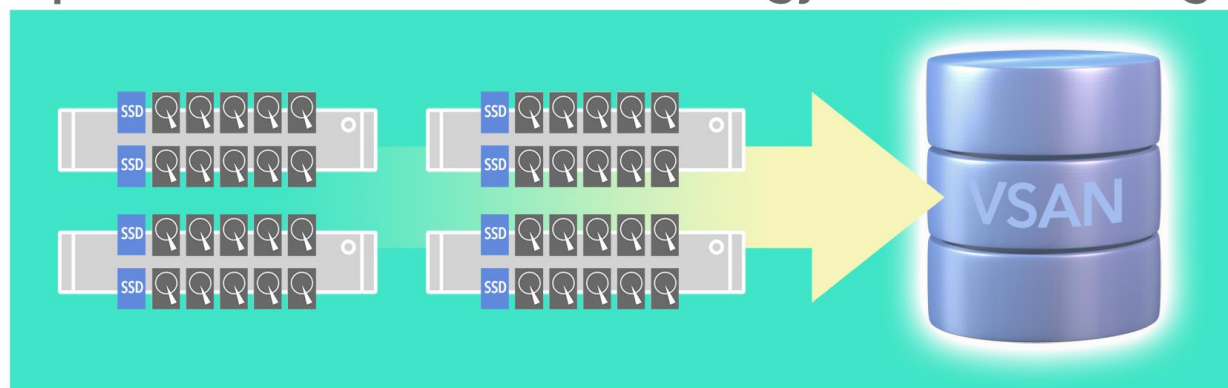


COST AND PERFORMANCE COMPARISON FOR OPENSTACK COMPUTE AND STORAGE INFRASTRUCTURE

OpenStack with VMware technology and VSAN storage



Performed 53%
more NoSQL operations
per second

Performed 159%
more IO operations
per second

projected to
cost 26% less
over 3 years

vmware®

versus Red Hat OpenStack on Red Hat Storage

Hyper-converged architectures are emerging that can bring increased performance and significant cost reduction to virtualized infrastructures, where compute, network, and storage coexist closely on physical resources. Another trend shows cloud frameworks, such as OpenStack, are emerging and maturing to offer APIs for more efficient self-service provisioning and consumption of such resources. The combination of these industry trends and technology developments provides customers with many options when choosing the underlying virtualization technology for their cloud.

Two key components of cloud infrastructure design factor into its performance—the hypervisor itself and its underlying storage—making the architecture of the storage a critical consideration for businesses. Two main approaches to storage architectures are converged and non-converged. Non-converged storage exists on its own “tier,” separate from the compute infrastructure, and is typically shared storage either provided by traditional storage arrays or built from commodity server hardware running specialized software. Red Hat Storage Server uses this model by aggregating disks on multiple commodity server nodes and running a Red Hat version of GlusterFS.

The converged storage model, which also presents shared storage to compute resources, collapses the storage and compute hardware tiers and leverages local disks in each compute node so that performance and capacity can be scaled out together. VMware Virtual SAN takes this converged approach, using disks in each compute node, while also presenting the aggregate storage pool to the vSphere cluster as a datastore.



In the Principled Technologies labs, we compared these two solutions in the context of an OpenStack environment: VMware vSphere and VMware Virtual SAN and Red Hat Enterprise Linux® and the KVM hypervisor with Red Hat Storage Server, based on the Gluster distributed file system.

We used two different approaches in our testing, both of which ran in the context of an OpenStack infrastructure. First, to test in a realistic cloud application, we used the Yahoo!® Cloud Serving Benchmark (YCSB) with an Apache Cassandra distributed database. We chose Cassandra because of its popular distributed NoSQL database platform and because the results could provide data to support the cloud trends happening in the market today.

Next, to test the raw capabilities of the hardware and software without an application framework, we used the Flexible I/O (FIO) benchmark on a test file in each VM. We chose the FIO benchmark to test storage performance without the potential overhead caused by the application and database layers.

We found that the VMware solution provided 53 percent more YCSB operations per second (OPS), as well as 159 percent more IOPS during a mixed read-write FIO workload than the Red Hat solution using the KVM hypervisor and Red Hat Storage Server. The VMware solution also had a 26 percent lower cost and used less physical space in the datacenter.

HYPER-CONVERGENCE FOR CLOUD PERFORMANCE

OpenStack as a cloud framework can enable many types of applications to be deployed on a large scale. One such application is a NoSQL distributed database, such as Apache Cassandra. Underneath this distributed database must be a storage solution that can provide high performance under load while also scaling out to accommodate new nodes. Furthermore, commonplace IO-intensive actions, such as deploying instances, uploading images, and moving virtual machines, often create bottlenecks in OpenStack cloud environments, so strong resource management solutions are crucial. In the relatively new and fast-growing world of software-defined storage, how does a business using OpenStack balance performance and cost considerations?

VMware Virtual SAN brings all the proven benefits of shared storage in a VMware environment, such as High Availability (HA) and vMotion, while using direct-attached disks on the compute hosts. Using VMware Virtual SAN with VMware vSphere has several crucial characteristics that make it attractive for hyper-converged storage management:

1. VMware vSphere is a small-footprint hypervisor that requires less frequent patching than to a general-purpose OS hypervisor.
2. Native HA minimizes downtime in the event of a host outage.

3. VMware Distributed Resource Scheduler (DRS) uses vMotion to balance workloads to accommodate shifting demand and has the ability to patch infrastructure with zero downtime.
4. More advanced memory management technologies lead to potentially higher VM density.
5. VMware vSphere has broad and proven guest OS support.

VMware Virtual SAN integrates tightly with the hypervisor, and scales easily by adding more hosts to a cluster or more storage to existing hosts. In addition, VMware Virtual SAN can be managed directly through the familiar vCenter Server™ Web client console, alongside everything else in a VMware vSphere environment.

In a VMware Virtual SAN-enabled host, every disk chosen for Virtual SAN storage belongs to a disk group. Each disk group has one solid-state drive (SSD) and up to seven hard drives (HDDs). The SSD in a disk group serves as a read and write cache, with writes acknowledged by the SSD and later de-staged to the HDDs. Storage policies allow administrators to control stripe width, failures-to-tolerate, and storage reservation percentage on a per-VM basis. Additional storage or hosts add to the capacity and performance of a VMware Virtual SAN datastore without disruption.

VMware Virtual SAN integrates advanced features of more expensive external storage into the hypervisor, retaining high performance characteristics and allowing for space and cost savings.

OUR ENVIRONMENT

In both the VMware and Red Hat environments, we used the same number of disk slots (12) per storage node for VM virtual disk storage. Our VMware Virtual SAN solution consisted of four Dell PowerEdge R720s, each running VMware vSphere 5.5. Each server contained two disk groups, with one SSD and five HDDs in each disk group, to total 12 drive bays used per server. We installed the hypervisor onto internal dual SD cards and dedicated all the drive bays to Virtual SAN storage. We managed the environment with a separate machine running VMware vCenter Server. We created an OpenStack environment using Ubuntu 12.04 LTS and the Havana release of open source OpenStack. Finally, we connected the OpenStack environment to the vCenter Server to complete the setup.

Red Hat recommends separate hardware for storage and compute, so our Red Hat Storage Server solution for testing OpenStack required twice as many servers for the non-converged architecture of KVM and Red Hat Storage Server. We set up the Red Hat recommended minimum of four Dell PowerEdge R720s running Red Hat Storage Server 2.1, and four Dell C8000 servers running Red Hat Enterprise Linux 6.5 with KVM to serve as compute nodes. We used a pair of disks for the OS in the compute nodes,

and we used 14 total drive bays in each server for Red Hat Storage Server – 12 bays were used for data storage, and two bays were used for the OS. Each Dell PowerEdge R720 contained 12 HDDs in a RAID6 configuration for data storage, plus two HDDs in RAID1 for the operating system (12 HDDs per node is the Red Hat recommendation for general-purpose use cases).¹ We used no SSDs, as there is no native support for tiering using different drive types in Red Hat Storage Server. On separate hardware, we created a Red Hat OpenStack 4.0 environment and connected it to the compute and storage nodes to complete the setup.

Both environments used the Dell PERC H710P RAID controller. The Red Hat Storage requires hardware RAID6, RAID10, and a flash-backed or battery-backed cache. In contrast, Virtual SAN works in pass-through or RAID0 mode, which eliminates the need for advanced hardware RAID (we used RAID0 on the PERC H710P). Both environments used a 10Gb Ethernet switch to handle all traffic, and each had an additional server to host the OpenStack controller VM and client VMs for running the benchmark.

Once we configured the OpenStack environment for each solution, we used Glance to upload a CentOS 6.5 image to each OpenStack environment, and then launched 16 instances using that image. We used the 16 VMs to exercise the storage by running our selected two workloads. The first workload we used ran YCSB to generate and measure load, which we installed onto four client VMs existing on a separate server. We installed the Cassandra distributed database onto the 16 test VMs and used YCSB to load 400 million records into the database. We then used YCSB to perform read and insert operations on the data, with a 70/30 read-to-insert ratio. The YCSB test then reported how many operations were performed each second against the Cassandra database. The second workload we used ran as a low-level workload from the Flexible I/O (FIO) 2.1.6.1 benchmark. We used a workload mix of 70/30 reads to writes, accessed randomly. For more details on our exact test parameters, see [Appendix B](#).

¹ access.redhat.com/site/articles/66206

WHAT WE FOUND

As Figure 1 shows, we found that the VMware solution was able to perform 53 percent more YCSB operations per second (OPS).

Figure 1: The amount of YCSB OPS achieved by the two solutions. Higher numbers are better.

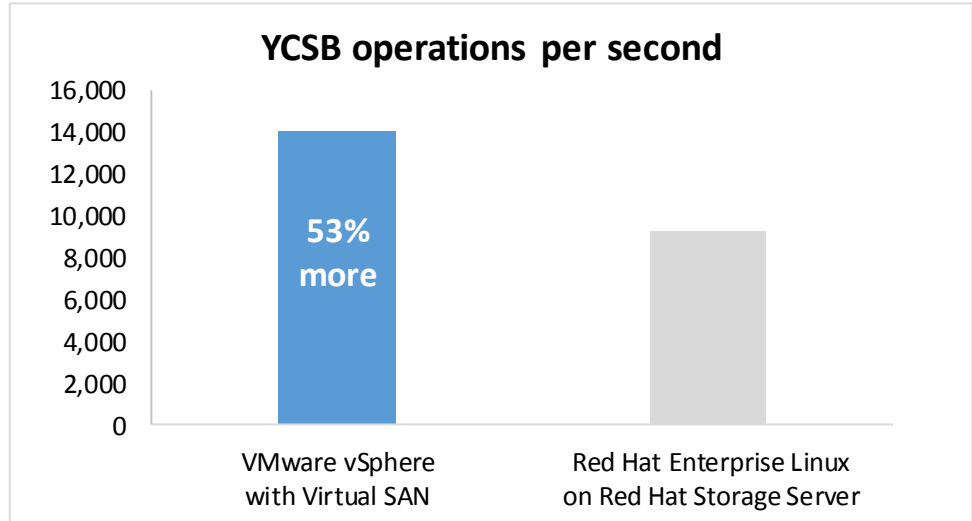
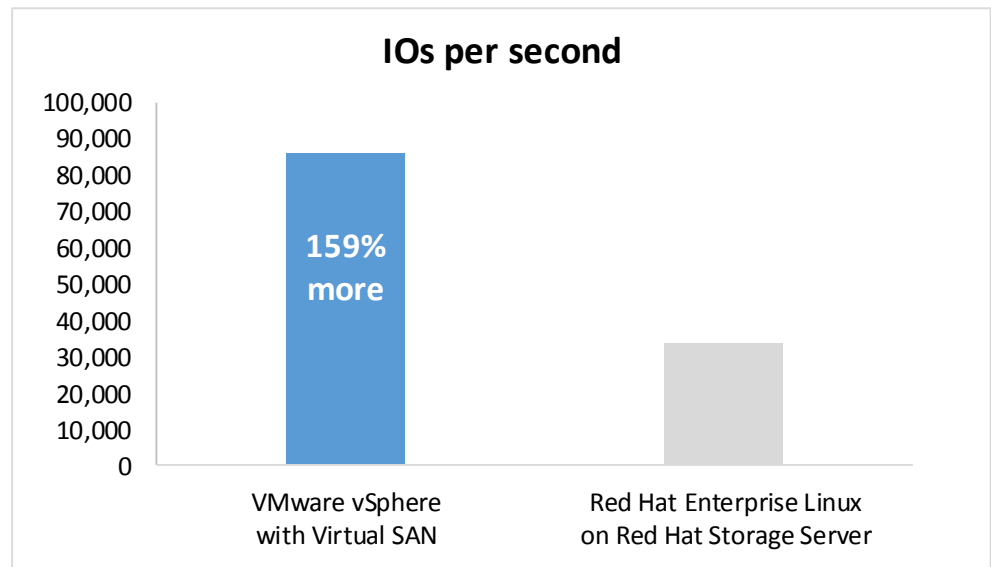


Figure 2 shows the IOPS numbers achieved by the two solutions. We found that the VMware vSphere with VMware Virtual SAN solution provided 159 percent more IOPS than the Red Hat solution.

Figure 2: The amount of IOPS achieved by the two solutions. Higher numbers are better. The workload was 70/30 R/W mix, random, and 4K block size.



DENSITY COMPARISON

By combining storage and compute resources into the same server, VMware vSphere with VMware Virtual SAN can substantially reduce the physical footprint that a cloud infrastructure requires. Advantages to space consolidation typically include significant savings in power, cooling, rack space, and switch ports. Whereas the VMware vSphere with VMware Virtual SAN solution used just four Dell PowerEdge R720 servers, the Red Hat Storage solution required an additional four Dell PowerEdge C8000 servers as compute nodes, which in turn required four additional 10Gb network ports on a single switch. VMware Virtual SAN eliminates the need for compute and storage resources to be on separate hardware, and allows administrators to add compute and storage resources in a linear fashion based on demand.

PRICE COMPARISON

Our VMware vSphere with Virtual SAN solution was not only faster and more space-efficient than the Red Hat solution—it was also cheaper. We compared prices for the VMware vSphere with Virtual SAN and the Red Hat Storage solutions we tested. Both solutions had similarities in hardware, such as four Dell PowerEdge R720s with 128GB of memory, but the Red Hat solution required four additional servers and more expensive software costs.

While we used the Dell PowerEdge C8000 servers for the performance testing, we did not use these servers in the pricing analysis, but instead substituted Dell PowerEdge R720 servers. The Dell PowerEdge C8000 servers used in testing for the Red Hat compute nodes required a chassis for power and networking. Rather than using the C8000 costs for the Red Hat solution, we compared the hardware costs of four (VMware) vs. eight (Red Hat) Dell PowerEdge R720s for consistency and equality in pricing.

Figure 3 shows the projected costs for the two tested solutions over three years. We include costs for compute hardware, storage hardware, and the hypervisor. Each OpenStack deployment and environment is different and support engagements vary widely from platform to platform; therefore, we do not include pricing for OpenStack support on either platform. In our configuration, the VMware vSphere with Virtual SAN solution can save up to \$44,398, or 26 percent, in hardware and software costs compared to the Red Hat Storage solution.

Figure 3: Projected three-year costs for the two solutions. Lower numbers are better

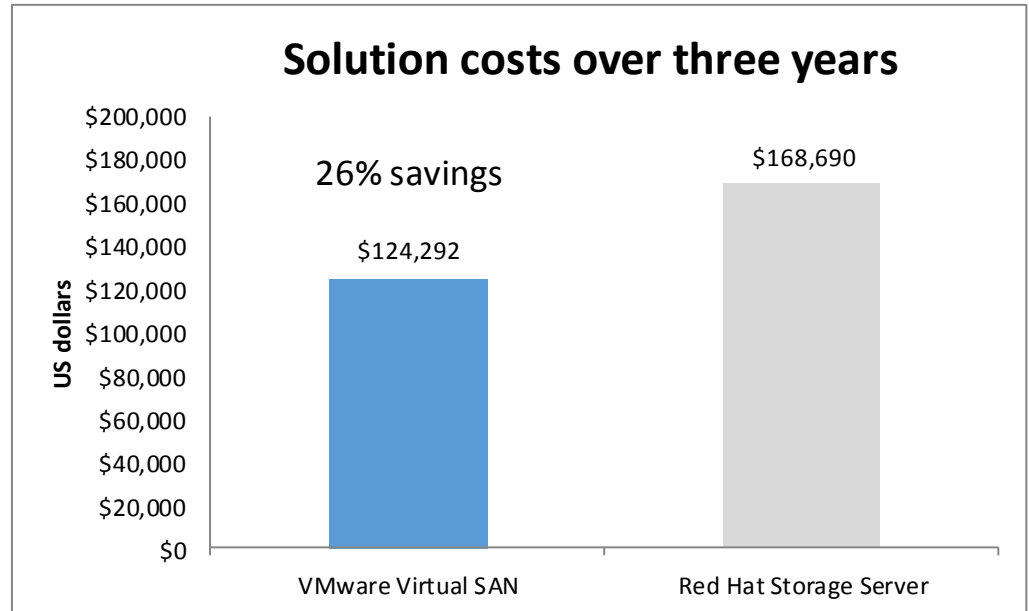


Figure 4 shows the costs for the VMware vSphere with VSAN solution, including solid-state drives (SSDs), hard disk drives (HDDs), VMware vSphere Enterprise, and VMware vCenter Server software. As we note above, each VMware vSphere node contained 10 HDDs, for a total of 40 HDDs in the solution. Each node contained two SSDs, for a total of 8 SSDs in the solution. We ran the hypervisor from the internal SD cards inside each server. For vSphere and vCenter Server, we included three years of Production support. Prices for vSphere and vCenter are from the VMware Web site; the Virtual SAN price is from a product announcement. For the Dell PowerEdge R720s, we used the online Dell configuration tool and selected the PERC H710P, a 16-drive bay, three years of Pro support, and no memory or hard drives.

VMware	Model	Quantity	Price	Total
Servers	Dell PowerEdge R720	4	\$6,438.87	\$25,755.48
SSDs	Intel SSD DC S3700 200GB	8	\$489.00 ²	\$3,912.00
HDDs	Dell/Seagate ST9300653SS	40	\$330.41 ³	\$13,216.40
VSAN	VMware VSAN 1-socket license (initial license)	8	\$2,495.00 ⁴	\$19,960.00
VSAN SnS	VMware VSAN Production support, 25% per socket (per year, 3 years)	24	\$623.75	\$14,970.00
VMware vSphere	VMware vSphere Enterprise with 3-year Production support for 1 CPU	8	\$4,773.16 ⁵	\$38,185.28
VMware vCenter Server	VMware vCenter Server Standard Edition with 3-year Production support (One instance of VMware vCenter Server is required for vSphere deployments.)	1	\$8,292.36 ⁶	\$8,292.36
Total				\$124,291.52

Figure 4: Detailed pricing of the components for the VMware vSphere with Virtual SAN solution, projected over three years.

Figure 5 shows the costs for the Red Hat solution, including HDDs, Red Hat Storage Server, and Red Hat Enterprise Linux software. While Red Hat does provide Red Hat Enterprise Linux OpenStack Platform, we left OpenStack support out of these calculations as stated above because each OpenStack environment and support engagement is so variable. As we note above, each Red Hat Enterprise Linux compute node contained two HDDs and each Red Hat Storage Server contained 14 HDDs, for a total of 64 HDDs in the solution. For this calculation, we chose to use Red Hat Enterprise Linux Premium, which provides continuous around-the-clock support like the Production support for VMware products.⁷ This solution did not include SSDs. For Red Hat software, we selected Premium support. The Dell PowerEdge R720 prices are as used in the VMware solution.

² www.serversdirect.com/components/drives?utm_source=bing_yahoo&utm_medium=cpc&utm_term=intel%20ssd&utm_content=3973866112&utm_campaign=Intel

³ www.memory4less.com/m4l_itemdetail.aspx?itemid=1463911782&partno=ST9300653SS&rid=99

⁴ www.vmware.com/ca/en/company/news/releases/VMware-Announces-General-Availability-of-VMware-Virtual-SAN

⁵ store.vmware.com/store/vmware/en_US/pd/productID.288070900&src=WWW_eBIZ_productpage_vSphere_Enterprise_Buy_US

⁶ *Ibid.*

⁷ Standard business hours (24x7 for Severity 1 and 2): access.redhat.com/site/support/offerings/production/sla

Red Hat	Model	Quantity	Price	Total
Servers	Dell PowerEdge R720	8	\$6,438.87	\$51,510.96
SSDs	(none)			
HDDs	Dell/Seagate ST9300653SS	64	\$330.41 ⁸	\$21,146.24
Red Hat Storage	Red Hat Storage Server for On-premise, Premium (4 nodes), three-year	1	\$81,225.00 ⁹	\$81,225.00
Red Hat Enterprise Linux Server	Red Hat Enterprise Linux Server, Premium (Physical or Virtual Nodes) (1 node), three-year	4	\$3,702 ¹⁰	\$14,808.00
Total				\$168,690.20

Figure 5: Detailed pricing of the components for the Red Hat Storage Server solution, projected over three years.

CONCLUSION

To be competitive in the cloud, businesses need more efficient hardware and software for their OpenStack environment and solutions that can pool physical resources efficiently for consumption and management through OpenStack. Software-defined storage has made it easier to create storage resource pools that spread across your datacenter, but external or separate distributed storage systems are still a challenge for many. VMware vSphere with Virtual SAN converges storage resources with the hypervisor, which can allow for management of an entire infrastructure through a single vCenter Server, increase performance, save space, and reduce costs.

Both solutions used the same twelve drive bays per storage node for virtual disk storage, however the tiered design of VMware Virtual SAN allowed for greater performance in our tests. We used a mix of two SSDs and 10 rotational drives for the VMware Virtual SAN solution, while we used 12 rotational drives for the Red Hat Storage Server solution behind a battery-backed RAID controller, the Red Hat recommended approach.

In our testing, the VMware vSphere with Virtual SAN solution performed better than the Red Hat Storage solution in both real world and raw performance testing by providing 53 percent more database OPS and 159 percent more IOPS. In addition, the vSphere with Virtual SAN solution can occupy less datacenter space, which can result in lower costs associated with density. A three-year cost projection for the two solutions

⁸ www.memory4less.com/m4l_itemdetail.aspx?itemid=1463911782&partno=ST9300653SS&rid=99

⁹ www.redhat.com/f/html/partners_us_skulist.html

¹⁰ www.redhat.com/about/news/archive/2013/7/red-hat-announces-general-availability-of-new-infrastructure-solutions-red-hat-openstack-certification-update

showed that VMware vSphere with Virtual SAN could save your business up to 26 percent in hardware and software costs when compared to the Red Hat Storage solution we tested.

APPENDIX A – SYSTEM CONFIGURATION INFORMATION

Figure 7 provides detailed configuration information for the test systems.

System	Dell PowerEdge R720	Dell PowerEdge C8000
Power supplies		
Total number	2	2
Vendor and model number	Dell D750E-S1	Dell D1200E-S2
Wattage of each (W)	750	1200
Cooling fans		
Total number	6	3 shared across chassis
Vendor and model number	AVC DBTC0638B2V	Delta Electronics PFC1212DE
Dimensions (h × w) of each	2.5" × 2.5"	5.5" × 5.5"
Volts	12	12
Amps	1.2	4.8
General		
Number of processor packages	2	2
Number of cores per processor	8	8
Number of hardware threads per core	2	2
System power management policy	Default	Default
CPU		
Vendor	Intel	Intel
Name	Xeon	Xeon
Model number	E5-2650	E5-2650
Socket type	LGA 2011	LGA 2011
Core frequency (GHz)	2.00	2.00
Bus frequency	8 GT/s	8 GT/s
L1 cache	32 KB + 32 KB (per core)	32 KB + 32 KB (per core)
L2 cache	256 KB (per core)	256 KB (per core)
L3 cache	20 MB	20 MB
Platform		
Vendor and model number	Dell PowerEdge R720	Dell PowerEdge C8000
Motherboard model number	00W9X3	0W6W6G
BIOS name and version	Dell 2.2.2	Dell 2.3.1
BIOS settings	Default	Default
Memory module(s)		
Total RAM in system (GB)	128	128
Vendor and model number	Samsung M393B1K70BH1-CH9	Samsung M393B1K70DH0-CK0
Type	PC3-10600R	PC3-12800R
Speed (MHz)	1,333	1,600
Speed running in the system (MHz)	1,333	1,600
Timing/Latency (tCL-tRCD-tRP-tRASmin)	9-9-9-36	11-11-11-27
Size (GB)	8	8
Number of RAM module(s)	16	16

System	Dell PowerEdge R720	Dell PowerEdge C8000
Chip organization	Double-sided	Double-sided
Rank	Dual	Dual
Operating system		
Name	VMware vSphere ESXi 5.5	Red Hat Enterprise Linux 6.5
Build number	1623387	2.6.32-431.5.1.e16.x86_64
File system	VMFS	Ext4
Language	English	English
RAID controller		
Vendor and model number	Dell PERC H710P	LSI SAS2008
Firmware version	21.2.0-0007	0.61
Driver version	00.00.05.34-9vmw	15.101.00.00
Cache size (MB)	1024	0
Solid-state drives (VMware and Red Hat compute nodes)		
Vendor and model number	Intel SSD DC S3700	Samsung MZ5EA100HMDR
Number of drives	2	2
Size (GB)	200	100
Type	SATA	SATA
Hard drives (VMware)		
Vendor and model number	Dell ST9300653SS	N/A
Number of drives	10	N/A
Size (GB)	300	N/A
RPM	15k	N/A
Type	SAS	N/A
Hard drives (Red Hat)		
Vendor and model number	Dell ST9300653SS	N/A
Number of drives	12 + 2 for OS	N/A
Size (GB)	300	N/A
RPM	15k	N/A
Type	SAS	N/A
Ethernet adapters		
Vendor and model number	Intel I350/X540-A2	Intel I350
Type	1/10Gb	1Gb
Driver	3.7.13.7.14iov-NAPI	5.0.5-k

Figure 7: System configuration information for the test systems.

APPENDIX B – HOW WE TESTED

These steps document how we constructed our test environments. In both the VMware and Red Hat solution, we used the four Dell PowerEdge R720 servers as our storage nodes and a separate Dell PowerEdge R410 to host VMs for OpenStack, vCenter Server, and a FIO client. In the Red Hat solution, we added an additional four Dell PowerEdge C8000 servers for compute nodes.

Constructing the VMware environment

Installing VMware vSphere onto the Dell PowerEdge R720 servers

Repeat these steps for each Dell PowerEdge R720.

1. Insert the vSphere installation media and power on the server, choosing to boot from the DVD drive.
2. Select the standard vSphere installer and allow the files to copy into memory.
3. Press F11 at the welcome screen.
4. Press Enter at the keyboard language selection screen.
5. Enter a password twice for the root user and press Enter.
6. Choose to install to the Dell Internal Dual SD Cards.
7. Allow the installer to finish installing vSphere and reboot the server.
8. At the home screen, press F2 and enter your credentials.
9. Select Configure Management Network.
10. Select IP Configuration.
11. Enter your static IP information and press Enter.
12. Select DNS Configuration.
13. Enter a hostname and press Enter.
14. Press Esc and then Y to restart the management network.

Installing VMware vCenter Server Virtual Appliance

Use the steps above to install vSphere onto a server outside of the testing environment. Navigate to the IP address and download the vSphere desktop client. Use the client to connect to the server, and follow these steps.

1. Click File→Deploy OVF Template...
2. Browse to the location of your VMware vCenter Server Appliance .ova or .ovf file and click Next.
3. Read over the template properties and click Next.
4. Enter a name for the virtual appliance and click Next.
5. Choose a datastore for the virtual appliance and click Next.
6. Choose hard disk configuration options for the virtual appliance and click Next.
7. Choose the networking configuration for the virtual appliance and click Next.
8. Click Finish and allow the virtual appliance to deploy.
9. Power on the virtual appliance and login to the console. The VMware default username is `root` and the default password is `vmware`.
10. Run `/opt/vmware/share/vami/vami_config_net` to configure the networking.

11. Using a Web browser on the same network, navigate to `https://<vCenter virtual appliance IP address>:5480` and log in with the default credentials.
12. Follow the wizard to configure vCenter. We used no time synchronization, embedded database, and no active directory.

Building the Virtual SAN environment

1. Using a Web browser on the vCenter network, navigate to `https://<vCenter virtual appliance IP address>:9443/vsphere-client` and login using the credentials you set during the virtual appliance configuration.
2. Create a new datacenter.
3. Add each host to the datacenter.
4. Select the first host, and add a new VMkernel adapter on a new virtual switch. Select the checkbox for Virtual SAN network traffic. Provide a 10GbE physical NIC and an IP address separate from the management network.
5. Repeat step 4 for the remaining three hosts.
6. Create a new cluster with DRS turned on and set to Fully Automated. Leave Virtual SAN turned off for now.
7. Move the four hosts into the cluster.
8. In the cluster settings, enable Virtual SAN and select Manual disk management.
9. In the cluster disk management, create 2 disk groups for each host. Each disk group should have 1 SSD and 5 HDDs.

Building the OpenStack environment

Our implementation of OpenStack was built upon two VMs (`controller`, and `compute`) that resided on physical hardware separate from our testbed. Each VM ran Ubuntu 12.04 LTS and had two NICs – one connected to the same network as vCenter (internal – 192.168.3.xxx) and one connected to our VM network (external – 10.0.0.xxx). Our setup closely follows the documentation found here: docs.openstack.org/havana/install-guide/install/apt/content/index.html. We used OpenStack patches made available by VMware to allow VMware Virtual SAN to work with the Havana release of OpenStack.

Building the test VMs

1. Using a separate server with the same build of vSphere, create a VM with the configuration desired. In our case, we used a minimal CentOS 6.5 installation on a VM with 4 vCPUs, 2GB of RAM, and a 60GB VMDK.
2. Shut down the VM and export it as an OVF template.
3. Retrieve the VMDK file and transfer it to the OpenStack controller VM.
4. Run the following command to upload the image to Glance:

```
glance image-create --name=base --disk-format=vmdk --container-format=bare \
--is-public=True --property vmware_linked_clone="False" \
--property vmware_disktype="streamOptimized" < ovf_exported.vmdk
```
5. Use a Web browser to navigate to the IP address of the controller VM and login to Dashboard.
6. Click Flavors and create a new flavor with the appropriate sizing.
7. Under admin, click Instances.
8. Launch an instance with the appropriate flavor and select the base image under Boot from Image.
9. Repeat step 8 until there are 16 instances created.

Constructing the Red Hat environment

Installing Red Hat Storage 2.1 onto the Dell PowerEdge R720 servers

Repeat these steps for each Dell PowerEdge R720.

1. Insert the Red Hat Storage installation media and power on the server, choosing to boot from the DVD drive.
2. Select Install or upgrade an existing system and press Enter.
3. Use the TCP/IP configuration screen to add a static IP address.
4. Press Next at the installation welcome screen.
5. Choose a language and press Enter.
6. Choose a keyboard layout and press Enter.
7. Select Basic Storage Devices and click Next.
8. Enter a hostname and click Next.
9. Select your time zone and click Next.
10. Enter a root password twice and click Next.
11. Select Use All Space and click Next.
12. At the Boot Loader screen, click Next.
13. At the Minimal Selection screen, click Next.
14. After the installation completes, reboot the server.

Installing Red Hat Enterprise Linux onto the Dell PowerEdge C8000 servers

Repeat these steps for each Dell PowerEdge C8000.

1. Insert the Red Hat Storage installation media and power on the server, choosing to boot from the DVD drive.
2. Select Install or upgrade an existing system and press Enter.
3. Use the TCP/IP configuration screen to add a static IP address.
4. Press Next at the installation welcome screen.
5. Choose a language and press Enter.
6. Choose a keyboard layout and press Enter.
7. Select Basic Storage Devices and click Next.
8. Enter a hostname and click Next.
9. Select your time zone and click Next.
10. Enter a root password twice and click Next.
11. Select Use All Space and click Next.
12. At the Boot Loader screen, click Next.
13. At the Minimal Selection screen, click Next.
14. After the installation completes, reboot the server.

Installing Red Hat OpenStack Platform 4

We installed Red Hat OpenStack onto a VM on a separate server. We placed all services on one VM running Red Hat Enterprise Linux 6.5 and had one NIC connected to the compute network. Our setup closely follows the documentation found here: https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux_OpenStack_Platform/4/html-single/Installation_and_Configuration_Guide/index.html

Building the test VMs

1. Upload the VMDK file used from the VMware solution to the Red Hat OpenStack controller.
2. Convert the VMDK file to qcow2 format:

```
qemu-img convert -f vmdk -O qcow2 ovf_exported.vmdk base.qcow2
```

3. Upload the converted disk to Glance:

```
glance image-create --name=base --disk-format=qcow2 --container-format=bare \
--is-public=True --file=base.qcow2
```

4. Use a Web browser to navigate to the IP address of the controller VM and login to Dashboard.
5. Click Flavors and create a new flavor with the appropriate sizing.
6. Under admin, click Instances.
7. Launch an instance with the appropriate flavor and select the base image under Boot from Image.
8. Repeat step 8 until there are 16 instances created.

Running the benchmarks

Building the YCSB and Cassandra environment

The YCSB and Cassandra environment consists of 16 VMs running the Cassandra back-end, and four client VMs running the YCSB benchmarking client. Each of the 20 VMs begins as a minimal CentOS installation with the Software Development Workstation package added.

1. Set a static IP address and turn off iptables for each VM.
2. Edit /etc/hosts on each VM to reflect the hostnames of your choice.
3. Download YCSB and Maven onto the client VMs and Cassandra onto the back end VMs. We used the following links:
 - YCSB: `git clone https://github.com/cmatser/YCSB.git`
 - Maven: <http://maven.apache.org/download.cgi>
 - Cassandra: <http://cassandra.apache.org/download/>
4. On the client VMs, run `mvn clean package` in the YCSB directory to build YCSB.
5. On the backend VMs, configure `cassandra.yaml` to reflect the following:
 - `listen_address: <local IP address>`
 - `broadcast_address: <local IP address>`
 - `rpc_address: <local IP address>`
 - `seed: "<node1 IP address>"`
6. Run `ant` in the Cassandra directory to build Cassandra.
7. Start Cassandra on the seed nodes (node 1). Cassandra is started by running `bin/cassandra`. Once the line "Startup completed! Now serving reads." appears, Cassandra has started and pressing Enter will return to the shell.
8. Start Cassandra on the subsequent nodes. Output should reflect a bootstrapping process as they receive configuration information from the seed node.
9. On one of the Cassandra nodes, run `bin/cqlsh <node1 IP address> 9160` to connect the Cassandra CLI. Run the following commands to create and configure a keyspace and table for YCSB to use. Take care to

include the semicolon after each command and keep the names `ycsb` and `usertable` as they are YCSB defaults.

- create keyspace `ycsb` with replication = { 'class' : 'SimpleStrategy', 'replication_factor' : 1 };
- use `ycsb`;
- create table `usertable(y_id varchar primary key, field0 varchar, field1 varchar, field2 varchar, field3 varchar, field4 varchar, field5 varchar, field6 varchar, field7 varchar, field8 varchar, field9 varchar)`;

10. Copy the Cassandra `.jar` files in `...apache-cassandra-2.0.4/lib/` to `/YCSB/cassandra/target/` on the client VMs.

Running the YCSB benchmark

We first used one of the YCSB clients to generate 400,000,000 records into the Cassandra database. We then edited one of the standard workload files to reflect a 70/30 read to insert ratio and ran the YCSB client from four VMs simultaneously.

To generate the dataset, run the following command from one YCSB client:

```
./bin/ycsb load cassandra-cql -P workloads/workloadc_70_30 -p  
host=192.168.1.101 -p port=9042 -p recordcount=400000000 -p  
operationcount=400000000 -p threads=16 -s > load.dat
```

To run the benchmark, run the following command from each YCSB client simultaneously:

```
./bin/ycsb load cassandra-cql -P workloads/workloadc_70_30 -p  
host=192.168.1.101 -p port=9042 -p recordcount=400000000 -p  
operationcount=8000000 -p threads=128 -s > run.dat
```

Installing FIO onto the test VMs

Our FIO environment consisted of 16 test VMs with the same specification as the Cassandra VMs. Repeat the following steps for each test VM.

1. Download FIO 2.1.6.1 from an appropriate mirror and extract the package.
2. Edit the "configure" file to change the "ext4 move extent" flag to `yes`
3. Run `configure`
4. Run `make install`
5. Create a 6.2 GB test file by running:

```
dd if=/dev/zero of=/fio_test_file oflag=direct bs=1M count=6200
```

6. Run FIO on the test file:

```
fio --ioengine=libaio --name=test1 --filename=/fio_test_file --overwrite=1 -  
-direct=1 --iodepth=64 --bs=4k --rw=randrw --rwmixread=70 --size=6200m --  
runtime=300
```

APPENDIX C – TEST RESULTS

Our testing showed that VMware vSphere with VSAN outperformed the Red Hat solution in the NoSQL Cassandra database testing, as well as raw disk performance. VMware vSphere with VSAN was also cheaper over three years. Figure 8 shows each result.

Solution	YCSB OPS	IOPS	Cost
VMware vSphere with VSAN	14,063	85,832	\$124,291.52
RHEL on Red Hat Storage	9,208	33,148	\$168,690.20

Figure 8: Performance results for the VMware vSphere with VSAN solution and the Red Hat Storage Server solution.

ABOUT PRINCIPLED TECHNOLOGIES



Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.

Principled Technologies is a registered trademark of Principled Technologies, Inc.
All other product names are the trademarks of their respective owners.

Disclaimer of Warranties; Limitation of Liability:

PRINCIPLED TECHNOLOGIES, INC. HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, PRINCIPLED TECHNOLOGIES, INC. SPECIFICALLY DISCLAIMS ANY WARRANTY, EXPRESSED OR IMPLIED, RELATING TO THE TEST RESULTS AND ANALYSIS, THEIR ACCURACY, COMPLETENESS OR QUALITY, INCLUDING ANY IMPLIED WARRANTY OF FITNESS FOR ANY PARTICULAR PURPOSE. ALL PERSONS OR ENTITIES RELYING ON THE RESULTS OF ANY TESTING DO SO AT THEIR OWN RISK, AND AGREE THAT PRINCIPLED TECHNOLOGIES, INC., ITS EMPLOYEES AND ITS SUBCONTRACTORS SHALL HAVE NO LIABILITY WHATSOEVER FROM ANY CLAIM OF LOSS OR DAMAGE ON ACCOUNT OF ANY ALLEGED ERROR OR DEFECT IN ANY TESTING PROCEDURE OR RESULT.

IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC. BE LIABLE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ITS TESTING, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC.'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH PRINCIPLED TECHNOLOGIES, INC.'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES ARE AS SET FORTH HEREIN.
