

Dell Compellent Storage Center SAN & VMware View 5 1,000 Desktop Reference Architecture

Dell Compellent Technical Solutions



Compellent

THIS REFERENCE ARCHITECTURE IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2012 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the DELL logo, and the DELL badge and Compellent are trademarks of Dell Inc. VMware, ESXi and View 5 are trademarks of VMware Inc. Microsoft are trademarks of Microsoft Corp. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

January 2012

Contents

- Introduction 5
- Executive summary 5
- Key findings 6
 - Dell Compellent Fast Track 8
 - Load Generation 8
 - Virtual Machine Properties 9
 - VMware ESXi 5 9
 - Windows 7 11
 - Antivirus 11
 - Storage Center Configuration 11
 - Data Progression, Replay Technology 13
- Testing Results 14
 - Boot test 14
 - Load Test 18
 - View Planner Scoring Methodology 18
 - VMware View Planner 2.1 Workload Results 19
 - Functional IOPS 20
- Conclusion 21
- Reference Materials 22

Figures

Figure 1. Tested Architecture 8

Figure 2. Virtual Disk Relationship per Volume 10

Figure 3. VMware ESXi 5 Datastore Layout 11

Figure 4. Storage Tiers SSD 12

Figure 5. “VDI-SteadyState” Storage Profile 13

Figure 6. Replay creation 14

Figure 7. Replay statistics post workload run 15

Figure 8. Total IOPS during boot storm 17

Figure 9. SSD IOPS during boot storm 16

Figure 10. 15K IOPS during Boot Storm..... 17

Figure 11. IOPS with medium workload 19

Introduction

The intention of this document is to provide technical professionals with information to assist in the planning, design and deployment of a Dell™ Compellent™ Storage Center SAN in a VMware View 5 VDI environment using VMware ESXi 5 as the hypervisor. This document is based on internal testing of 1,000 virtual desktops on 16 Dell M710 HD physical blade servers. The data for this document was gathered and validated at the Dell Compellent lab and design center.

This document focuses on the impact of a virtual desktop infrastructure (VDI) on the Dell Compellent Storage Center SAN. As such, detailed information on host, network and other supporting infrastructure is out of the scope of this document. Host, network and supporting infrastructure is important to a VDI deployment and should also be carefully evaluated when planning a VDI project.

The information in this document is to serve as a reference to implementing VDI with the Dell Compellent Storage Center SAN. Workloads will be different for each organization, depending on applications and infrastructure. As such, it is strongly recommended that each organization run a VDI pilot to gather sizing and IOPS data as a basis for purchasing and provisioning the correct amount of storage.

Executive summary

Proper storage configuration is important to the success of a VDI project. Dell Fluid Data architecture and Dell Compellent integrated software features, such as automated tiered storage, thin provisioning, continuous snapshots and remote replication, provide a highly effective storage option for VMware View.

Dell Compellent Dynamic Capacity software separates storage allocation from utilization. While administrators can create storage volumes of any size for virtualized applications, physical capacity is consumed only when data is written to the disk. With thin provisioning, IT groups can avoid purchasing excess capacity upfront. When applications require more storage, the system provisions it from the pool of unused capacity. For instance, thin provisioning can present 10 virtual storage volumes, 10TB each, for use by 16 unique servers, but only require the physical capacity that is actually written for each host.

In addition, automated tiered storage from Dell Compellent, called Data Progression, can help organizations optimize the use of storage in a VMware View® environment with VMware ESXi 5. Data Progression can statically or dynamically move blocks of data between tiers of storage based on performance characteristics. This allows the intelligent use of all tiers of storage, without the write penalty typically associated with traditional storage arrays. Utilizing Data Progression, all writes into the storage array occur at Tier 1, RAID 10. Over time, the software dynamically moves those blocks to more efficient RAID levels, and the less frequently accessed blocks to lower tiers of storage. Data Progression reduces the need for large numbers of high-performance, high-cost disks by moving frequently used data to higher performance tiers of storage while moving infrequently used data to lower cost, higher-density disks. The migration of data is done on the block level, so data within a volume can be moved based on performance characteristics. These technologies help to create an efficient storage platform for VMware View 5 with VMware ESXi 5.

Key findings

The Dell Compellent Technical Solutions team tested a sample of 1,000 virtual desktops and simulated real-world workloads using VMware View Planner 2.1 with VMware ESXi 5 and VMware View 5. The environment included a Dell Compellent Storage Center SAN, Dell M1000e Blade enclosure, (16) M710HD blade servers, 10Gb iSCSI connectivity. VMware View Planner 2.1 uses standard office applications with periodic wait times to generate the simulated workload of a typical knowledge worker.

During the deployment of VMware View 5 virtual desktops, a selection of what datastore the replica disks, delta disks, and persistent data disks reside becomes imperative to a successful VDI implementation. Selecting different datastores for different types of VMware View disks allow a rich end-user experience. Utilizing a Dell Compellent Storage Center SAN and 6 SSD drives to create a volume that stores the VMware View replica disks images helps offload the heavy read IOPS of a boot and login storm. Offloading the heavy read IOPS to the SSD disks allows the higher capacity, lower IOPS 15k drives to handle the VMware View delta and persistent data disks.

Using VMware View Planner 2.1 and the full user workload, a peak of 25,000+ total IOPS and 400+ MBPS throughput was observed during View Planner boot-up initiation. Read/write percentage during boot-up sequence was 90% read and 10% write. After all 1,000 VDI desktops were booted and the workload commenced with a steady state of approximately 5,000 IOPS was observed. The read/write percentage during the steady state workload testing was 10% read, and 90% write.

In addition to the load test, a boot storm was emulated by powering on 1,000 virtual desktops simultaneously via VMware vCenter. This generated IOPS as high as 36,000. During the boot storm an average of 70% of IOPS came from the SSD tier of storage, while the remaining 30% came from the 15K disk tier. While measuring sustained high IOPS, the length of the 1,000 virtual desktop boot storm was approximately six minutes.

Tested architecture

The tested environment was comprised of 16 VMware ESXi servers running in two clusters with eight nodes per cluster. Each cluster was assigned to a single VMware View 5 Desktop Pool. Each VMware View desktop pool consisted of 500 VDI desktops. Eleven volumes were provisioned from the Dell Compellent Storage Center for each pool and corresponding VMware VMFS datastores were created on those volumes: Five VMware datastores stored the delta disks, five VMware datastores stored the persistent disks and one VMware datastore stored the replica disks.

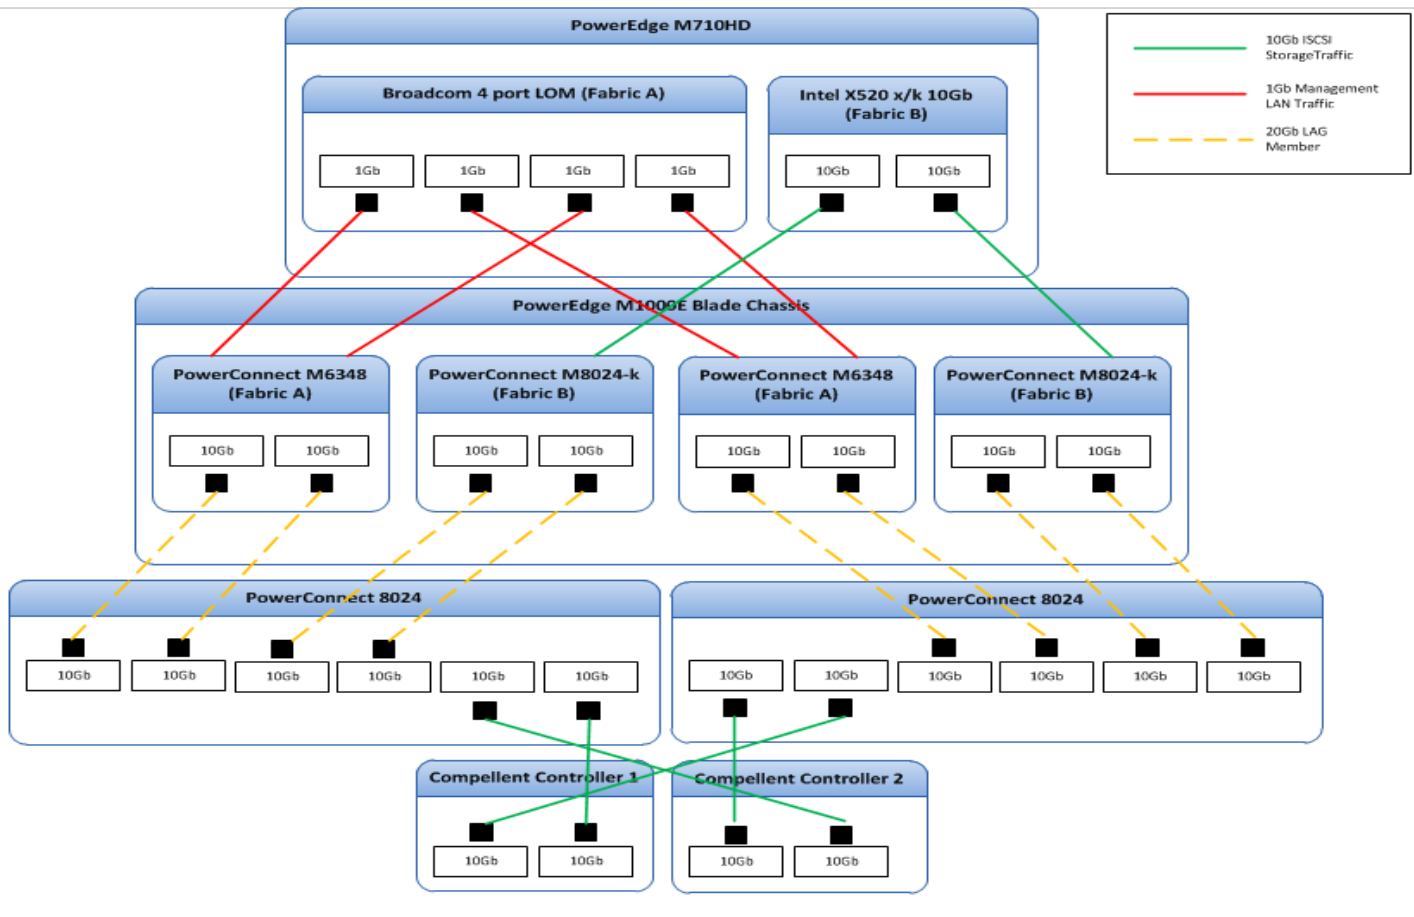
Powershell scripts were utilized to efficiently and accurately provision each volume in the Dell Compellent array, and present and create the VMware VMFS datastores. (*See appendix for example CSV answer file*) Each volume was 10TB in size; the disk size of 10TB was beyond the requirement for the test. However, with thin provisioning from the Dell Compellent Storage Center, no space was wasted on the SAN.

Additional infrastructure virtual machines were used to support the tested environment. This included VMs for infrastructure tasks, such as Active Directory, SQL Server, Dell Compellent Enterprise Manager, VMware vCenter, VMware View 5, and VMware View Planner 2.1.

Test infrastructure


- Dell M1000e Blade Chassis
 - (2) - Dell PowerConnect M6348 - Fabric A
 - (2) - Dell PowerConnect M8024-k - Fabric B
- (16) Dell M710HD blade servers as virtual hosts for VDI desktops
 - Dual Six Core Intel Xeon X5650 2.7GHz processor
 - 96GB 1333MHz DDR3 SDRAM
 - Quad-Port Broadcom 1-GbE Mezzanine Card (LAN) - Fabric A
 - Intel Ethernet X520 10 GbE x/k Mezzanine Card (iSCSI) - Fabric B
 - VMware ESXi 5 software iSCSI initiator
- (2) Dell PowerEdge R610 virtual hosts for infrastructure servers
 - Dual Quad Core Intel Xeon X5506 2.13GHz processor
 - 72GB RAM
 - Dual-port Intel 10GbE NIC, Model x520-DA2 for network and iSCSI
 - Virtualized infrastructure including:
 - 2, AD Domain controllers
 - 1, VMware vCenter 5 server
 - 1, VMware View 5 server
 - 1, VMware View Planner 2.1 Virtual Appliance
 - 1, Dell Compellent Enterprise Manager Data Collector server
 - 1, Windows Server 2008 R2 file server
 - 1, SQL Server 2008R2 server
- (2) Dell Compellent Storage Center Series 40 controllers
 - 48, 2.5-inch 146GB 15K SAS drives
 - 6, 2.5-inch 200GB SSD SAS drives
 - Firmware Version 5.5.3
 - 10GB iSCSI
 - Dell Compellent Fast Track-enabled

Figure 1. Tested architecture



Dell Compellent Fast Track

The desktop deployment utilized the Dell Compellent Storage Center SAN feature Fast Track, which dynamically places the most frequently used data on the fastest, or outer, tracks of each drive. By continually optimizing data placement on each disk, Fast Track eliminates wasted space on the fastest portion of the drive. In this test, Fast Track was used on every drive within the system.



Load generation

To accurately represent a real-world scenario, a load generation utility was used throughout the test. The load was generated using VMware View Planner 2.1 application. VMware View planner replicated a workload of 1,000 virtual desktops logging in and used a variety of office applications such as Microsoft Office, Internet Explorer, Mozilla Firefox, archiving software, and video playback. VMware View Planner 2.1 initialized a workload run that consisted of 8 total workload iterations. To pass this test, the 95th percentile of application latency must be below 1.5 seconds.

Virtual machine properties


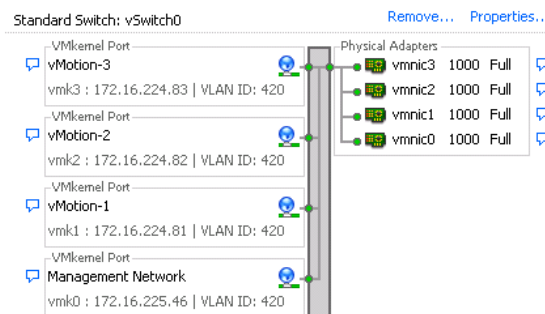
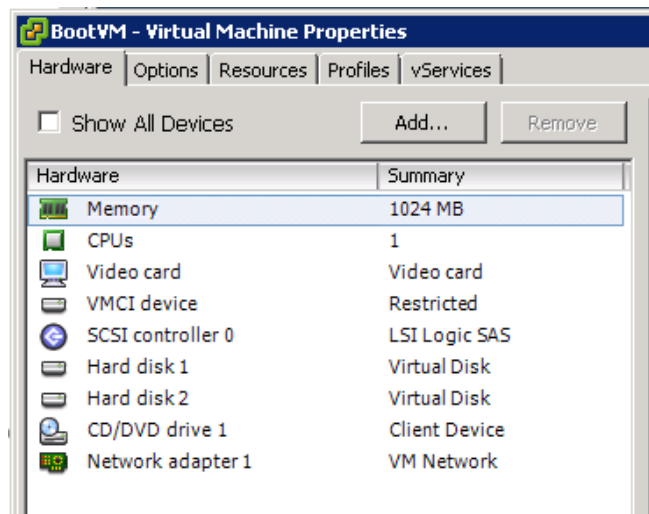
Following VMware View Planner 2.1 guidelines a generic "Base VM" was utilized to create all 1,000 VMware View Composer Linked Clone desktops. The template consisted of VMware Hardware v.8, 1 CPU, 1024mb of RAM, LSI Logic SAS Controller, 1 40GB Eager-Zeroed Hard Disk, 1GB Eager-Zeroed Hard Disk, CD-ROM Device and Network adapter (VMXNET 3). Microsoft Windows 7 32-bit was installed as the desktop Operating System.

VMware ESXi 5

VMware ESXi 5 Enterprise edition was utilized as the hypervisor for the guest virtual machines. ESXi storage connectivity was configured utilizing the VMware ESXi software iSCSI initiator and the dual port Intel 10Gb mezzanine cards in each Dell M710 blade server. (Figure.1)

vSwitch1 was assigned both Intel 10Gb ports. An "iSCSI-1" port group and "iSCSI-2" port group for each Intel 10Gb port was a created. Each port group had a dedicated physical NIC, IP address and VMkernel port. iSCSI-1 and iSCSI-2 port groups were then bound to the VMware ESXi software iSCSI adapter. All paths to the storage array were configured to utilize the native VMware NMP of Round Robin to distribute the load of the VDI desktops over all paths for all volumes. A single "VM Network" port group was added for virtual machine connectivity.

VMware ESXi management and vMotion networks utilized vSwitch0 and the 4 port Broadcom onboard NICs. NIC0 was active for management and NIC1, 2, 3 were set to standby in case of NIC0 failure. The vMotion network consisted of the remaining



3 NICs. NIC1, 2 and 3 all had separate port groups and VMkernel ports setup and dedicated for vMotion. By setting all 3 NICs, VMware ESXi 5 allows aggregation of the 1Gb NICs to provide maximum throughput for vMotion activities.

VMware View 5

To prepare for the desktop deployment, a single virtual machine (VM) was created as the base image for VMware View 5. This image was optimized and configured in accordance with the VMware View Planner 2.1 installation and user guide. Windows 7 SP1, Microsoft Office 2007 and other business applications required by VMware View Planner 2.1 were installed. The VMware View Desktop Agent was installed for management within Virtual Machine. Lastly, the VMware View Planner 2.1 setup package was installed to prepare the image for the VMware View Planner 2.1 Load Generation tool.

The base image then was shutdown and a VMware snapshot of the virtual machine was created. A VMware snapshot is a requirement of VMware View Composer 5 linked clones. A VMware View 5 automated resource pool (Pool1) with 500 VMware View Composer linked clone virtual machines was created. An identical automated pool of 500 virtual machines was created for (Pool2). The snapshot image was then automatically cloned to each specific replica disk for each pool. After the replica was copied to its respective VMware Datastore, the VMware View Composer linked clone virtual machine creation process began.

Figure 2. Virtual disk relationship per volume





Figure 3. VMware ESXi 5 datastore layout per pool



Windows 7

The virtual desktops platform for all tests was Windows 7 Professional 32-bit with SP1. Each VM was configured with one vCPU and 1024Mb of memory. Microsoft Office 2007 was installed, as well as the applications required for VMware View Planner. The template desktop was optimized using the best practices from the View Planner installation guide and VMware View and Windows 7 optimization guide.

Antivirus

No antivirus product was utilized within the VMs. Introducing traditional antivirus software may dramatically increase the I/O load of the VM. Special consideration should be taken when implementing antivirus programs in a virtual environment. Using virtualization-aware antivirus applications such as VMware vShield Endpoint or Trend Micro may reduce the I/O requirements within the virtual environment.

Storage Center configuration

The Dell Compellent Storage Center SAN was designed to maximize performance in a VDI environment. A pool of disks was created with two tiers of storage within the SAN. The first tier consisted of six 200GB SSDs, and the second tier consisted of forty-eight 146GB SAS drives. Making 1 SDD and 2 15k

drives hot spares left a total of 5 active 200Gb SSD drives and 46 active 15k 146Gb drives. Storage Center was running firmware version 5.5.3 during the test.

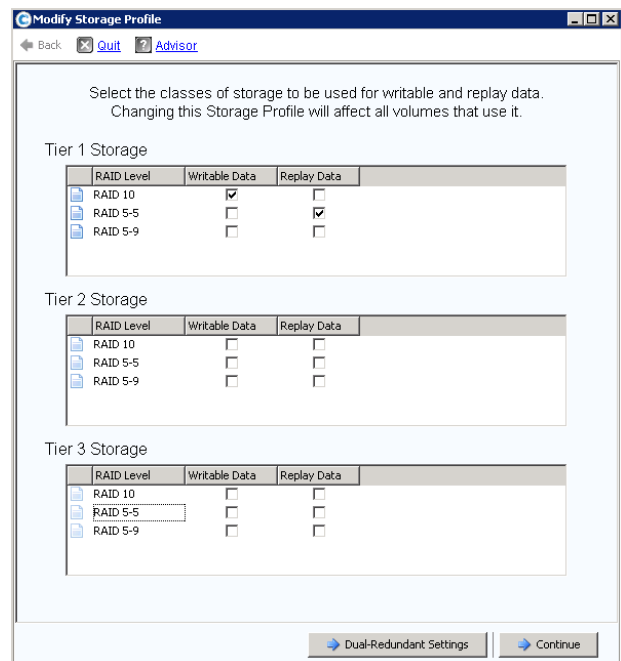
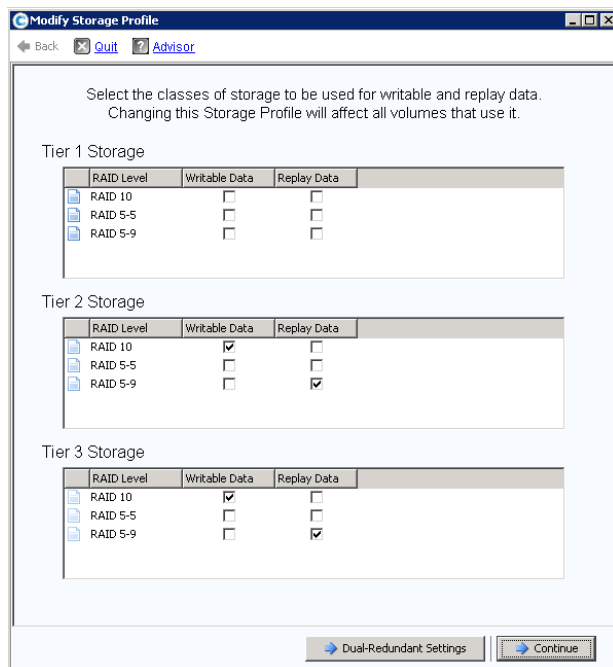
Figure 4. Storage tiers SSD



Two custom storage profiles were created to accommodate the performance characteristics of the two drive types used in testing. The first custom storage profile created was the “VDI-Replica” profile (Figure 5). It was applied to all replica volumes to “lock in” the replica images on the Tier1 SSD. The second custom storage profile that was created was “VDI-SteadyState” profile (Figure 5). This profile bypasses the SSD tier1 and directs all writes to the 15k tier of disks. (For details on creating custom storage profiles, see the Dell Compellent Administration Guide).

Figure 5. “VDI-SteadyState” storage profile

“VDI-Replica Storage” profile



Data Progression, replay technology

Data Progression is the technology that allows the Dell Compellent Storage Center to dynamically move 2Mb pages of storage between tiers of storage or between different RAID types among the same tier level. Data Instant Replays are the snapshot component of a Dell Compellent Storage Array. Data Instant Replays are unique because they utilize a technology of redirect at write instead of the traditional copy on write technology. At a very high level when a Replay is taken of a volume, all active 2Mb pages are marked as read-only. All pages that are marked as read-only will stay read-only until either a new write is directed at a specific page or the data in the specific page is deleted.

Data Progression is dependent upon replays to maximize storage efficiency. Data Progression, which is a daily low priority scheduled activity, analyzes all volumes and determines which pages have the read/-only attributes. Data Progression dynamically moves the read-only pages to a more efficient RAID level typically RAID 5-9 or RAID 6-10.

Data Progression is a powerful tool for VDI efficiency. Because VDI is traditionally write-intensive, the Dell Compellent Storage Array is a perfect fit. Dell Compellent is very efficient with writes because the default RAID type of all volumes is RAID 10. RAID 10 offers significant write performance because of a minimal write penalty. After a replay is taken of the volume, Data Progression will then move the newly created pages to RAID 5 or RAID 6. Because reads from RAID 5 or RAID 6 not incur a read penalty read performance from RAID 5 or RAID 6 are comparable to reading data from RAID 10.

A replay profile named VDI was created in the Dell Compellent Storage Center GUI. The replay profile schedule was setup to take a replay daily at 6 p.m. (See Figure 6) The VDI replay profile also expired after only 25 hours. Setting an expiration of 25 hours allows for maximum performance from Data Progression and allows for minimal disk space to be consumed. The replay profile was then associated to all VMware View Datastores. (See Figure 3)

Figure 6. Replay creation


Name VDI
Index 3
Type Replay Profile (Custom)
Replay Creation Serial

Schedule	Expiration
Daily at 6:00 PM	25 hours

Figure 7 shows one of the VMware View Delta volumes used for testing. Each Delta volume was utilized by approximately 100 VDI desktops. The chart shows the volume consumption statistics after a View Planner 2.1 workload test had finished. Within this volume, 54.07GB of data resided on the 15k Tier3 drives configured as RAID 5-9. The remaining data resided on RAID 10 with Fast Track (the outer 20% of the disk) or RAID 10 standard. Because of Dell Compellent ability to ingest all writes at RAID10, 33GB of

data was written throughout the day during the workload test on this single volume with 100 VMs running on it. After the View Planner workload completed, another replay was to be taken of this volume. Data Progression will move those RAID 10 pages onto RAID 5 and old expired RAID 5 pages will expire and be reallocated to the free page pool.

Figure 7. Replay statistics post workload run



Testing results

Testing for this reference architecture involves measuring two performance indicators. The first indicator is the amount of IOPS consumed during a boot storm. The second indicator measures the performance of all 1,000 desktops during simulated user load.

Boot test

The boot test represents a worst-case scenario where all 1,000 virtual desktops are powered on simultaneously. This situation would not take place during normal VMware View operation due to VMware View’s power management functionality. However, it is plausible that many desktops will need to be started quickly in a production environment, possibly when recovering from an outage. A boot storm also provides a good measurement of overall back-end storage performance.

In our testing, the length of the boot storm was measured by examining the duration of sustained high IOPS during a simultaneous power on of all 1,000 virtual desktops. In the configuration outlined above, with the replica image residing on SSDs, and all other read-write operation coming from 15k SAS drives, the average length of the boot storm was six minutes. During the boot storm, total front end IOPS ranged from 32,000 to 35,000. Seventy percent of all IOPS came from the read-only SSD tier, while the remaining 30% came from the 15K SAS tier. Read latency on the delta volumes spiked to 30ms. Write latency averaged 35ms during the same time. This high latency is common because of the limited

number of 15k drives required to meet the steady state requirements of the user workload. Because only 30% of the total IOPS required for the boot storm came from these 15k drives, the latency during the boot storm was acceptable. There was no measurable read latency on the SSD tier of storage.

The graphs below show an example of one of the boot tests performed with 1,000 virtual desktops. As the graphs show, the majority of reads came from SSD/replica disk during the boot storm. This is due to the read-only replica data on the master image residing on SSD disks. There is also an initial read spike on the 15K disk tier with the number of writes increasing as the systems are booted. After the boot storm, the number of writes surpasses the number of reads.

(Figure 8) shows the overall IOPS/sec during the 1,000 virtual desktop boot storm. The yellow line is overall IOPS/sec, the purple line is read IOPS/sec and black line is write IOPS/sec. Notice the dependency on reads during the boot storm process.

(Figure 9) shows the SSD disks IOPS/sec during the 1,000 virtual desktop boot storm. The yellow line is overall IOPS/sec, the purple line is read IOPS/sec and black line is write IOPS/sec. Notice that the SSD disks are providing approximately 70% of the IOPS/sec during the boot storm period. Also note that no writes were occurring at the SSD tier of disks.

(Figure 10) shows the 15k disks IOPS/sec during the 1,000 virtual desktop boot storm. The yellow line is overall IOPS/sec, the purple line is read IOPS/sec and black line is write IOPS/sec. Notice that the 15k disks are providing approximately 30% of the IOPS/sec during the boot storm period. Although 30% may seem like a small dependency, under sizing the 15k disks may present a bottleneck during boot and login storms. In this scenario the 15k disks do provide adequate IOPS/sec, but latency on the 15k disks was relatively high at 30-40ms during the boot storm.

Figure 8. Total IOPS during boot storm




Figure 9. SSD IOPS during boot storm





Figure 10. 15K IOPS during boot storm



Load test

The goal of the load test is to identify the impact of the user workload on the Dell Compellent Storage array. The load test was performed against the 1,000 desktops using VMware View Planner 2.1. VMware View Planner 2.1 was run utilizing the local mode option. The local mode operation starts 64 virtual desktops per minute until the required (1,000) virtual machines check-in with VMware View Planner. Once all 1,000 virtual desktops successfully check-in, the load test begins on the 1,000 virtual desktops.

View Planner scoring methodology

When the View Planner workload generator is run according to specific requirements, it produces a View Planner score. This score indicates how many users concurrently running a standardized set of operations a particular storage infrastructure can support.

While this standardized set of operations is representative of many real-world deployments, some users might have different needs. The custom applications feature allows these varying loads to be modeled and is designed to address this exact use case. However, because of the lack of standardization, runs using the custom applications feature may not be used to generate a compliant View Planner score.

View Planner workload run phases

During a workload run each desktop virtual machine goes through three phases:

- Ramp-Up (one iteration)
- Steady-State (at least three iterations)
- Ramp-Down (one iteration)

During an iteration of the user workload, View Planner reports the latencies for each operation performed within each virtual machine.

Requirements for compliant run

In order for a workload run to be considered compliant, it must meet the following requirements:

- It must run all of the standard View Planner applications and no other applications.
- The think time for View Planner must be set to 20 seconds.
- It must contain at least five iterations (including Ramp-Up and Ramp-Down).
- At least 95% of the response times during Steady-State must be 1.5 seconds or less.

View Planner Workload Configuration

- Users - 1,000
- Workload type - Local Mode
- Iterations - 8
- Ramp-up Time - 1 second
- Think Time - 20 seconds
- All VMware View Applications - MS Outlook, MS Excel, MS Word, MS PowerPoint, MS Internet Explorer, Adobe Acrobat, 7zip, Mozilla Firefox, Windows Media Player

VMware View Planner 2.1 workload results


The workload test was a local mode run with VMware View planner set to 8 iterations, a think time of 20 seconds, and a ramp up time of 1 second. All required VMware View Planner applications were tested during this workload. The 8 iterations that made up the workload took approximately 10 hours to complete.

The average IOPS/sec for the entire test including boot and login was approximately 6,450 IOPS. The read-write % during the workload was 10% read 90% write.

The chart below shows the throughput, read, write IOPS during the duration of the workload test. The initial IOPS burst symbolized VMware View Planner starting all 1,000 virtual desktops. VMware View planner steady-state workload begins at the approximately 16:00 mark.

VMware View Planner running in local mode for 1,000 desktops reported a successful passing score. The 95th percentile response time for the test was only .94 seconds. VMware View Planner 2.1 requires the 95th percentile application response time be at or below 1.5 seconds.

Figure 11. IOPS with medium workload



Functional IOPS

Functional IOPS are important with VDI due to the high write requirement of virtual desktops. Some estimates put the number of writes in a VDI environment up to 90% of all IOPS. As shown in the example above, the steady-state IOPS read-write percentage was 10/90. This can affect the number of IOPS available due to the write penalties associated with different RAID types. Write penalties are assigned to RAID levels based on characteristics of their write activity. For example, with RAID 1 or 10, data is written once and read once to verify each write. This consumes two raw IOPS, so RAID 1 and 10 are assigned a write penalty of two. Likewise, RAID 5 with 5 drives has a penalty of 5. The write penalty increases with the number of drives in the RAID set.

In a planning phase of our workload testing we ran a single virtual machine through the VMware View Planner workload. An average of 4.5 IOPS/sec was observed while running the VMware View Planner workload. $4.5 \text{ IOPS} \times 1000 \text{ VM's} = 4500 \text{ IOPS}$.

Using the functional IOPS calculation, a total of 48 drives was determined to provide adequate disk IOPS.

To calculate functional IOPS, the number of raw IOPS are factored with the percent of reads and writes along with the RAID penalty. The formula for calculating functional IOPS is expressed as:

$$\text{Functional IOPS} = ((\text{Raw Storage IOPS} \times \text{Write \%}) / \text{RAID Write Penalty}) + (\text{Raw Storage IOPS} \times \text{Read \%})$$

In the example below we will use the formula to figure out the needed disks for our VMware View Planner 2.1 workload. Forty eight 146GB 15K drives in a RAID 10 configuration with 90% write to 10% read ratio will provide 4752 functional IOPS.

$$\text{Functional IOPS} = ((8640 \times 90\%) / 2) + (8640 \times 10\%)$$

In comparison, if the above were changed to RAID 5 instead of RAID 10, the functional IOPS would be 2419.

$$\text{Functional IOPS} = ((8640 \times 90\%) / 5) + (8640 \times 10\%)$$

For more information on functional IOPS in a VDI environment, see the Dell Compellent document "Functional IOPS in VDI with Dell Compellent."

Conclusion

Based on the test results, the storage requirements for 1,000 desktops should be built to support 4,500 IOPS. However, to accurately size a storage system the percentage of write IOPS and the RAID penalty for those writes must be factored in. Adjusting for RAID penalty will give the functional IOPS required when provisioning data storage.

During the load test, averages of 10% of the IOPS were read; while the remaining 90% were write. With the configuration outlined in this document, all write IOPS occur at RAID 10. This is the default for all writes on a Dell Compellent Storage Center SAN. RAID 10 has a RAID penalty of two, meaning it takes two raw IOPS to perform a single write. Factoring the RAID penalty in the calculation below, we can conclude that it will take 8,640 raw IOPS to support 1,000 desktops.

$$\text{Functional IOPS} = ((8640 * 90\%) * 2) + (8640 * 10\%)$$

To meet the read-write requirements using 15K SAS drives, the test team estimates IOPS per drive to be 180. Given 180 IOPS per 15K SAS drive, it would take 48 drives to support a 1,000 user environment during normal operation.

In some environments, there may be a requirement to support aggressive boot storms, such as what may be experienced at the beginning of the workday while users log in. In our tests, the IOPS requirement increases up to 36,000 during a boot storm with the majority of IOPS being read. At 180 IOPS, it would take 200 15K drives to support this number of IOPS during a boot storm. The testing team was able to reduce the number of drives to support a boot storm by placing the replica image onto SSDs during the tests. The number of drives can be determined by total capacity of required replica images, as the SSDs can handle nearly 25,000 read IOPS per drive.

Latency is also an important consideration in a VDI environment. During the load tests, there was no measurable read latency on the SSD drives. The baseline write latency for the 15K disk pool was 3ms, and the read latency baseline was 8ms for the same pool of 15k disks.

The preceding tests used VMware View 5 with VMware View Composer Linked Clone desktops. These desktops are not persistent and typically rely on profile management strategy to deliver user-specific data to the desktop. As such, we did not factor in space requirements for user data in our testing. However, this should be considered when implementing a VDI environment.

Reference materials

The Dell Compellent Administration Guide is available on the Dell Compellent Customer Portal

VMware

<http://www.vmware.com>

VMware View 5 Documentation

http://www.vmware.com/support/pubs/view_pubs.html

VMware View Optimization Guide for Windows 7

<http://www.vmware.com/resources/techresources/10157>

Functional IOPS in VDI with Dell Compellent

http://www.compellent.com/~media/www/Files/White_Papers/Functional_IOPS_in_VDI.ashx

Dell Compellent Storage Center Best Practices Guide with VMware vSphere 5.x

<http://kcint.compellent.com/Published%20Documents/Dell%20Compellent%20Best%20Practices%20with%20VMware%20vSphere%205.x.pdf>

Dell Compellent Storage Center Users Guide

<http://kcint.compellent.com/Published%20Documents/680-019-009.pdf>

Appendix A

Example Powershell Answer file

VolName	StorageType	ParentFolder	StorageProfileName	ReplayProfile	Size	ActiveController	LUN	Cluster
VDI-P1-User-1	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	14	201	VDI-Cluster1
VDI-P1-User-2	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	14	202	VDI-Cluster1
VDI-P1-User-3	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	14	203	VDI-Cluster1
VDI-P1-User-4	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	14	204	VDI-Cluster1
VDI-P1-User-5	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	14	205	VDI-Cluster1
VDI-P2-User-1	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	15	206	VDI-Cluster2
VDI-P2-User-2	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	15	207	VDI-Cluster2
VDI-P2-User-3	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	15	208	VDI-Cluster2
VDI-P2-User-4	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	15	209	VDI-Cluster2
VDI-P2-User-5	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	15	210	VDI-Cluster2
VDI-P1-OS-1	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	14	211	VDI-Cluster1
VDI-P1-OS-2	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	14	212	VDI-Cluster1
VDI-P1-OS-3	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	14	213	VDI-Cluster1
VDI-P1-OS-4	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	14	214	VDI-Cluster1
VDI-P1-OS-5	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	14	215	VDI-Cluster1
VDI-P2-OS-1	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	15	216	VDI-Cluster2
VDI-P2-OS-2	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	15	217	VDI-Cluster2
VDI-P2-OS-3	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	15	218	VDI-Cluster2
VDI-P2-OS-4	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	15	219	VDI-Cluster2
VDI-P2-OS-5	Redundant - 2MB	VDI\DSKTPS\	VDI-SteadyState	VDI	10240g	15	220	VDI-Cluster2
VDI-P1-Replica-1	Redundant - 2MB	VDI\DSKTPS\	VDI-Replica	VDI	10240g	14	221	VDI-Cluster1
VDI-P2-Replica-2	Redundant - 2MB	VDI\DSKTPS\	VDI-Replica	VDI	10240g	15	222	VDI-Cluster2

Test System Component Details

Hardware	Version	Quantity
DELL M1000e Blade Enclosure	3.21	1
DELL M8024-k Blade I/O Module	4.1.0.6	2
DELL M8038 Blade I/O Module	3.1.3.9	2
DELL M710HD Blade Server	2.0.0	16
DELL Compellent SC-040 Controller	5.5.3	2
DELL Compellent 24 bay 2.5 inch Enclosure	N/A	3
DELL Compellent 15k 146Gb 6Gb SAS Hard Drives	N/A	48
DELL Compellent SSD 200Gb 6Gb SAS Hard Drives	N/A	6
DELL 8024 10gb Ethernet Switch	3.1.4.5	2