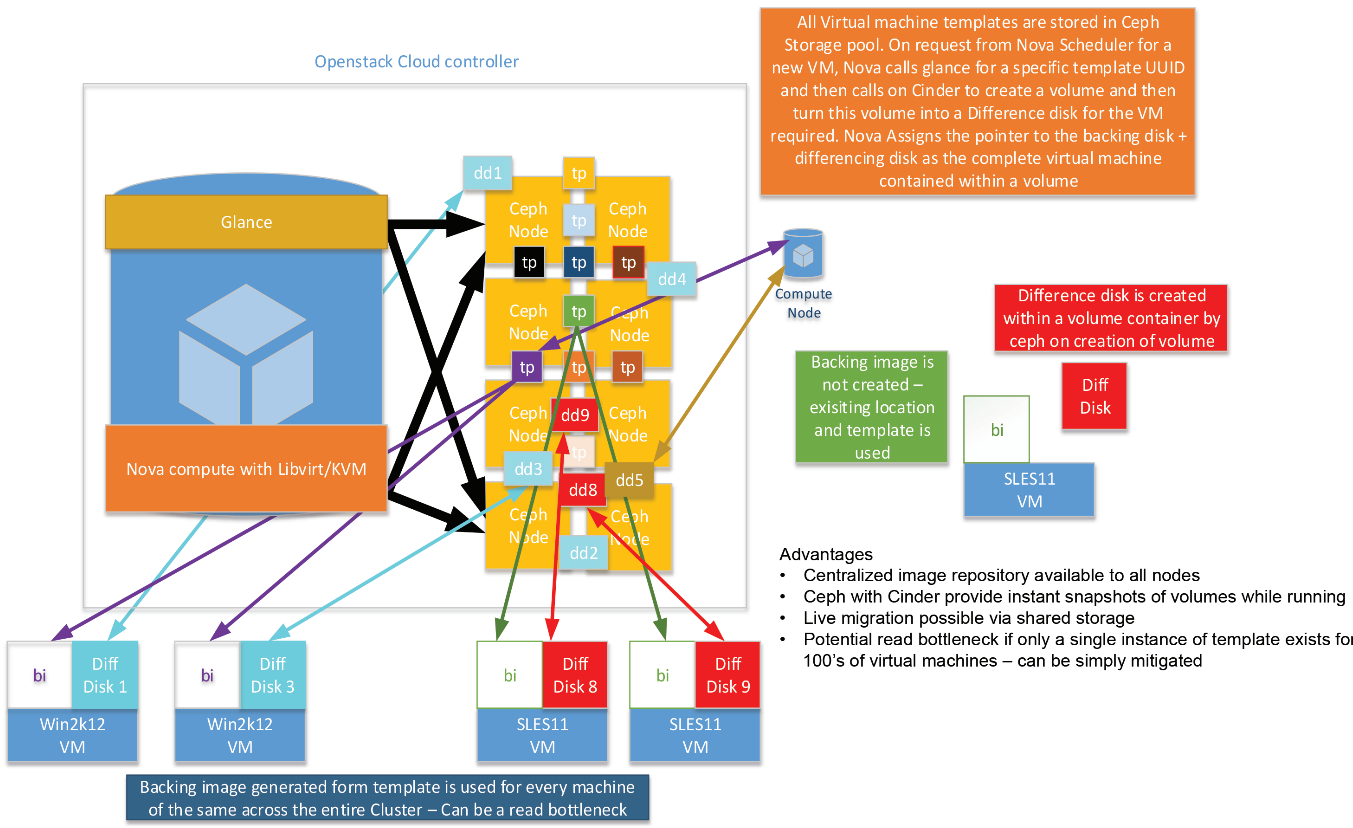

Figure 17 - VM Deployment with Ceph integrated Cinder-glance-libvirt

For more detailed information regarding operating OpenStack platform software, please refer to the operations documents linked here:

Cloud administration: http://docs.openstack.org/admin-guide-cloud/content/

Operations guide: http://docs.openstack.org/ops/

Virtual machine imaging guide: http://docs.openstack.org/image-guide/content/

Key Considerations for a System Platform Approach

The ETERNUS CD10000 is delivered as a complete platform. Based on the required customer configurations, all hardware and software components are delivered as a system level product. Fujitsu's development teams ensure that the internal server, storage, networking and software functions work seamlessly together. This also includes Ceph open source software. Benchmark and tuning tests are performed for the right sizing of all components, which is key to running a petabyte scale storage system without performance bottlenecks and availability issues.

ETERNUS CD10000 maintenance services provide support for eliminating potential technical errors of all components. Consistent upgrades are delivered to keep the overall system up-to-date whilst ensuring the interoperability of all components. This enables users to benefit from future enhancement of the Ceph software without comprising the stability of the system, reducing customer resource requirements by elimination of the need to “keep up” with various software updates and code drops.

For a distributed storage system which is software based, lifecycle management is paramount as it allows in principle a mix of generations of hardware nodes as the core functionality lies with the software. Through the inherent data distribution mechanism, hardware refreshes can be conducted completely online with automated data migration. Thus a very long system lifetime can be achieved by helping to reduce big migration projects, which are not trivial in Petabyte dimensions. Nevertheless it is all about compatibility and right sizing of hardware and software, which is a key advantage of the ETERNUS CD10000 storage system. Users can fully benefit from open source software and receive professional enterprise-class support, while they are released and insulated from complex evaluation, integration and maintenance tasks.

In Conclusion

Intel’s Xeon E5-2600 v2 processor family and Intel 1GbE and 10GbE networking devices form the hardware “backbone” of the Eternus CD10000 storage solution, that offer highly scalable and reliable IO for block, object, and file use cases including VM image and applications hosting. Using the Intel P3700 PCIe SSDs further enhances the performance of the platform. Specific customer use cases and testing would be required to determine maximum performance capabilities under varying use cases and storage array sizes.

For Ceph benchmarking and performance data on Intel Xeon processors with Intel P3700 PCIe SSDs and Intel networking adapters, Intel performance data is expected to be available as a follow-on to this Solution Reference Architecture with use cases clearly identified.

Disclaimers