Deploying Microsoft RDS with Storage Spaces and Dell PowerVault MD1220 storage arrays

Bala Chandrasekaran
Dell Client Cloud Computing Engineering
May 2014

A Dell Technical White Paper

Revisions

Date	Description
May 2014	Initial release

Table of contents

Executive summary

Microsoft Storage Spaces enables administrators to create a cost efficient and highly scalable storage solution leveraging industry standard servers, storage controllers, and hard drives. This white paper includes a reference architecture for deploying Microsoft Remote Desktop Services with Storage Spaces and Dell PowerVault MD1220 storage arrays. Best practices for array management, multi-pathing, data tiering, data deduplication, and caching are also provided. Finally, two disk configurations are evaluated, as part of the reference architecture:

- SAS Storage Configuration (single tier 24 SAS drives)
- Tiered Storage Configuration (4 SSDs + 20 SAS drives)

Key findings from the evaluation show that:

- For low cost and smaller deployments, the SAS Storage Configuration provides a good starting point that supports up to 150 enhanced desktop users
- For highly scalable and low latency deployments, Tiered Storage Configuration is recommended that supports up to 700 enhanced pooled (shared OS images) desktop users
- With data deduplication, Tiered Storage Configuration can support up to 700 personal (dedicated OS images) desktops

1 Introduction

Microsoft Storage Spaces is a new capability available in Windows Server 2012. Storage Spaces provides resilient, highly scalable and highly available storage system utilizing cost-efficient commodity JBOD (Just a Bunch Of Disks) enclosures, and storage controllers. In this white paper, we provide a recommended architecture and best practices deploying for Microsoft Remote Desktop Services (RDS) using Storage Spaces and Dell PowerVault MD1220 array. We provide sizing data for both pooled and personal desktop collections with two PowerVault MD storage array configurations. We evaluate and provide recommendations for Storage Spaces capabilities such as resiliency, data tiering, data deduplication, and file server caching.

2 Audience

Data center administrators who are planning to deploy RDS using Storage Spaces can use this document to better understand Dell's best practices and sizing guidelines. It is assumed that the readers have an understanding of RDS. We also recommend reading Deploying Windows Server 2012 R2 Storage Spaces on Dell PowerVault, which provides information on selecting the Storage Spaces compatible hardware, firmware, and information on supported storage configurations. This document also provides the detailed process for firmware updates and JBOD monitoring features.

In addition to Microsoft Storage Spaces, Dell has other solution architectures with Microsoft RDS, as well as vWorkspace. Refer to reference architecture for Dell Wyse Datacenter for Microsoft VDI and vWorkspace for more information.

3 Overview

This section provides an overview of Storage Spaces and PowerVault MD1220 array.

3.1 Storage Spaces

Microsoft Storage Spaces is a feature set first introduced in Windows Server 2012 and enhanced in Windows Server 2012 R2. Storage Spaces enables to create resilient, cost-efficient, and highly scalable storage system by incorporating industry standard servers, JBOD enclosures, and storage controllers. Storage pools can be created as an aggregate of inexpensive commodity hard drives and presented as a resilient virtual disk. Data is striped across these drives to ensure high performance and reliability. The virtual disk is presented to cluster servers. Cluster Shared Volumes (CSV) are created over the virtual disk to provide a unified name space and access to the underlying storage. By enabling Scale-Out File Server roles on the clustered servers, the storage can be exposed to Hyper-V nodes as Server Message Block (SMB) 3.0 shares and the storage can be used to host virtual machines. In this RA, the clustered servers are the same physical hosts that also function as Storage Spaces servers and Scale-Out File Server roles enabled.

The main advantages of Storage Spaces over a traditional SAN are:

- **Cost:** Traditional SAN typically incorporates proprietary hardware and software, increasing the cost of the solution. Storage Spaces uses commodity servers, SAS adapters and hard drivers thereby decreasing the cost.
- **Manageability:** Storage Spaces reduces the complexity of the storage infrastructure. It can be easily deployed and managed through standard Microsoft tools. Little special training or expertise is required. Storage can be deployed from Microsoft Server Manager or using PowerShell commands (refer to this link for deployment PowerShell commands).

- **Scalability:** Storage infrastructure can be scaled seamlessly and on demand using Storage Spaces. File servers can be added, if more compute resources are needed or disk drives can be added, if more storage capacity or IOPS performance is needed.

The following Microsoft TechEd presentation provides a good overview of Storage Spaces.

Storage Spaces is particularly suitable for VDI workloads because:

- VDI workloads have high random I/O operations. Storage Spaces support tiered storage model with SSD drives, which provides better performance for random I/O operations.
- Storage Spaces provides a good starting point for VDI administrators to deploy a storage solution without requiring advanced expertise with storage technologies like Fibre Channel and iSCSI.
- Cost per user is an important measure in VDI workloads. By using industry standard components, Storage Spaces reduces the storage cost per user for VDI workloads.

3.2 PowerVault MD1220 arrays

PowerVault MD1220 arrays are direct-attached storage enclosures that offer versatile and high-capacity storage for mainstream applications. The arrays provide seamless server and primary storage expansion using 6Gb/s SAS storage adapters. PowerVault MD1220 arrays support mainstream and capacity-intensive applications such as email, backup to disk and video streaming with high capacity SAS drives, as well as high-performance applications with high speed Solid State Drives (SSD). Different drive types and speeds can be mixed to support data tiering. The MD1220 array has dual controllers for high availability and support up to 24 drives (Figure 1).

Figure 1 PowerVault MD1220 with 24 drives

4 Architecture

Figure 2 provides the high-level reference architecture for the solution and shows the logical connectivity among the various solution components.

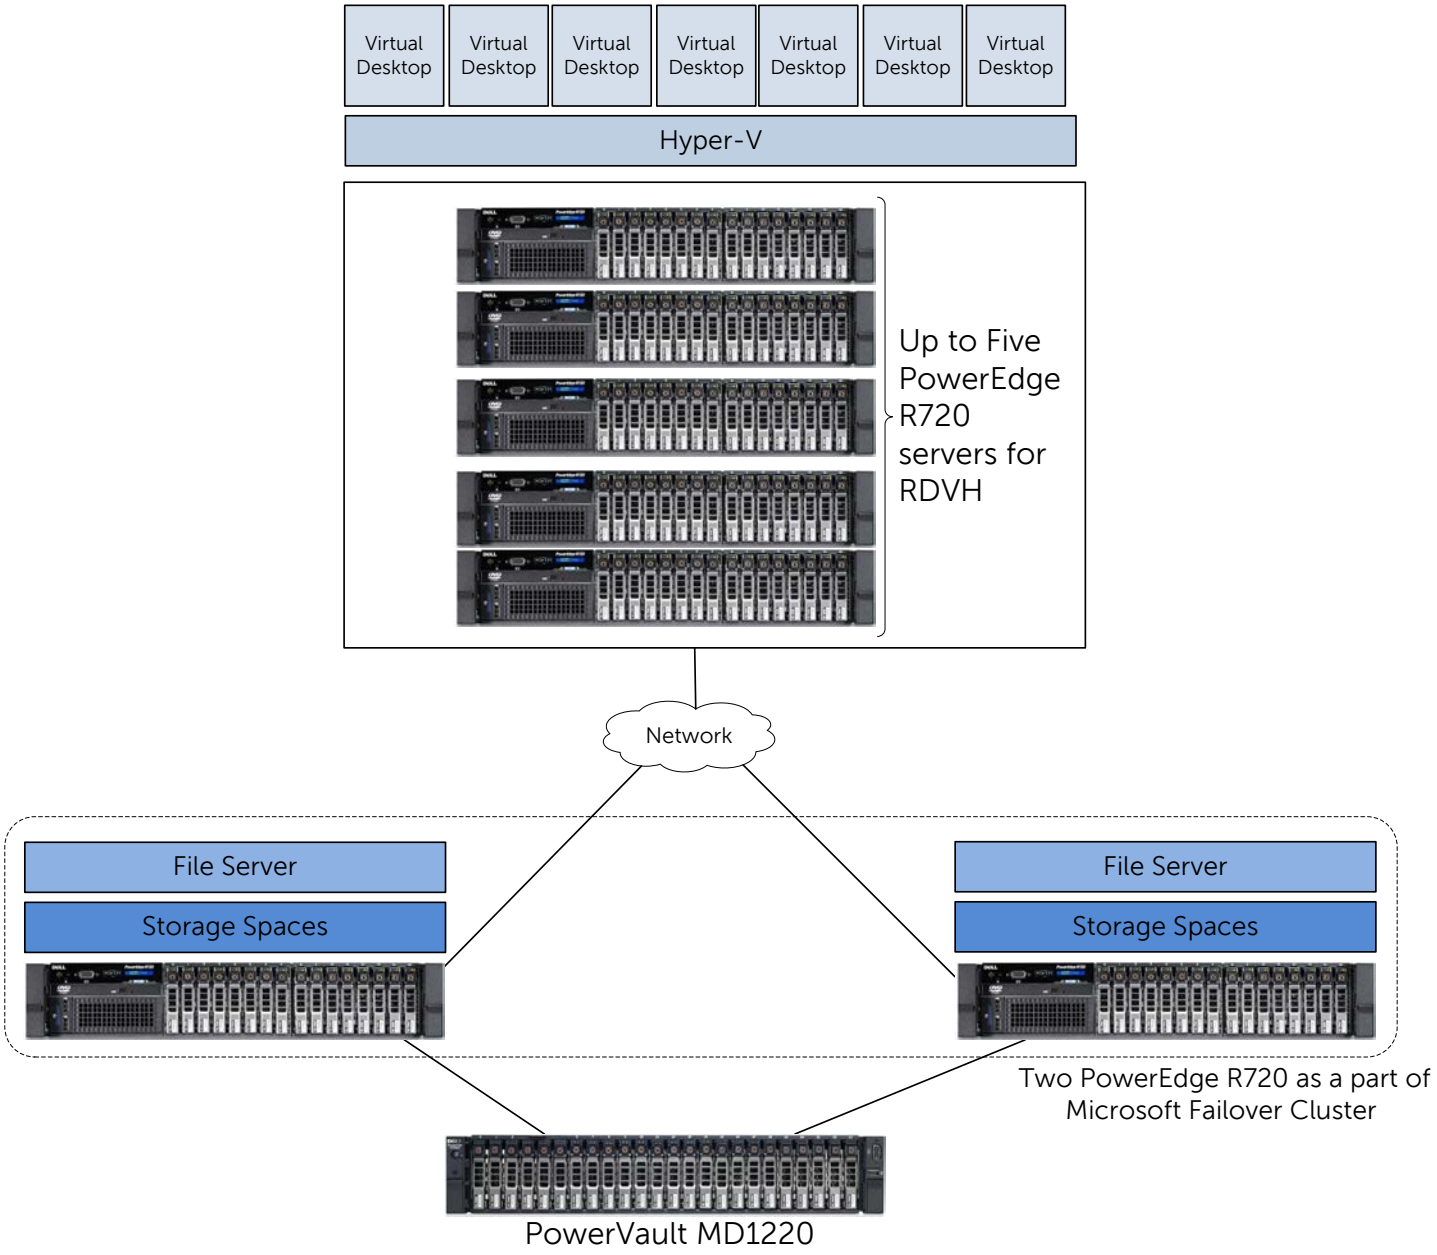

Figure 2 High level Reference Architecture for Microsoft RDS with Microsoft Storage Spaces and PowerVault MD1220 storage array

The solution components include:

- **MD Storage Array:** MD1220 storage array with 24 drives is used to host desktop virtual machines storage. The hard disk configuration of the storage array and its connectivity to the servers are provided in Sections 4.1 and 4.2 respectively.
- **Storage Space Servers (File Servers):** The MD1220 storage array is connected to two Dell PowerEdge R720 servers as Direct Attached Storage (DAS). Using Storage Spaces, storage array is configured to create storage volumes.

The two servers are configured as part of a failover cluster to provide high availability. SMB File Server role is enabled in the server and the storage volumes are exposed to the Hyper-V hosts as SMB shares. These servers are referred as 'File Servers' throughout the rest of the document.

- **Remote Desktop Virtualization Hosts**: Dell PowerEdge R720 servers act as Remote Desktop Virtualization Hosts (RDVH) with Hyper-V roles enabled. The number of PowerEdge R720 servers depends on the user environments, such as application requirement and user profiles. In the configuration described in this white paper, five PowerEdge R720 servers are deployed as RDVH nodes.
- **Other Infrastructure Components** (not shown in the figure)**:** In addition to the above components, RDS requires the following components which can be configured as virtual machines.
    - **Active Directory** services along with DNS and DHCP servers
    - **Remote Desktop Service** - Connection Broker (RDCB) and Web Access (RDWA)

4.1 Storage Configuration

Two types of storage configuration are discussed in this white paper. They are:

1. **SAS Storage Configuration** - (24) 15K 300 GB SAS drives in (1) PowerVault MD1220 storage enclosure
2. **Tiered Storage Configuration** – (4) 400GB Dell Wyse Datacenter qualified SSD drives and (20) 7.2 K 1 TB SAS drive in (1) PowerVault MD1220 storage enclosure

Throughout the rest of the white paper, the various best practices and sizing guidelines are provided for these two configurations.

4.2 Storage Connectivity

The PowerVault MD1220 storage arrays have two storage controllers to support availability. Each controller has two SAS ports. The storage array is connected to each file server using two SAS cables and two Dell 6 Gbps SAS controllers as shown in Figure 3 below.

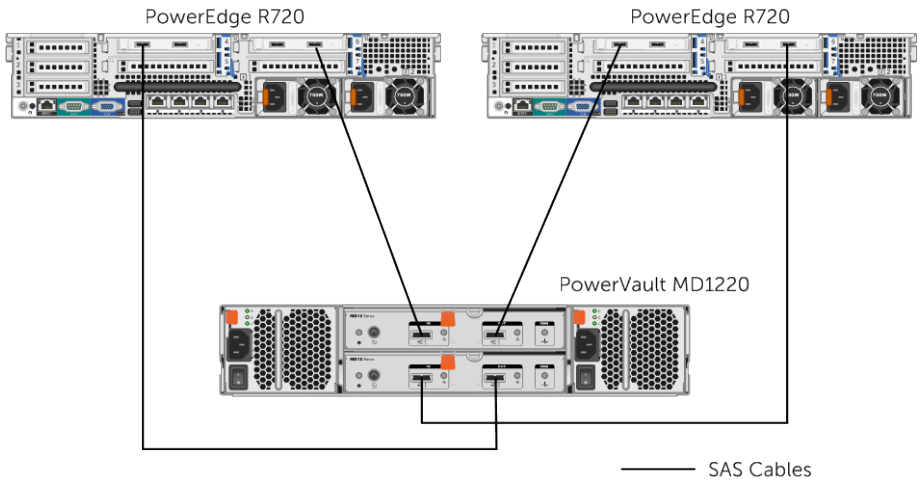

Figure 3 Storage Connectivity between PowerEdge R720 server and PowerVault MD1220

5 Use Cases

Microsoft RDS supports two types of desktop virtual machines:

1. Pooled desktops – Administrators deploy virtual desktops and users are randomly assigned virtual desktops. At log off virtual desktops are generally rolled back.
2. Personal desktops (also called "Persistent VMs") – Each user is assigned one virtual desktop. The assignment is persistent, i.e., when the user logs again, they login to the same virtual desktop.

Pooled desktops storage characteristics: Pooled desktops have one master image. The master image contains the operating system and other applications installation. Each virtual desktop image is simply a differential disk of the master image. Typically, the master image of a Windows 8.1 virtual machine with a few applications installed would be around 10 to 15 GB. The differential disk will be around 2 to 4 GB for each virtual desktop. The differential disk may grow when the users log on and utilize the virtual desktop. The changes are typically rolled back when the users log off. Since most of the storage is already differential disks, data deduplication may provide little benefit for pooled desktops.

Personal desktops storage characteristics: Each virtual desktop image includes the operating system installation as well as the applications and user data. The size of each virtual desktop really depends on the user environments (such as installed applications and user data) and can typically range anywhere from 20GB to 40 GB per virtual desktop. Data deduplication may provide significant benefits, since most of the storage has duplicate data (OS and application installation bits).

6 Capabilities

Storage Spaces provides several capabilities to deliver a highly available and scalable storage solution. Some of the key capabilities that are relevant to the solution discussed in this white paper are:

- **Failover Clustering:** Failover Clustering is a Microsoft Windows Server feature that enables multiple servers (known as cluster nodes) to be formed into a cluster which provides high availability to the underlying storage and supported workloads, such as File Server and Hyper-V. Storage Spaces supports failover clustering. In this white paper, two file servers are configured as a part of failover cluster.

- **Resilience:** Storage Spaces provides three types of resiliency:

    - **Simple:** Simple resiliency is just striping of data across the available physical disks with no data redundancy. Since there is no data redundancy, drive failures will cause data loss. This option is recommended for scenarios in which resiliency is provided through other mechanisms. Simple resiliency would be comparable to traditional RAID 0 configuration.
    - **Parity:** In this option, parity information is stored in disks for data resiliency. Either a single copy (single parity) or two copies (dual parity) of the parity information can be stored. Single parity helps protecting against single disk failures and dual parity helps to protect against two simultaneous disk failures. Parity provides resiliency compared to Simple, at the cost of storage capacity. Parity resiliency would be comparable to traditional RAID 5 or RAID 6 configuration.
    - **Mirror:** In this option, data and parity is stripped and mirrored across physical disks. A two-way mirror stores two copies of data to protect against single disk failures. A three-way mirror stores three copies of

data to protect against two simultaneous disk failures. Since data is mirrored across the drives, this option provides better performance than parity, at the cost of storage capacity. Mirroring would be comparable to traditional RAID 1+0 configuration. Mirror resiliency is recommended by Microsoft for VDI applications.

- **Data Tiering:** Storage Spaces supports data tiering. Virtual disks can be created that are comprised of two storage tiers:
    - Solid-state drive tier for frequently accessed data
    - Hard disk drive tier for less-frequently accessed data

  Storage Spaces transparently moves data at a sub-file level between the two tiers, based on how frequently data is accessed. Storage tiers allow the solution to combine the best attributes of solid-state drives and traditional spinning media hard disk drives. To increase solution performance the most-frequently used ("hot") data is stored on SSD while less-frequently accessed (""cold") data is stored on less expensive traditional drives. The tiering service is automated and runs at scheduled time. IT administrators can also pin certain (assign) files to the standard hard disk drive tier or to the faster solid-state drive tier to ensure that some files are always accessed from the appropriate tier.

- **Data Deduplication:** Windows Server 2012 includes Data Deduplication as a service that can be installed and managed by using Server Manager. It reduces the amount of storage required in a solution by eliminating redundant data. The goal is to store more data in less space by segmenting files into small variable-sized chunks identifying duplicate blocks, and maintaining a single copy of each block. Redundant blocks are replaced by a reference to the single copy. Data deduplication for remote storage of Virtual Desktop Infrastructure (VDI) workloads was introduced in Windows Server 2012 R2. It can optimize active virtual hard disks (VHDs) for VDI workloads on Cluster Shared Volumes (CSVs).

  Storage Spaces supports the native data deduplication feature available in Windows Server. The deduplication algorithm can be scheduled to run automatically at off-peak hours to consolidate the data and free up space. Data deduplication works in tandem when applied with data tiering can yield both improved capacity utilization and performance. Data deduplication would compress duplicate data, thereby freeing space on the faster SSD drives. As a result, SSD drives can subsequently support more hot data yielding better solution performance.

- **CSV Caching:** Cluster Shared Volumes (CSV Cache is a feature in Windows Server 2012 which allocates system memory (RAM) as a write-through cache. The CSV Cache provides caching of read-only unbuffered I/O. This can improve performance for applications such as Hyper-V, which uses unbuffered I/O when accessing a Virtual Hard Disk (VHD) file. CSV Cache is completely integrated into the Failover Clustering feature and handles orchestration across the sets of nodes in the cluster.

Storage Spaces have additional capabilities such as enclosure awareness, which are outside the scope of this document. Refer to Deploying Windows Server 2012 R2 Storage Spaces on Dell PowerVault for more information.

7 Test Configuration

In this section, the test configuration used in this white paper is elaborated. This configuration is used to evaluate the performance of Storage Spaces and PowerVault MD1220 storage array and provide sizing guidelines.

Table 1 Test Configuration

Component	Configuration
Storage	Two configurations are used: 1. SAS Storage Configuration - (24) 15K 300 GB SAS drives in (1) PowerVault MD1220 storage enclosure 2. Tiered Storage Configuration – (4) 400 GB SSD drives and (20) 7.2 K 1 TB SAS drives in (1) PowerVault MD1220 storage enclosure
Storage Space Server (File Server)	(2) PowerEdge R720 Servers with the following configuration: • Processor: (2) Intel Xeon E5-2690 3.0 GHz processor • Memory: 128 GB • Disk: (2) disks configured as RAID 1 for OS image • Network: (4) 1Gb Ethernet ports • Storage adapater: (2) dual port 6 Gbps Dell SAS HBA (to connect to PowerVault MD1220) • OS: Windows Server 2012 R2 with Storage Spaces and File Server roles
Remote Desktop Virtualization Hosts	(5) PowerEdge R720 Servers with the following configuration: • Processor: (2) Intel Xeon E5-2690 3.0 GHz processor • Memory: 256 GB • Disk: (2) disks configured as RAID 1 for OS image • Network: (4) 1Gb Ethernet ports • OS: Windows Server 2012 R2
Network Infrastructure	Dell Force10 S55 1Gb managed Ethernet switches
Other Infrastructure components	The following components were installed as virtual machines to support the infrastructure: • Active Directory with DNS and DHCP roles • Remote Desktop Connection Broker (RDCB) • Remote Desktop Web Access (RDWA)
Virtual Desktops	• Processor: (2) vCPU • Memory: 1 GB (Startup), 512 MB (Minimum), 3 GB (Max) • Disk: o For pooled: One 14GB Master HDD and 2.3 GB diff disk per virtual desktop o For personal: One 10GB HDD per virtual desktop • Network: One virtual NIC • OS: Windows 8.1 • Applications: Microsoft Office 2010, Adobe Acrobat, Internet Explorer

All the components have the latest BIOS, firmware, OS patches and updates.

Evaluation tool: LoginVSI, a software-based load-testing tool, was used to evaluate the performance and provide sizing guidelines. LoginVSI, which can test the performance and scalability of VMware View, Citrix XenDesktop, XenApp, Microsoft RDS and most other VDI solutions, automatically logs in to the virtual desktops for a set number of users and simulates user operations such as opening, utilizing and closing applications. For this evaluation, LoginVSI *medium* user profile is used to simulate the Enhanced User workload.

Please note that the observations and results are based on this simulated workload run in a lab environment. The actual results may vary depending upon the customer environment and workload.

8 Best Practices

Following are some of the best practices for deploying RDS with Storage Spaces and PowerVault MD1220 storage arrays.

8.1 Array Management

By default, Windows Server 2012 R2 implementation of Storage Spaces doesn't provide monitoring functionality for JBOD hardware elements. As a result, power supply, temperature sensor, fan, enclosure identification and overall health status cannot be monitored. Dell and Microsoft worked together to create a management module to monitor the health status of the JBOD. No special installation is required to access the new capabilities. As long the latest Microsoft patches and Dell firmware has been installed, the new health monitoring capabilities will be available.

Following PowerShell commands can be used on the file server to manage the PowerVault storage array:

- Identifying Disk Firmware: *Get-PhysicalDisk | ft -Autosize -Property FriendlyName, SlotNumber, FirmwareVersion, Size, Model, SerialNumber*
- Enclosure Status: *Get-StorageEnclosure | fl*

For more information refer to the document Deploying Windows Server 2012 R2 Storage Spaces on Dell PowerVault.

8.2 Microsoft Multipath I/O

Each file has two paths to the PowerVault MD1220 storage enclosure. It is highly recommended to have at least two paths to the array to avoid any single point of failures. In order for both paths to be active, Microsoft Multipath I/O (MPIO) must be installed in the server. This functionality is available as a role in Windows Server 2012 R2.

To ensure that MPIO is configured and both paths are active, run the following PowerShell command on the file server:

List all the disks in the server:
mpclaim –s –d
List all the paths for a particular desk
mpclaim –s –d <disk number>

For more information on Microsoft MPIO refer to this user guide.

8.3 Planning Storage Pools, Virtual Disks and Volumes

A **Storage pool** is a collection of disk drives. Storage pools can be created to segregate applications with different access pattern and requirements for performance and capacity. For example, a storage pool can be created with SSD drives (or mix of SSD and SAS) for virtual desktops which require higher performance. A separate pool with low speed SAS drives can be created for applications that require low performance and high capacity.

Virtual disks are created on the storage pools. Multiple virtual disks can be created over the same storage pool for multiple desktop collections. Please note that applications using one virtual disk may impact applications in another virtual disk, if they are created over the same storage pool. Data tiering is enabled at the virtual disk level.

Volumes are created on the virtual disks. The configuration described in this white paper has two file servers to provide high availability. In order to take advantage of both the file server compute resources, at least two volumes must be created. Each volume can be manually assigned to one file server from Failover Cluster Manager.

Storage Configuration: For the SAS Storage Configuration, one storage pool was created encompassing all 24 drives. The storage pool was configured with single mirror resiliency. A virtual disk and a volume were created on the storage pool.

For the Tiered Storage Configuration, one storage pool was created encompassing all 24 drives. The storage pool was configured with single mirror resiliency. A virtual disk was created with data tiering enabled. A volume was created on the storage pool.

8.4 Data Tiering

The Tiered Storage Configuration has (4) SSD drives and (20) SAS drives. After running data deduplication, for 700 pooled desktops we found that the percent of total I/Os serviced from the SSD tier is 78%. This is significant considering that the SSD tier is only 10% of the capacity of the total volume. Hot data is serviced from the SSD while the slower SAS drives are used for capacity and holds cold data.

> Data tiering can be manually started by running the following command on the file servers:
> *defrag /c /g /h*

8.5 Data Deduplication

Data deduplication reduces the consumed storage by finding and eliminating redundant data. Data deduplication can be configured in the volume. For the Tiered Storage Configuration, and 150 personal desktops, the deduplication rate was 93%. For 700 personal desktops, the deduplication rate was 76% with a deduplication savings of almost 10 TB. Because of the significant savings, it is recommended to enable data deduplication for personal desktops.

For pooled desktops, the deduplication rate was significantly lower (10%). Pooled desktops derive space efficiency through differential disks and does not benefit from data deduplication. It is recommended not to enable data deduplication for pooled desktops.

Data deduplication is configured to run automatically. The following PowerShell commands can be used on the file servers as well:

- To configure data deduplication on a volume: *Enable-DedupVolume –Volume <VolumeID>*
- To start a data deduplication job: *Start-DedupJob -Volume <VolumeID> -Type Optimization -Memory 50*

8.6 Impact of CSV caching

Cluster Shared Volumes (CSV) Cache is a feature in Windows Server 2012 which allocates system memory (RAM) as a write-through cache. CSV caching improves read performance and reduces the number of read IOPS going to the storage.

CSV cache is enabled by default. The size must be configured by running the following PowerShell command on the file servers:
(Get-Cluster). BlockCacheSize = 2048

Figure 4 shows the result for SAS Storage Configuration and 150 pooled desktops, with CSV cache size of 2GB. "Login" refers to the time period when all the simulated users login to the pooled desktops. "Steady state" refers to the time after all the users have logged in and are now simulating user activies, such as opening and closing applications.

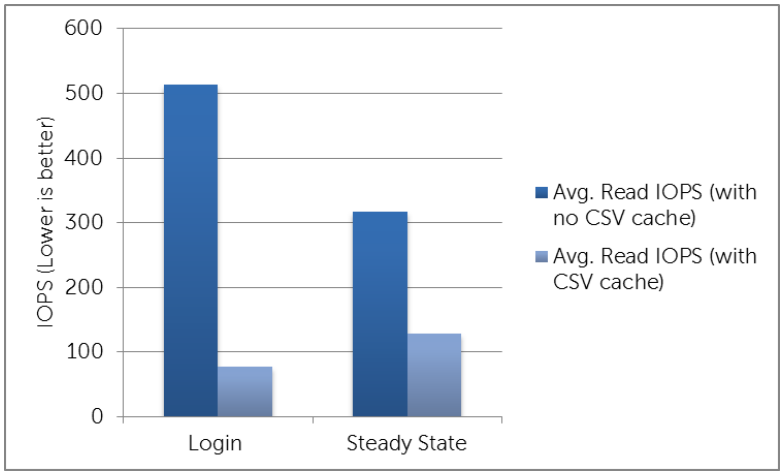

Figure 4 Impact of CSV caching on Read IOPS

Read IOPS are measured at the physical drives. It can be seen that the read IOPS are significantly lower when CSV caching is enabled (85% lower in login and 60% lower at steady state). These IOPS are served from the cache, thereby improving performance. It is recommended to enable CSV caching on the file servers.

Storage Spaces does not support CSV caching for tiered storage. Therefore in the Tiered Storage Configuration, enabling CSV caching will not have an impact. CSV caching is recommended to be enabled only on non-tiered configuration, like the SAS Storage Configuration, to improve the read performance.

8.7 File Server Sizing

Table 1 in section 7 provides the configuration that was used for file server with PowerEdge R720. However, it was observed that the file servers required less compute and memory resources. For 700 users, the file servers consumed less that 10% of CPU and less than 24 GB of RAM. Therefore, it is recommended to deploy the file servers on a PowerEdge R620 with two Intel Xeon E5-2650v2 Processor and 64 GB of RAM for infrastructure for up to 700 virtual desktops, as listed in the latest solution architecture for Dell Wyse Datacenter for Microsoft VDI and vWorkspace.

8.8 Measuring and Understanding Disk Performance

Each I/O operation from the virtual machine gets translated into multiple disk I/O operations. The number of disk IOPS depends on number of disks, resiliency (simple, parity, or mirror) configuration, storage tiering and other factors. The total disk IOPS delivered by the configuration is the key factor in determining the number of desktops supported by the configuration. Therefore, it is important to understand how to measure and interpret the disk performance of the configuration.

Measuring Disk Performance: IOPS can be measured through Windows Performance Monitor on the file servers. Two key objects for measuring disk performance are:

- **Physical Disk Performance Object:** This monitors disk drives on the file server. Each disk in the PowerVault MD1220 enclosure is represented by an instance. The instances are typically numbered. For example, the configurations described in this white paper have 24 disks. There will be 24 instances (numbered from 0 to 23) in this object representing the physical disks. There may be other instances, representing the local drive as well as the logical drive. For this reason, it is recommended NOT to use the instance *_Total* for measuring the aggregate performance of the drives. Each drive must be individually monitored and then the total should be aggregated. The following counters are useful for monitoring the performance of the storage array:
  - Disk Transfers/sec – Measure the total IOPS of the drive
  - Disk Reads/sec – Measure the read IOPS of the drive
  - Disk Writes/sec – Measure the write IOPS of the drive
  - Avg. Disk sec/Transfer – Measure the disk latency for both read and write operations
  - Avg. Disk sec/Read – Measure the disk latency for read operations
  - Avg. Disk sec/Write – Measure the disk latency for write operations
- **Cluster CSV File System:** This object monitors logical CSV volume. Each volume created on the virtual disks is represented by an instance. This instance would provide performance information at the logical level.

Physical Disk IOPS vs Logical Disk IOPS: The IOPS measured at the Volume level (Cluster CSV File System) represent the logical IOPS, which would be significantly lower than the actual IOPS. The disk operations are abstracted at the volume level and hence do not provide a true measure of the performance of the system. To measure the actual disk performance, it is recommended to aggregate the performance of all the physical disk instances. For example, for 700 pooled desktops in Tiered Storage configuration, total physical IOPS is 15624, while the logical IOPS is 4457.

Refer to this blog for more information on understanding of Windows Performance Monitor disk counters.

9 Sizing

This section provides the maximum number of virtual desktops that can be deployed for both SAS Storage Configuration and Tiered Storage Configuration.

SAS Storage Configuration: SAS Storage configuration includes (24) 15K 300 GB SAS drives in (1) PowerVault MD1220 storage enclosure. These drives are configured with Mirror resiliency. For this configuration, up to 150 pooled desktop users could be deployed. The users log in at an interval of 60 seconds. Figure 5 shows the total IOPS measured on the SAS drives for the pooled desktops.

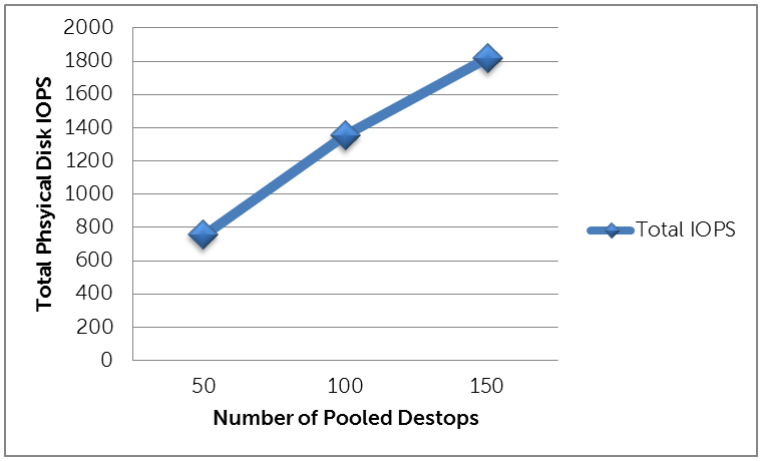

Figure 5 Pooled desktops in SAS Storage Configuration

SAS storage configuration is an ideal configuration for small environments or pilot projects that require low cost storage. To scale this configuration beyond 150 users, new enclosures must be added. If a more scalable and low latency solution is required, it is recommended to use the tiered storage configuration.

Tiered Storage Configuration: Tiered Storage Configuration includes (4) 400 GB SSD drives and (20) 7.2 K 1 TB SAS drives in (1) PowerVault MD1220 storage enclosure. These drives are configured with Mirror resiliency. Data tiering is enabled. Figure 6 shows the total IOPS measured on the SSD and SAS drives for the pooled desktops. For tiered storage configuration, we could run up to 700 pooled desktop users. With data tiering, 78% of the IOPS are served by the SSD drives at average disk latency (Avg. Disk sec/Transfer) of less than 0.5 milliseconds.

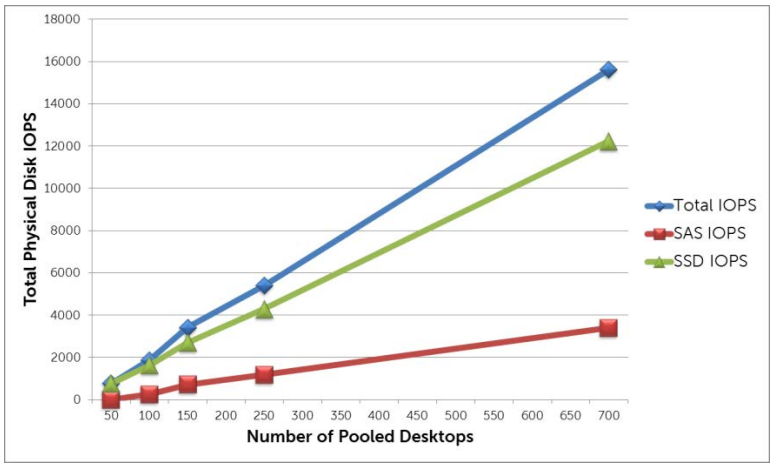

Figure 6 Personal desktops in Tiered Storage Configuration

Similar results for personal desktops are observed. For personal desktops, data deduplication was enabled. We can support up to 700 personal desktop users in Tiered Storage Configuration.

Historically, cost per user for personal desktops is higher compared to pooled desktops. Personal desktops have higher capacity requirements on traditional storage because each desktop require dedicated OS and application installation image. With data deduplication and data tiering, that cost is reduced. Duplicate and cold data are compressed and stored in slower cost efficient drives. SSD drives can service more hot data. With Storage spaces, we have determined that the Tiered Storage Configuration can provide the same number of pooled and personal desktops.

10 Conclusion

Dell provides the most complete Storage Spaces implementation – server, controller, JBOD and disks to form an end-to-end solution. Indeed, the PowerVault MD1220 storage array includes new capabilities for monitoring and managing the storage array. This means that we can provide the best hardware solution for IT departments who recognize that commodity and cost efficient servers, SAS controllers and disk drives can be used to build highly scalable and resilient storage using Microsoft Storage Spaces. Dell is Microsoft's largest OEM partner to support Storage Spaces.

Low cost SAS drives in PowerVault storage array provide a good entry point for desktop virtualization. For solutions requiring more scalability and higher performance, a mix of SSD and SAS drives in PowerVault MD1220 enclosure can be used. Microsoft data deduplication and data tiering enhance the performance and scalability of a tiered storage solution. But whatever the combination of drives and components, Dell's Cloud Client Computing division can work with your IT department to benchmark, recommend, and deploy an appropriately-sized desktop virtualization solution, whatever your pain points and whatever your requirements. Out tested and trusted solutions will enhance your enterprise's data security, user flexibility, and IT manageability and align with your datacenter's requirements and goals for return on investment and time to value.

In addition to the configuration described in this white paper, Dell offers complete end to end solution for desktop virtualization. Dell Wyse Cloud Client Solutions provide thin clients for remote desktop access. Dell vWorkspace software provides in-depth monitoring and diagnostics, advanced VM management capabilities, and enhanced end

user experience. Dell services are available consultation, deployment and support. For more information on end to end solution for Microsoft RDS, visit Dell Wyse Datacenter for Microsoft VDI and vWorkspace.

Acknowledgement

The author would like to acknowledge the following people for their valuable contribution to the white paper: Senthil Bala, Nicolas Cuendet, Steven Hunt, Farzad Khosrowpour, Reed Martin, Syama Poluri, Shruthin Reddy, Shawn Salisbury, Scott Stanford, and Jerry Van Blaricom.