

Dell Compellent Storage Center SAN & Citrix XenDesktop 1,000 Desktop Reference Architecture

Dell Compellent Technical Solutions

This document has been archived and will no longer be maintained or updated. For more information go to the [Storage Solutions Technical Documents page on Dell TechCenter](#) or contact support.



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.

© 2011 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. Citrix, XenServer and XenDesktop are trademarks of Citrix Systems Inc. Microsoft and Hyper-V 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.

August 2011

Contents

Figures..... 4

Introduction 5

Executive Summary 5

Key Findings 6

Tested Architecture..... 6

 Fast Track..... 8

 Load Generation..... 8

 System Center Virtual Machine Manager 9

 XenDesktop 5..... 9

 Windows 710

 Antivirus.....10

 Storage Center Configuration10

XenDesktop “Master” Image Management10

 Machine Create Service and Hyper-V Configuration11

 Data Progression and Replay technology and the Master Image.....12

Testing Results14

 Boot Test.....14

 Load Test16

 Medium Workload Results.....17

 Heavy Workload Results.....18

Dynamic Memory Consideration.....19

Functional IOPS20

Conclusion21

Reference Materials.....22

Figures

Figure 1. Tested Architecture..... 8

Figure 2. Virtual Disk Relationship per Volume..... 9

Figure 3. Storage Tiers 11

Figure 4. Custom Storage Profile..... 11

Figure 5. CSV Tiers..... 13

Figure 6. Replay Statistics 13

Figure 7. SSD IOPS During Boot Storm..... 15

Figure 8. 15K IOPS During Boot Storm..... 16

Figure 9. IOPS with Medium Workload 17

Figure 10. KBPS with Medium Workload..... 17

Figure 11. IOPS with Heavy Workload..... 18

Figure 12. KBPS with Heavy 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 Citrix® XenDesktop® VDI environment using Microsoft® Hyper-V® 2008 R2 as the hypervisor. This document is based on internal testing of 1,000 virtual desktops on 12 physical 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. A 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 Citrix XenDesktop and Hyper-V.

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, 1TB each, for use by 10 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 Citrix XenServer® environment with Microsoft Hyper-V. 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 Citrix XenDesktop with Microsoft Hyper-V.

Key findings

The Dell Compellent Technical Solutions team tested a sample of 1,000 virtual desktops and simulated real-world workloads using Citrix XenDesktop with Microsoft Hyper-V. The environment included a Dell Compellent Storage Center SAN, 10GB iSCSI connectivity and a Login Consultants Login VSI scripted load-generation solution. Login VSI uses standard office applications with periodic wait times to generate the workload of a typical knowledge worker. Two load tests were performed, one with a medium workload and one with a heavy workload emulating 1,000 power users.

During the deployment of XenDesktop virtual desktops, an instance of the template VM is copied to the volume where the virtual desktops will reside. This is a shared, read-only copy of the template that new virtual machine differencing and identity information will reference. Using Dell Compellent Data Instant Replay and Data Progression, this read-only portion of the volume was moved to the SSD tier of storage while the read-write virtual machine differencing and identity information stayed on the 15K tier of storage.

Using a medium workload generated a peak of 6,800+ total IOPS and 125,000+ KBPS, while the heavy workload generated a peak of 7,500+ IOPS and 160,000+ KBPS. During the medium test, average read IOPS were 37 percent and write IOPS were 63 percent. The heavy workload generated 33 percent read and 67 percent write IOPS on average.

In addition to the load test, a boot storm was emulated by powering on 1,000 virtual desktops simultaneously. This generated IOPS as high as 40,000 with a data rate of over 700 MBPS. During the boot storm an average of 70 percent of IOPS came from the SSD tier of storage, while the remaining 30 percent came from the 15K disk tier. The length of the boot storm while measuring sustained high IOPS was six minutes.

Tested architecture

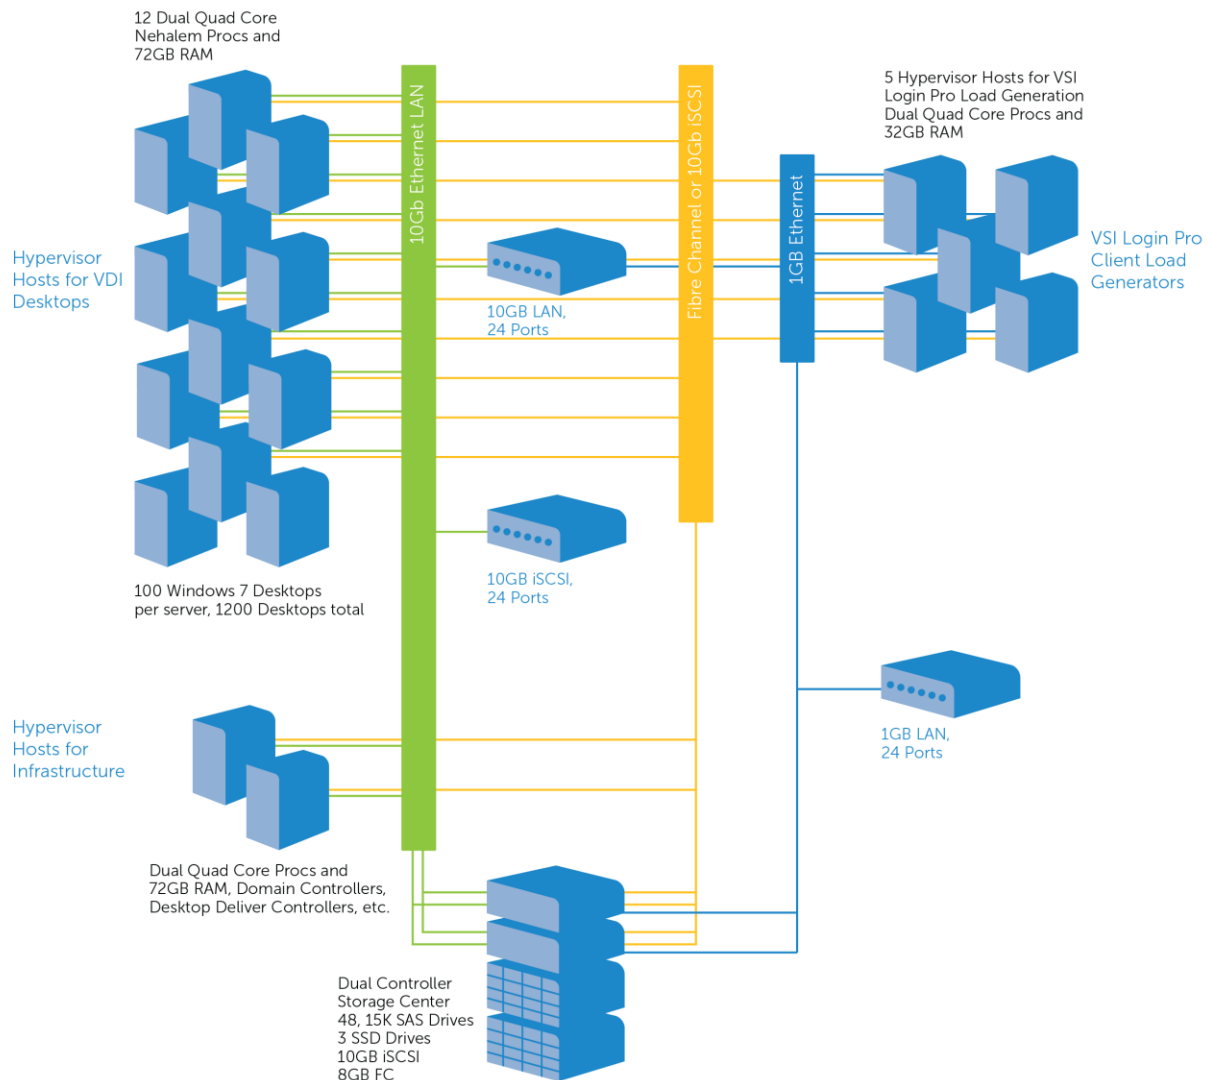
The tested environment was comprised of 12 Hyper-V servers running Windows Server 2008 R2 Enterprise edition in two clusters with six nodes per cluster. Each cluster was assigned six, 2TB Cluster Shared Volume (CSV) volumes for hosting the virtual desktops. Six CSVs were created to increase the number of command queues Microsoft Windows has to access the virtual hard disks (VHDs). The disk size of 2TB was beyond the requirement for the test. However, with thin provisioning little space was wasted on the SAN.

An additional seven virtual host servers were used to support the tested environment. This included two servers for infrastructure tasks, such as Active Directory, SQL Server, System Center Virtual Machine Manager (SCVMM) and XenDesktop. The remaining five were used for load generation. Following are the specifics on the hardware used.

- 12 Windows Server 2008 R2 Enterprise servers as virtual hosts for tested environment
 - Dual Quad Core Intel Xeon X5550 2.6GHz processor
 - 72GB RAM
 - Dual-port Intel 10GbE NIC, Model x520-DA2 for network and iSCSI
 - Microsoft software iSCSI initiator

- All recent updates and hot fix KB2263829
- 2 Windows Server 2008 R2 Enterprise servers as virtual hosts for infrastructure
 - 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, SCVMM server
 - 1, XenDesktop 5 server
 - 1, Dell Compellent Enterprise Manager Data Collector server
 - 1, Windows Server 2008 R2 file server
 - 1, SQL Server 2008 server
- 5 Windows Server 2008 R2 Enterprise servers as virtual hosts for load generation
 - Dual Quad Core Intel Xeon X5506 2.13GHz processor
 - 32GB RAM
 - Dual Broadcom 1GB NIC
 - 5 virtual servers per host used for load generation
- 2 Dell Compellent Storage Center SAN Series 40 controllers
 - 48, 2.5-inch 146GB 15K SAS drives
 - 3, 2.5-inch 200GB SSD SAS drives
 - Firmware Version 5.5.2
 - 10GB iSCSI
 - Dell Compellent Fast Track-enabled

Figure 1. Tested Architecture



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 Login Consultants Login VSI application. One test was run with a load set to medium and a second with a setting to heavy. With Login VSI, the tests replicated a workload of all 1,000 virtual desktops logging in and used a variety of office applications such as Microsoft Outlook, Office and Internet Explorer, as well as PDF creation and Flash movies.

System Center Virtual Machine Manager

System Center Virtual Machine Manager 2010 R2 (SCVMM) was used to manage the hosts and all virtual desktops. SCVMM also provides the interface from XenDesktop and Machine Creation services to deploy and manage desktops within the clusters.


XenDesktop 5

To prepare for the desktop deployment, a single virtual machine (VM) was created on an infrastructure server as the base image for XenDesktop. This image was configured with Windows 7 SP1, Microsoft Office and other business applications. The Virtual XenDesktop Desktop Agent was installed for management within XenDesktop. Lastly, the Login VSI Target installer was initiated to prepare the image for the Login VSI Load Generation.

Once the base image was ready, it was cloned to each cluster using SCVMM. Cloning the base image to each cluster makes it available to XenDesktop to use as a master image. Machine Creation Services copied this master image to each volume on the cluster when the new machine catalog was created. XenDesktop then used the master disk and a unique identity and differencing disk to create each VM. The configuration to deploy the virtual desktops is as follows:

- Desktops created with Machine Creation Services
- Machine type pooled, randomly assigned
- Two host groups, one for each cluster
- Two catalogs of 500 virtual desktops, one for each cluster
- One desktop group with all 1,000 virtual desktops

Figure 2. Virtual Disk Relationship per Volume



Windows 7

The virtual desktops platform for all tests was Windows 7 Professional 32-bit. Each VM was configured with one vCPU and 768MB of memory. The Windows page file was set to twice the size of available RAM. Microsoft Office 2007 was installed, as well as the applications required for Login VSI. The template desktop was optimized using the Citrix XenConvert Optimizer application.

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 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 three 200GB SSDs, and the second tier consisted of forty-eight 146GB SAS drives. Storage Center was running firmware version 5.5.2 during the test.

Figure 3. Storage Tiers



A custom storage profile was created to accommodate the performance characteristics of the two drive types used in testing. The storage profile definition allowed writes to come in at RAID 10 on Tier 2, while allowing Replay (snapshot) data to migrate up to the Tier 1 SSD level. (For details on creating custom storage profiles, see the Dell Compellent Administration Guide).

XenDesktop “Master” Image Management

One objective of the testing was to split read-only data that is shared between all virtual desktops from the read-write data that is unique to each virtual desktop. By splitting this data, the Dell Compellent Storage Center SAN can better utilize the performance characteristics of both the SSD and

SAS drives within the system. Master image management allows read-only data to reside on the SSD tier, while read-write transactions take place on the 15K SAS tier of storage.

Figure 4. Custom Storage Profile

Select the classes of storage to be used for writable and replay data.

Tier 1 Storage

RAID Level	Writable Data	Replay Data
RAID 10	<input type="checkbox"/>	<input type="checkbox"/>
RAID 5-5	<input type="checkbox"/>	<input checked="" type="checkbox"/>
RAID 5-9	<input type="checkbox"/>	<input type="checkbox"/>

Tier 2 Storage

RAID Level	Writable Data	Replay Data
RAID 10	<input checked="" type="checkbox"/>	<input type="checkbox"/>
RAID 5-5	<input type="checkbox"/>	<input type="checkbox"/>
RAID 5-9	<input type="checkbox"/>	<input type="checkbox"/>

Tier 3 Storage

RAID Level	Writable ... ▾	Replay Data
RAID 5-5	<input type="checkbox"/>	<input type="checkbox"/>
RAID 5-9	<input type="checkbox"/>	<input type="checkbox"/>
RAID 10	<input type="checkbox"/>	<input type="checkbox"/>

Machine Create Service and Hyper-V Configuration

Virtual desktops were created using XenDesktop Machine Creation Services (MCS) on Hyper-V 2008 R2 with Cluster Shared Volumes (CSV). As mentioned previously, MCS copies the “master” or base image onto each CSV during virtual desktop creation. MCS then creates a differencing and identity disk that each VM uses to store unique data. The master image is a read-only copy of the virtual desktop template used in conjunction with the differencing and identity disk to create the complete virtual desktop image. (See Figure 2).

The space that the VMs consume must be considered with the configuration used during testing because of the affect it has on the SSD tier of storage. After initial deployment of the virtual desktops and before they are powered on, the differencing and identity data consume approximately 16MB of data. With a base image of 8.4GB and the differencing and identity information from 82 to 84 virtual desktops on each CSV, the total amount of space consumed on a CSV before any virtual desktops are turned on is approximately 10GB. With 12 CSVs in the testing infrastructure, total space consumed before the virtual desktops were started was approximately 120GB.

Once the virtual desktops are powered on and users have logged in, the differencing and identity data will grow. In our testing, the space used by the differencing and identity data after a Login VSI load test was approximately 2.2GB. Total amount of space used on a CSV after the load test was approximately 185GB. The total space consumed for all CSVs was 2.3TB.

* All drive space figures are approximate and will change depending on the environment. For best results, use a pilot program to determine drive space needs within your environment.

Data Progression, Replay Technology and the Master Image

Natively with XenDesktop MCS, it is not possible to split the read-only and read-write data on a volume basis. This is because the shared master image is stored with the identity and differencing data on each volume. However, in the tested infrastructure, Dell Compellent Replay technology and Data Progression were used to overcome this and move the read-only master image data to SSDs while directing all virtual desktop specific read-writes to the 15K SAS drives. The steps involved to split the master image reads from virtual desktop specific read-writes were:

- Create disk pools and storage profile as outlined above.
- Create new volumes, attach them to the cluster and configure as CSV.
- Create the virtual desktops in XenDesktop, placing the virtual desktops on the new CSVs. Do not power on the virtual machines.
- Run configuration scripts to set VLAN or startup configuration.
- Create a Manual Replay set to Never Expire on the new CSVs.
- Continue configuration specific to your environment, starting the virtual desktops as necessary.

Following the steps above, the disk pools and storage profile were configured as outlined in the Storage Center Configuration section. This set all writes to come in as RAID 10 on Tier 2, while moving Replay data to RAID 5 on Tier 1. Next, new CSVs for the virtual desktops were created. (It is important to create a new volume to guarantee that no data resides on the volumes. When finalized, the new volumes should be mapped and configured as CSVs on the hypervisor).

The virtual desktops were then deployed using XenDesktop. As noted above, at the point where the desktops have been deployed but before they are powered on, they consume a limited amount of data. In the test environment, this included 8.4GB for the master image and 16MB of space for differencing and identity information per virtual desktop. This data is primarily read-only and will be locked using Dell Compellent Replay technology and migrated to the SSD tier using Data Progression. Any configuration scripts, such as setting VLAN tagging or modifying power-on settings, can be run at this point, providing they don't power on the virtual desktops. Also, be cautious of assigning the virtual desktops to a desktop pool, as XenDesktop may power on the virtual desktop automatically.

After the virtual desktops are deployed, a manual Replay set to Never Expire is created on the volumes. This will lock all blocks on the volume and allow them to migrate to Tier 1 when Data Progression runs. By locking these blocks before the virtual desktops are powered on, the read-only data (the master image) can be split from the read-write data (differencing and identity). The small amount of identity and differencing data locked in the Replay will not have an effect on performance of the virtual desktops. Figure 5 illustrates how the data resides on each tier of storage within a volume.

Figure 5. CSV Tiers





Figure 6 shows one of the CSV volumes used for testing. Within this volume, 9.2GB of data was moved to Tier 1 on the SSD drives configured as RAID 5. The rest of the data resided on the lower tier as RAID 10 with Fast Track (the outer 20 percent of the disk) or RAID 10 standard. Total data on the volume, not including RAID overhead, was 111GB.

Figure 6. Replay Statistics



Using Dell Compellent Replay and Data Progression technology, each volume is divided in a way that allows read-intensive data (in this case, the master image) to reside on the SSD tier, while read-write (differencing and identity) data stays on the SAS tier. This leads to cost savings because the SSDs are only required for read-intensive data within the volume. With this configuration there is no need to

size SSDs for the entire VDI infrastructure, only the portion of the data that will benefit from the high read IOPS of the SSDs.

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 XenDesktop operation due to XenDesktop'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 master image residing on SSDs, and all other read-write operation coming in on SAS drives, the average length of the boot storm was six minutes. During the boot storm, total IOPS ranged from 35,500 to 39,500. Seventy percent of all IOPS came from the read-only SSD tier, while the remaining 30 percent came from the 15K SAS tier. Read latency on the 15K SAS tier spiked to 9ms with an average of 4.5ms. Write latency averaged 6ms during the same time. 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 SSDs during the boot. This is due to the OS data on the master image residing on Tier 1. There is also an initial read spike on the 15K SSD 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 7. SSD IOPS During Boot Storm





Figure 8. 15K IOPS During Boot Storm



Load Test

The goal of the load test is to identify the impact of the load on the Dell Compellent SAN. Two load tests were performed against the 1,000 desktops using Login Consultants Login VSI Pro. One test was conducted with the load set to medium with flash, and one ran with a load set to heavy. (For details, see the Login Consultants Wiki.)

Medium Workload with Flash

- Emulates a knowledge worker using Office, Internet Explorer and a PDF
- Sessions repeat every 12 minutes
- Response time measured every 2 minutes
- Opens 5 applications simultaneously
- Type rate is set to 160ms for each character
- Approximately 2 minutes of idle time included to simulate real-world users

Heavy Workload

- Based on medium workload
- More memory and CPU-intensive
- Simulates a power user

- Type rate is 130ms per character
- Idle time is only 40 seconds
- Opens 8 applications simultaneously


With both the medium and heavy workloads the Login VSI application was configured to run in parallel mode. Parallel mode evenly distributes the sessions across all the Login VSI clients that generate the login activity. The test was set to launch client logons over 10,000 seconds. In this configuration the Login VSI client launched one remote connection every 10 seconds. The test continued to run for one hour after the last session logged in, creating a sustained workload of 1,000 sessions for one hour.

Medium Workload Results

This test was run with Login VSI set to a medium workload with flash enabled. The average IOPS for the entire test was 4,817. This includes the virtual desktop login and the hour of sustained workload. The average of the 12 highest data points representing the most active hour of testing was 6,088 IOPS. The average IOPS during the last hour of testing when all virtual desktops were logged in and running was 5,363.


The chart below shows the read, write and total IOPS during the duration of the medium workload test. The first portion of the test leading up to the highest data point represents the login of the 1,000 virtual desktops. The remaining portion of the chart shows the final hour of sustained workload.

Figure 9. IOPS with Medium Workload



The next chart shows the data throughput during the test represented by kilobits per second. As with the previous chart, the data leading up to the highest data point represents the VM login of one virtual desktop every 10 seconds. The remaining portion represents the final hour of sustained workload. During the last hour of the test with all machines logged in and the load running, 16 percent of all IOPS were read and the remaining 84 percent were write.

Figure 10. KBPS with Medium Workload




Heavy Workload Results

Running a heavy workload created an average IOPS of 5,281 during the duration of the test. This includes the virtual desktop login and the hour of sustained workload. The average of the highest 12 data points during the test was 6,652. The average IOPS produced during the last hour of testing, the point after virtual desktop login while all sessions were running, was 6,118.


The chart below shows the read, write and total IOPS during the heavy workload test. The first portion of the test leading up to the highest data point represents the login of the 1,000 VMs. The remaining part of the chart shows the hour of sustained workload.

Figure 11. IOPS with Heavy Workload



The next chart shows the data throughput during the test represented by kilobit per second. As with the previous chart, the data leading up to the highest data point represents the VM login of one virtual desktop every 10 seconds. The remaining portion represents the final hour of sustained workload. During the last hour of testing while all machines were logged in and the load running, 19 percent of the IOPS were read while the remaining 83 percent were write.

Figure 12. KBPS with Heavy Workload



Dynamic Memory Consideration

When XenDesktop creates virtual desktops in Hyper-V, memory is set to a static amount. This ties the number of virtual desktops to the amount of RAM on the host despite actual memory requirements. Dynamic Memory can improve virtual desktop density by assigning RAM to each guest according to its need. Based on documents published by Dell and Microsoft, it is possible to increase density up to 40 percent per server using Dynamic Memory.

The tests in this document were conducted with memory statically set to 768MB on each virtual desktop. Each host had a total of 72GB of RAM. The total amount of available RAM on the host was 3GB with all 84 virtual desktops started. With 3GB of free space it would have been possible to add three more desktops on each host, leaving some room for additional host overhead. This would have resulted in a total of 88 virtual desktops per host, or a total of 1,056 desktops over the 12 hosts.

To increase the density in the test environment, Dynamic Memory was enabled on all 1,000 virtual desktops. Starting RAM was set to 384MB, and maximum usable memory was set to 1GB. All virtual desktops were powered on, and Login VSI was used to place the virtual desktops under a medium workload. During the medium workload, the amount of assigned memory ranged from 512MB to 670MB. The average amount of RAM used by the host was 56GB, or 80 percent of available RAM.

The hosts had approximately 14.5GB of RAM available with Dynamic Memory enabled when placed under a medium load. Estimating memory usage to be 600MB average under a load and leaving some space for additional overhead, an additional 24 virtual desktops were added per host. This brought the total number of desktops per the 12 hosts to 1,296. This configuration allowed for the running of 240 more desktops than with static memory, a 23 percent gain over the number of virtual desktops supported using static memory.

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 percent of all IOPS. 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 4 drives has a penalty of 4 because each write operation requires four raw IOPS, one to read the existing data, one to read existing parity, one to write new data, and one to write new parity. The write penalty increases with the number of drives in the RAID set.

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:

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

In the example below, four 146GB 15K drives in a RAID 5 configuration with 80% write to 20% read ratio will provide 288 functional IOPS.

$$\text{Functional IOPS} = ((720 * 80\%) / 4) + (720 * 20\%)$$

In comparison, if the above were changed to RAID 10 instead of RAID 5, the functional IOPS would be 648.

$$\text{Functional IOPS} = ((720 * 80\%) / 2) + (720 * 20\%)$$

For more information on functional IOPS in a XenDesktop 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 7,000 IOPS. However, to accurately size a storage system the percent 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 two load tests, an average of 30 percent of the IOPS were read, while the remaining 70 percent 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 11,900 raw IOPS to support 1,000 desktops.

$$\text{Functional IOPS} = ((7000 \times 70\%) \times 2) + (7000 \times 30\%)$$

To meet the read-write requirements using 15K SAS drives, the test team estimates IOPS per drive to be 250. Given 250 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 40,000 during a boot storm with the majority of IOPS being read. At 250 IOPS, it would take 160 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 master image onto SSDs during the tests. The number of drives can be determined by space requirements, as the SSDs can handle nearly 30,000 read IOPS.

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 7ms, and the read latency baseline was 5ms for the same pool of disks.

The preceding tests used XenDesktop with Pooled-Random desktops. These desktops are not persistent and typically rely on Citrix Profile Management or other 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

Citrix

<http://www.citrix.com>

Dell Microsoft Dynamic Memory Solution Overview

<http://www.microsoft.com/casestudies/Windows-Server-2008-R2-Datacenter/Dell/Dell-Boosts-Virtualization-Density-for-Customers-by-40-Percent-with-New-Software/4000009858>

Desktop Virtualization Top 10 Mistakes

<http://support.citrix.com/servlet/KbServlet/download/24559-102-647700/XD%20-%20Top%2010%20Mistakes%20Identified%20When%20Doing%20Desktop%20Virtualization.pdf>

Hyper-V Network Connectivity hot fix

<http://support.microsoft.com/kb/2263829>

Login Consultants

<http://www.loginconsultants.com/>

Login Consultants WIKI

<http://vsi.wikispaces.com/>

Windows 7 Optimization

<http://blogs.citrix.com/2010/01/15/optimizing-windows-7-for-flexcast-delivery/>

Functional IOPS in VDI with Dell Compellent

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