

NexentaStor on Cisco UCS C-Series Rack Mount Servers: Storage Platform for Desktop Virtualization

What You Will Learn

This document provides guidance on implementing the Cisco® and Nexenta solution to optimize storage space to support virtual desktop implementations. This document also discusses a service provider case study based on the work performed for Korea Telecom (KT). This document assumes that the reader is familiar with desktop virtualization and implementation concepts in general.

Introduction

The rapid increase in data in the data center is causing dramatic growth in centralized data stores. Server, networking, and storage infrastructures are evolving to address many of the challenges associated with this growth. An increasing trend in the data center is the use of standard x86-based servers with internal storage, which then are presented as the storage layer.

The second generation of Cisco UCS™ C-Series Rack-Mount Servers supports up to 16 internal disk drives in a small form factor (SFF). They provide a robust platform for both enterprise applications and storage server platforms. Cisco UCS C-Series servers, deployed with NexentaStor software from Nexenta Systems, provide a flexible, network-based storage environment that is optimized for various workloads, such as virtual desktops, and cloud-based applications.

Challenges of Virtual Desktop Infrastructure Storage Workloads

Desktop users can present challenges for IT departments. Users may request deployment of unique applications on their desktop that run nonstandard and preproduction applications beyond standard OS and application patch levels. Virtual desktop infrastructure (VDI) offers the promise of central administrative control over user desktops. With VDI, IT now can apply best practices to application image management and system management while still meeting the unique demands of individual users. Leading VDI solutions, such as Citrix XenDesktop and VMware View, are examples of this technology. However, one of the challenges of VDI solutions has been the inability of traditional storage systems to handle the challenging I/O requirements for this workload cost effectively.

From a storage I/O perspective, VDI workloads produce a higher percentage of small, random write operations, which can overwhelm traditional storage solutions, but which server-based direct-attach storage (DAS) solutions can handle more economically than traditional SAN-based storage arrays that depend on a large number of short-stroked disks.

NexentaStor is a storage software product from Nexenta Systems. Based on the innovative ZFS architecture, NexentaStor is designed to address the limitations of traditional storage systems. NexentaStor, running on Cisco Unified Computing System™ (Cisco UCS) servers, offers a high-performance and cost-effective solution for the most demanding VDI applications.

Cisco UCS Overview

Cisco UCS C-Series Rack-Mount Servers reduce total cost of ownership (TCO) and increase business agility by extending unified computing innovations and benefits to rack-mount servers and offering a wide range of options that can be balanced for various application workloads.

The Cisco UCS C210 M2 General-Purpose Rack-Mount Server and C260 M2 Rack-Mount Server are ideal platforms for storage nodes. In addition to the flexibility of the platform and configuration options, Cisco UCS C-Series servers have unique benefits as part of the Cisco UCS infrastructure:

- Cisco Integrated Management Controller (IMC) integrates server manageability that provides Simple Network Management Protocol (SNMP)-based monitoring, XML API-based BIOS configuration and server management, and internal storage management through OS-agnostic means.
- Cisco UCS Manager provides unified management through the Cisco UCS 6100 and 6200 Series Fabric Interconnects and Cisco Nexus® 2200 fabric extender platform.
- Scale-out interconnect topology is provided through the Cisco Nexus 5500 switch platform and Cisco Nexus 2200 fabric extender platform.

The Cisco UCS C210 M2 server is a general-purpose, 2-socket, 2-rack unit (2RU) rack-mount server that balances performance, density, and efficiency for storage-intensive operations. Some of the main features are:

- Up to two Intel Xeon 5500 or 5600 series multicore processors
- Up to 192 GB of industry-standard, double-data-rate-3 (DDR3) main memory
- Up to 16 internal SFF, SAS, or SATA disk drives; up to 16 terabytes (TB) total
- RAID 0, 1, 5, 6, 10, 50, and 60 support for up to 16 SAS, SATA, or SSDs with up to two optional LSI MegaRAID controllers
- Five full-height PCI Express (PCIe) slots: two full-height, full-length x8 PCIe card slots and three full-height, half-length x8 PCI card slots, all with x16 connectors

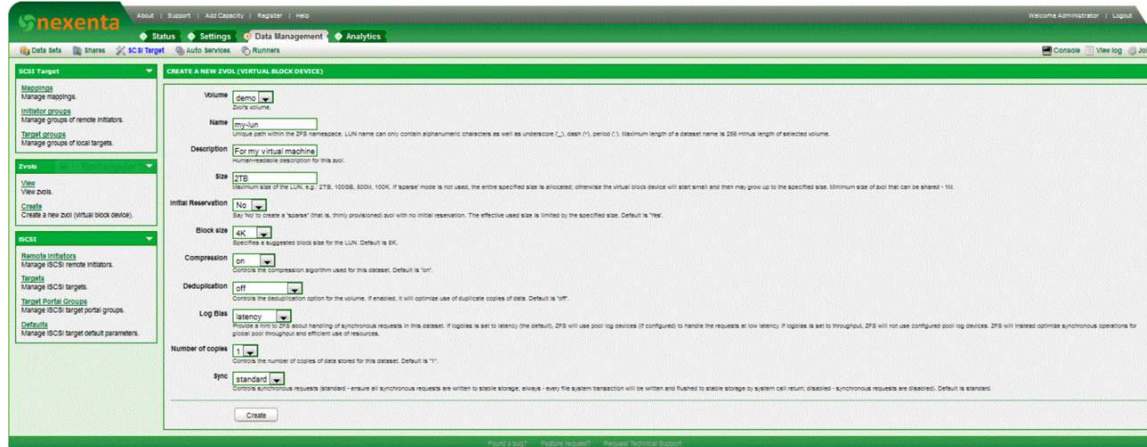
The Cisco UCS C260 M2 Rack-Mount Server is one of the industry's highest-density 2-socket rack-server platforms. This server platform offers compact performance for enterprise-critical applications within this architecture as well as CPU enhancements for greater performance, expandability, and security, plus reliability, availability, and serviceability (RAS) features. Some of the other main features are:

- Large memory capacity and 16 drives, making it an ideal platform for memory-bound or disk-intensive applications
- Exceptional memory, with up to 64 DIMM slots and up to 1 TB, based on DDR3 technology
- Intel Xeon processor E7-2800 product family
- 16 internal SFF, SAS, SATA-II, and SSD drives; up to 16 TB total
- RAID 0, 1, 5, 6, 50, and 60 support for up to 16 SAS, SATA, or SSDs with up to two optional LSI MegaRAID controllers
- Support for up to six PCIe cards in multiple low-profile and standard high slots

NexentaStor Overview

NexentaStor is a unified storage offering file and block access. File support includes Common Internet File System (CIFS), Network File System (NFS), FTP, rsync, and WebDAV. Block support includes Small Computer System Interface over IP (iSCSI) and Fibre Channel. NexentaStor integrates volume management capabilities and presents a single shared storage pool from which data sets can be configured to consume the pooled space as needed. The 128-bit file system architecture helps enable nearly unlimited scalability. Users are no longer limited by pool sizes, volume sizes, or file system sizes, and capacity can be adjusted on demand (Figure 1).

Figure 1: Nexenta's Data Management View



NexentaStor offers a number of features to make the best use of storage space and deliver efficiency:

- **Compression:** Compression is built in to provide capacity savings.
- **Hypervisor integration:** In VMware ESX, when virtual machines are destroyed, the associated storage space is released. In addition, VMware vStorage APIs for Array Integration (VAAI) improves replication performance.
- **Inline deduplication:** This feature is especially beneficial for virtualized workloads.
- **Thin provisioning:** Thin provisioning provides space efficiencies. Dynamic and nondisruptive physical space allocation are supported.
- **Scalable 128-bit architecture:** The architecture supports 2^{64} snapshots and provides search capabilities to help manage the snapshots.
- **Management:** Nexenta VM DataCenter (VMDC) offers simplified management in a virtualized environment with hypervisor visibility and context. Manage also can be performed through a web GUI.

NexentaStor for VDI

With VDI, the user's desktop experience is highly dependent on I/O latency. However, VDI introduces severe I/O challenges for traditional storage systems:

- 90 percent or more of VDI I/O consists of write operations.
- The I/O traffic has sporadic bursts as the virtual desktops page to storage.
- The I/O bursts from paging activity are small and random.
- Users may have similar usage patterns, such as logging in at the same times in the morning. These log-in storms can put tremendous pressure on the storage system.
- I/O alignment problems can affect storage performance.

Critical features of NexentaStor are amplified in the virtual desktop environment:

- ZFS performance for small random write operations is improved through the allocate-on-flush architecture.
- Variable block sizes are supported to allow performance to be optimized by matching application record sizes. Aligning the block size with the smallest record size typically used by the application avoids the performance penalty that results from reading the blocks, modifying the blocks with the changed data, and then writing the blocks back to disk.
- NexentaStor provides inline compression. Since compression is performed before the data is written, I/O is reduced if the data is compressible. For the desktop operating system virtual disk, the compression ratio is 1.5.
- The hybrid storage pool architecture of ZFS is critical to achieving optimal price-performance ratios for VDI. ZFS can target solid-state devices to be used for specific purposes within a storage pool. For example, a device can be targeted for use as a read cache. This capability allows you to achieve the low-latency benefits of solid-state devices without requiring the creation of a whole pool of the devices. Use of SSDs for read caches can be particularly effective for managing the boot storms inherent in virtual desktop deployments. ZFS intent logs can buffer write operations to reduce latencies for protocols using synchronous transactions, such as NFS.

Collectively, these built-in NexentaStor features can improve overall VDI performance significantly.

Storage Node Considerations

Cisco UCS C210 Storage Nodes

The Cisco UCS C210 storage nodes are equipped with two Intel Xeon X5500 or X5600 series processors, 48 GB of memory (up to 192 GB is supported), one dual-port 10 Gigabit Ethernet Cisco UCS P81E Virtual Interface Card (VIC), and 16 internal SFF disk drives. The performance-optimized configurations consist of 600-GB 10,000-RPM SAS disk drives, and capacity-optimized configurations consist of 1-TB 17,200-RPM SATA disk drives. In addition, a combination of performance- and capacity-optimized configurations on the storage node is supported.

Cisco UCS C260 Storage Nodes

The Cisco UCS C260 storage nodes are equipped with two Intel Xeon E7 processors, 256 GB of memory (upto 1 TB is supported), one dual-port 10 Gigabit Ethernet Cisco UCS P81E VIC, and 16 internal SFF disk drives. The extreme performance-optimized configuration consists of 16 300-GB SSD disk drives, the performance-optimized configuration consists of 16 600-GB 10,000-RPM SAS disk drives, and the capacity-optimized configuration consists of 16 1-TB 17,200-RPM SATA disk drives. In addition, a combination of configurations optimized for extreme performance, performance, and capacity on the storage node is supported.

The deployment topologies based on Cisco UCS 6100 or 6200 Series Fabric Interconnects and the Cisco Nexus 5500 switch platform are shown in Figures 2, 3, and 4.¹

¹ NexentaStor on UCS with file based storage was demonstrated, not block based storage. At the time of writing this document, NexentaStor block based interoperability with UCS was not completed.

Figure 2: Deployment Topology Using Cisco UCS 6100 or 6200 Series Fabric Interconnects

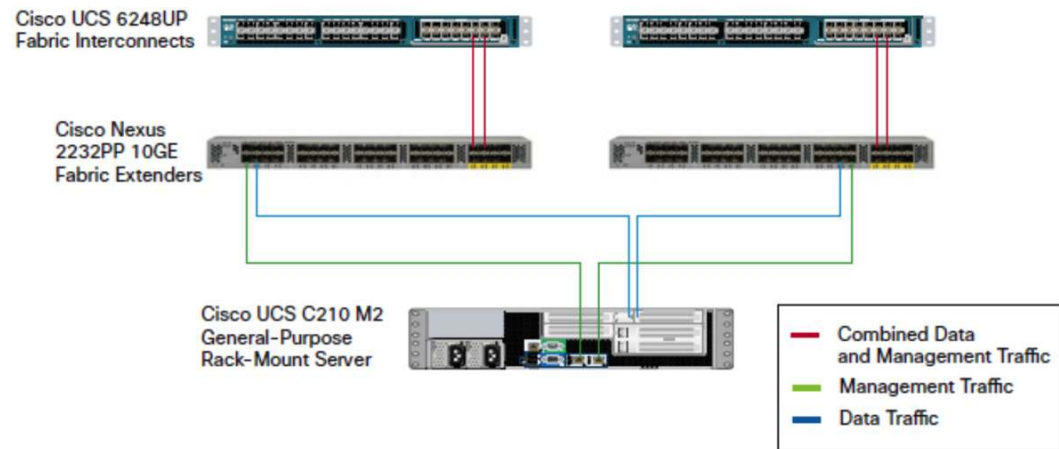


Figure 3: Deployment Topology Using Cisco Nexus 2232 and Cisco Nexus 5500 Platform Switches


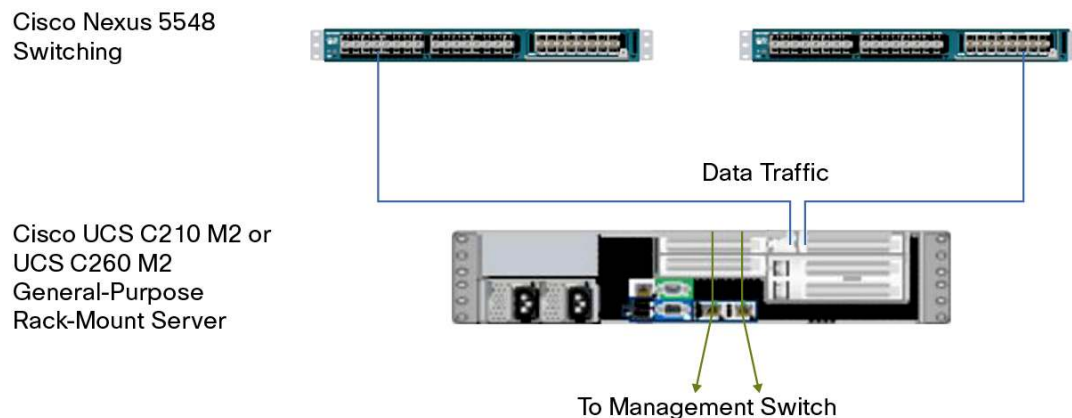


Figure 4: Deployment Topology Using Cisco Nexus 5500 Switch Platform



When considering how to size and configure the storage system, you should consider the different deployment options. For stateless workloads, you may want to use a hypervisor with NexentaStor deployed as a virtual machine using local disks. This configuration can benefit from the caching efficiencies mentioned earlier. For the hypervisor host, load the server with additional RAM because RAM tends to be the limiting resource as the infrastructure scales to support more desktops.

To support persistent (also known as stateful) desktop sessions, you can use dual-ported disk arrays, connected to multiple systems, to protect the data. The hypervisor itself may provide a fault-tolerant architecture with storage migration features, or NexentaStor can be deployed on separate storage servers in a high-availability cluster configuration.

Paging files can be stored in storage pools of solid-state devices within the storage system. For this approach, it is important to make sure that the system has sufficient drive bays to support SSDs.

Best Practices for Configuring Storage for VDI

Recommendations for deploying NexentaStor applications combine commonsense principles with best practices from existing customer deployments. Cisco and Nexenta jointly have validated the desktop virtualization deployments and recommend the following:

- Use NexentaStor's ZFS architecture for end-to-end data integrity, self-healing, and unlimited scalability.
- For optimal performance, use solid-state disks and configure compression, mirrored storage pools, and multiple I/O paths from virtual machines to storage.
- For optimal availability, use high-availability clusters to protect against storage server failures, and use the replication options to protect against disasters.
- Use Nexenta VMDC to get visibility into each virtual machine's use of storage.

Compression can improve performance by reducing I/O at the cost of some processor overhead. Processors on the storage systems typically are underutilized, so compression always should be enabled. With a variety of compression algorithms available, you may want to experiment with different options and perform testing to see if more-complex algorithms yield more compression while also helping ensure that the processor does not become a bottleneck.

You should turn off deduplication for virtual desktop environments.

ZFS supports a variety of RAID options that offer varying levels of space efficiency and redundancy. Mirrored configurations, however, can achieve better performance. Therefore, for VDI, you should configure storage pools with striped mirrors for optimal performance.

NFS often is recommended as the protocol for accessing the virtual backing store. NFS is a synchronous protocol, so you should use solid-state disks for the ZFS intent log to get the best performance.

Additional considerations include the following:

- Use persistent storage for OS images.
- Use temporary storage to cache paging files for desktop sessions.
- Set ZFS log bias = latency, because users are latency sensitive for VDI.
- For XenDesktop, use 4000-block sizes to reduce opportunities for misaligned blocks.
- Use the standard synchronous mode.
- Disable the ZFS prefetch algorithms, because much of the I/O is random.

Service Provider Case Study

Korea Telecom (KT) currently has a public cloud offering that uses NexentaStor and Citrix UCS servers. KT wanted to extend its cloud offering to include a VDI service. The company engaged with Citrix, Cisco, and Nexenta to determine the hardware and software configuration that would achieve optimal performance.

Hardware Configuration

For VDI, both the hypervisor host system and the storage nodes can be performance bottlenecks, so you need to choose these systems wisely. Cisco UCS servers were chosen for both. A summary of the hardware used is shown in Table 1.

Table 1: Hardware Configuration

Server	Cisco UCS C210 M2 (storage node) and Cisco UCS C200 M2 (client)
Processor	Intel Xeon (Westmere-based)
RAM	192 GB
Networking	Intel 10-Gbps Ethernet network interface cards (NICs) Cisco Nexus 5000 Series (Cisco Nexus 5500 platform) 10 Gbps switch or Cisco 6100 or 6200 Series Fabric Interconnect
Host bus adapter (HBA)	LSI 9200-4e4i

Software Configuration

NexentaStor on Cisco UCS servers as storage nodes can be used with all major hypervisors. In this specific case study, Citrix XenServer was deployed. The software used and clients tested are shown in Table 2.

Table 2: Software Configuration

Virtual host	Citrix XenServer 5.6
	Citrix XenDesktop 5
Client systems	Microsoft Windows Server 2008 R2
	Microsoft Windows 7 Enterprise Edition, 32-bit

Client systems were configured with 1 GB of RAM and a 24-GB virtual disk.

Storage Configuration

NexentaStor was deployed on the storage node to provide the high-performance file sharing and data protection required for VDI. Storage was split into two independent pools: one for document storage and another for the OS and paging file storage. Storage was shared with the hypervisor host using NFS.

The storage for the OS and paging files was the focus of much of the performance study (Table 3).

Table 3: Storage Configuration


Pool redundancy	RAIDz2 (double parity)
Disks for main pool	10x 600-GB SAS disks, 10,000 RPM
ZFS intent log	2x 100-GB SSDs, mirrored
Read cache	None
Spare disks	None

Performance Results

KT developed specific test scenarios for VDI user workloads. These scenarios tested the processes of logging in, accessing Microsoft Word and PowerPoint files, and logging out. Microsoft Word files were 16 MB, and Microsoft PowerPoint files were 42 MB.

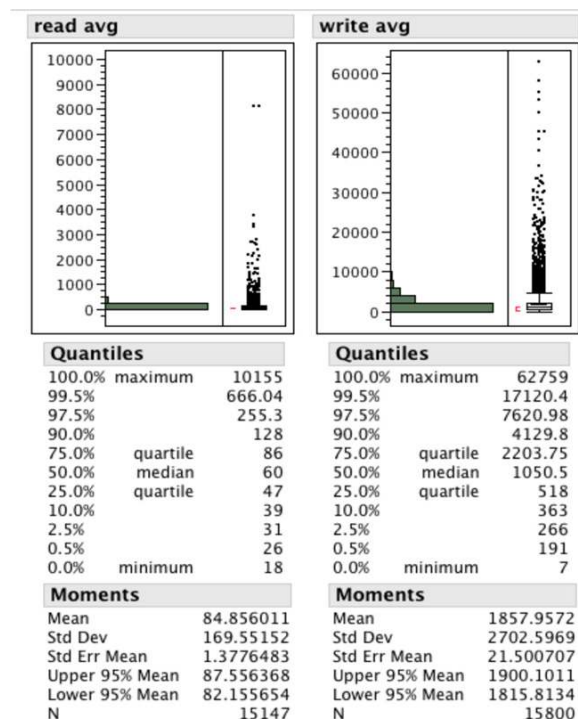
Figure 5 shows how the number of NFS operations varied during the test run. At the beginning of the run, you can see an increase in read operations as the desktops are booting. Later, the I/O characteristics change to be more write-intensive.

Figure 5: NFS Operations: Performance over Time



The corresponding read and write response times can be seen in the charts from one of the client systems (Figure 6). The median response time for writes is approximately 1 millisecond (ms). This time largely was dominated by the write time of the separate log device and could improve if higher-performing SSDs are used. The read operations benefited from the system cache and had a median response time of 60 ms.

Figure 6: Read and Write Response Times



Based on the results of the study, it was concluded that a single Cisco UCS 210 storage system could support a minimum of 360 virtual desktop users.

Conclusion

Using the enterprise architecture of Cisco UCS, the NexentaStor-based storage solution can offer exceptional performance for desktop virtualization solutions at an effective price point.

NexentaStor, deployed on Cisco UCS C-Series Rack-Mount Servers, offers exceptional price-performance advantages for VDI. NexentaStor uses the enterprise architecture of the Cisco UCS C-Series servers to deliver reliable and optimal storage for desktop virtualization solutions.

For More Information

<http://www.cisco.com/go/vdi>

© 2012 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.co/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.(1110R)

© 2012 Nexenta® Systems, All Rights Reserved. Nexenta Systems is the leading supplier of enterprise-class OpenStorage solutions. Its flagship software-only platform, NexentaStor, delivers high-performance, ultra-scalable, cloud- and virtualization-optimized storage solutions.