Quest Rapid Recovery Sizing Guide 4.0

Sizing for Rapid Recovery 6.x, AppAssure 5.x and DL Series Appliances

DISCLAIMER: This document is subject to change without notice and is provided “as is” without warranty of any kind, express or implied. Quest does not make any representations regarding the use, validity, accuracy or reliability of the information in this article or the results of the use of this document. The entire risk arising out of the use of the information in this article remains solely with the customer. In no event shall Quest be liable for any direct, consequential, incidental, special, punitive or other damages, even if Quest is negligent or has been advised of the possibility of such damages, arising from use of the information provided herein. This information is not for use by or for individual/consumer purposes.

Contents

Introduction 4
Mileage May Vary 4
General Sizing Considerations 5
Assessing Your Environment 5
Protected Clients 6
Environment Capacity 6
Change Rate per Hour 7
Repository Overview 8
DVM Repository Capacity 8
Deduplication Cache 9
Rapid Recovery Repository Server Architecture 10
RAID Considerations 10
Rapid Recovery Volumes 10
Rapid Recovery Sizing Recommendations 12
DVM Repository Spindle Recommendations 13
Rapid Recovery Core Server RAM Recommendations 14
Calculating Core Server RAM 14
Rapid Recovery Core Server Processor Recommendations 15
Rapid Recovery Core Server Networking Recommendations 15
Rapid Recovery Client Requirements 16
Rapid Recovery Performance Optimization 17
Rapid Recovery as a Virtual Machine 17
Verified Recovery Limits 17
Scheduling Rapid Recovery Snapshots 18
Number of Concurrent Tasks 19
Managing Large Amounts of Clients 19
Rapid Recovery Sizing Troubleshooting 20
Dell Performance Analysis Collection Kit 20
DL Appliance Sizing 20
Quick Sizing Reference 20
Best Practice Summary 22
Repository Sizing 22

Repository Hardware 22
Repository Volumes 22
Dedupe Cache Sizing 23
DVM Repository Optimization Job 24

Introduction

Proper sizing of the Rapid Recovery Core server is vital in allowing Rapid Recovery to properly perform its many tasks. When calculating sizing resources for Rapid Recovery, it is important to ensure that multiple critical factors are taken into account to avoid exhaustion of system resources. The exhaustion of one or more system resources will result in degraded system performance, leading to a poor product experience. The purpose of this document is to provide Rapid Recovery customers with general guidelines to obtain the best performance from Rapid Recovery in their given environment.

When sizing a Rapid Recovery Core, RAM, CPU, available network throughput, storage speed, and capacity are all critical elements. Assigning the correct resources to the Rapid Recovery Core will ensure proper performance and a good user experience. While this guide is designed to help you properly size a Core for your desired backup intervals, you should approach your Rapid Recovery testing in a staged manner. Start with moderate settings (such as once an hour backups and no validation checks turned on), and then adjust the settings as resources allow. By starting with moderate settings and taking a phased approach to adding some of the more robust capabilities (such as high frequency backups and validation checks), you can more readily assess which resources you may need to modify to take full advantage of all of the high powered and advanced features that Rapid Recovery offers.

Mileage May Vary

It is extremely important to note that recommendations provided in this document are general guidelines developed from testing and data obtained from Rapid Recovery customers. Every environment is different, and performance in your environment may vary. This document does not guarantee the optimal performance of Rapid Recovery in every condition, but rather provides guidelines and best practices to help obtain optimal performance. Quest recommends that you read this guide and familiarize yourself with the contents. If you have questions after thoroughly reading this document, please contact a sales or support representative.

Sizing is an ongoing requirement. The Rapid Recovery hardware configuration you build for your current environment may not fit if the environment changes. If you add additional protected machines or additional mail or database servers to your protection regime, or increase other demands on the system such as virtual standby machines, you must revisit the parameters included in the sizing matrix to ensure you haven't exceeded the guidelines for the current tier for which you are configured. Perform this analysis whenever your environment changes.

At a minimum, Quest recommends you review your sizing annually.

General Sizing Considerations

The proper sizing of Rapid Recovery is dependent upon three data points. The accuracy of this data will greatly impact the accuracy of the sizing results provided by this guide. As such, always strive to provide the most accurate data points possible when attempting to size Rapid Recovery. When not sure how accurate a data point is, err on the side of caution and increase its value by 10% to 20%. It is always better to allocate more resources than necessary rather than not enough.

In this sizing guide you will need to know the following data points:

1. The total amount of data that will be protected by a Rapid Recovery Core server.
2. The Change Rate per Hour of the servers that will be protected by Rapid Recovery. If you are not sure of this number, general guidelines are provided in this document.
3. The total number of physical and virtual clients that will be protected by Rapid Recovery.

In addition to the aforementioned data points, you should also be aware of a few key guiding principles. Failure to adhere to these principals can result in adverse Core performance. As such, those who choose not to adhere to these principals do so at their own risk and without support or warranty from Quest.

1. Quest does not recommend running the Rapid Recovery Core on the same physical machine that serves as a Hyper-V host (This recommendation does not apply to DL series appliances).
2. The Rapid Recovery Core may be installed on a VM. When installed on a VM, be sure to allocate sufficient disk, CPU and RAM resources to allow for adequate backup and restore performance. Follow guidelines outlined in the Rapid Recovery as a Virtual Machine section.
3. For the utmost data protection, both the primary and secondary deduplication cache should be stored on a separate disk or a RAID configuration when possible. This mitigates the chances of losing the deduplication cache in the event of disk failure.
4. A Network Attached Storage (NAS) device may be used to host a DVM repository only if it adheres to the sizing guideline presented in this document. See Rapid Recovery Volumes for more details.
5. When increasing the number of protected clients, always resize the Core server based on recommendations in this guide.
6. A Core used as a replication target should be sized using the same criteria as sizing a primary (source) Core.
7. Always maintain at least 20% free space in the DVM repository at all times.

Assessing Your Environment

In order to properly size a Rapid Recovery environment certain key information must first be gathered. You first need to know how many clients you wish to protect, and what level of protection you wish to provide for your systems. Next, you will need to know the total amount of data that you wish to protect, which will aid in selecting the proper size repository. Lastly, you must know how frequently your data changes on an hourly basis, which determines the type of hardware your Core will require and has a bearing on your repository sizing.

Protected Clients

The Table 1 below is designed to assist in determining the number of clients you wish to protect and the protection level that you wish to provide those clients. Rapid Recovery can protect both Microsoft Windows and Linux clients, and supports both agent and agentless protection. A definition of each protection class can be found in the Rapid Recovery Core Server RAM Recommendations section.

When selecting clients to be protected by the Rapid Recovery server, use the following recommendations for scale.

Scale up to 100 clients:

- The Core can only have up to 100 clients if using Class 1-4 clients

Scale up to 130 clients:

- Core can contain only Class 1, 2 and 3 clients
- Snapshots cannot be more frequent than every 60 minutes

Scale up to 150 clients:

- Core contains only Class 1 (Snapshot only) clients
- Snapshots cannot be more frequent than every 60 minutes

Below is an example of how to determine protection for an environment using Rapid Recovery. In this example, protection is determined based on protection type, a server's duty and criticality to the business. Servers that are the most critical should receive the highest levels of protection.

- **Exchange, SQL and SharePoint servers (Agent)**: Class 4
- **Mission critical servers (Agent)**: Class 4
- **Standard application servers (Agent)**: Class 2 or 3
- **VMware VM's (Agentless)**: Class 1

Protection Type	Class 1 Protection	Class 2 Protection	Class 3 Protection	Class 4 Protection	Total
***Example**: Agent quantity*	*10*	*5*	*2*	*35*	*52*
***Example**: Agentless quantity*	*60*	*0*	*0*	*0*	*60*
Agent quantity					
Agentless quantity					

Table 1: Client protection count

Environment Capacity

After you have identified the machines in your environment to be protected with Rapid Recovery, you should then determine their capacity (raw amount of data to be protected). If you do not already have a tool in your environment to assess the size of all the machines to be protected, then consider the use of Dell Performance Analysis Collection Kit (DPACK). This tool

can be obtained through your Quest Support representative to quickly assess the raw capacity of the systems to be protected.

Change Rate per Hour

To properly size a Rapid Recovery Core server, it is important to have an understanding of how much data will be sent to the Core on an hourly basis. The change rate per hour (Cr/h) is the amount of data the Core server must process per hour, before compression, deduplication and encryption. This data is a combination of both new data (data growth) and changes to existing data. The amount of data sent to a Core server has a significant impact on how much I/O (Input/Output) is generated on the repository disk, CPU usage, and RAM usage. Common methods of determining Cr/h are listed below. When calculating Cr/h, it is important to have an understanding of the hours during which the bulk of your change rate will take place e.g., 8am to 6pm.

Method 1: Use data from existing backups

If you currently have a backup product in place this is by far the easiest tool with which to estimate your Cr/h. In a typical backup scheme companies take a weekly full backup and daily incremental backups. Since Rapid Recovery is incremental-forever technology, you can use the data from current incremental backups to calculate you Cr/h. To do this take a look at reports from your current backup and see how large your typical incremental backup is **prior** to compression and deduplication. Then take the backup size and divide it by your backup cadence. For example: If reports show that your incremental backup is 16GB in size and your backup interval is every 10 hours, then you will need to divide your backup size by 10 hours (16GB/10hours = 1.6GB change per hour). Doing this for all protected machines, assuming there are 30 machines with similar Cr/h, would yield a total Cr/h of 48GB.

Method 2: Install a trial of Rapid Recovery

The change rate of protected machines can easily be determined either by looking at the size of each snapshot or by running a Rapid Recovery Core summary report. To check the size of a recovery point, first select an agent and click the Recovery Points tab. From there, look at the Size column for recovery points. Figure 1 shows that the protected agent is taking snapshots every 5 minutes with a change rate of about 50MB. If we multiply 50MB by 12 (the number of snapshots taken every hour) we get a 600MB per hour change rate for this machine. This method works when seeking the change rate for just a few machines but can be time consuming when working with many machines.

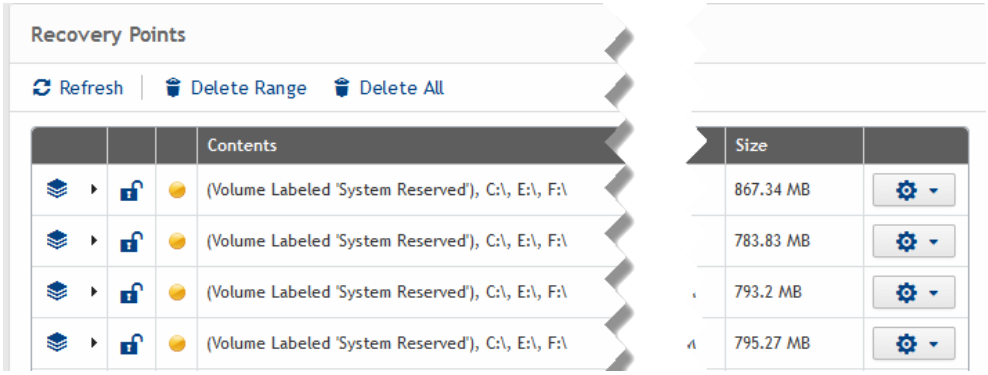

Figure 1: Rapid Recovery snapshots

To determine the change rate of all machines protected by the Rapid Recovery Core you can run a Core Summary report for your company's hours of operation (e.g., 8 or 12 hours). Export the report as an .xls or .xlsx file and then total the median change rate for all of the protected machines. Simply run the report to show the last 8 or 12 hours and then divide the total median change rate per day by the time frame of the report. This should yield your change rate per hour, allowing you to verify that your Rapid Recovery Core is sized properly.

Method 3: Use standard assumptions

Using standard assumptions means that Rapid Recovery will choose a Cr/h for you based on average data seen by customers. As this is a general number; it may not be the most accurate for your environment, but it will reduce the risk in improper sizing. This method should be used as a last resort if you are unable to use method 1 or 2. Please see Table 2 for standard assumptions.

Client Quantity	CR/h
10	10GB
20	20GB
50	100GB
80	300GB
100	400GB
150	500GB

Table 2: Stand change rate assumptions

Repository Overview

Rapid Recovery utilizes the Deduplication Volume Manager (DVM) repository. This repository provides inline deduplication, compression and encryption of incoming data for Windows and Linux clients, both physical and virtual. Quest makes an assumption of 50% savings for compression and deduplication. It is possible to see more or less savings than 50% depending on the type of data sent to the Rapid Recovery repositories. See Table 3 for a list of factors that impact your compression and deduplication ratios.

Higher Ratios	Lower Ratios
Human generated data	Multimedia (Images, videos, sound clips etc.)
Uncompressed and Unencrypted Data	Compressed and encrypted data
Applications with low data change and transfer rates	Applications with low data change and transfer rates
Inactive and infrequently access data	Active and frequently accessed data
Long retention periods	Short retention periods

Table 3: Compression and deduplication ratios

DVM Repository Capacity

To calculate the total space required for a DVM repository multiply the total amount of data to be protected by 1.4x (where x = the total amount of data to be protected). This provides the DVM repository with enough space to protect 1 year of data using the default retention policy, while assuming a 50% deduplication ratio. If your data is expected to yield lower ratios, then increase 1.4x by .1 for every 10% below 50% deduplication you expect to obtain. For example, if you

expect to achieve a 30% deduplication ratio use 1.6x for your calculation. If you expect to obtain more than 50% deduplication, then reduce 1.4x by .1 for every 10% above 50% you expect to obtain.

To retain data for longer than a year, add an additional .5 making the 2-year calculation 1.9x, 3 years 2.4x, and so on. Additionally, Rapid Recovery replication targets will need to be sized based on the data being replicated, as well as any data being stored by local clients. It is highly recommended that when calculating storage requirements, you maintain a 20% storage buffer at all times. Failure to do so many result in adverse repository performance and the slowdown of vital DVM maintenance tasks.

Example: You have calculated that you will need 3TB of repository space for machines at site A and you will need 4TB repository space to for machines at site B. In order to replicate data from site A to site B, site B must have a 7TB repository to hold all of your data.

Note: When sizing the Rapid Recovery repository be sure to take note of clients that are significantly larger than your other clients. In an environment with three agents, totaling 3TB of data, with one agent being 2.5TB in size, 1.4x sizing will not cover an extra base image for the large agent. Ensure that you provide enough space to hold at least two base images for your largest machine(s).

Deduplication Cache

One of the most important concepts to understand regarding the DVM repository is deduplication cache (DC), which is memory allocated and used specifically for deduplication. The DC holds the reference database of pointers for the unique data in the Rapid Recovery repository. The cache is loaded into RAM when the Rapid Recovery Core service starts, and remains there until the Core service is stopped. The DC is periodically persisted (written to disk) in the user-specified primary and secondary cache locations for safekeeping. This persist job takes place every 60 minutes. Because it is not possible to perform snapshots while the DC is persisted to disk, it is important that the DC persist job take place as quickly as possible. Ideally, the HW for the location where the primary and secondary dedupe cache reside should be fast enough to complete the DC persist job within 10 minutes or less. To mitigate the impact that the DC persist job has on the system and the time it takes, there are two options:

1. When possible, strive to keep the DC below 20GB in size. A DC of 20GB or below can typically be written to disk very quickly.
2. In environments where the DC must be larger than 20GB, it is highly recommended to:
    a. Move the DC to a dedicated disk separate from the OS and repository volumes.
    b. Use faster, dedicated SSD storage for the primary and secondary DC. The use of SSD drives will ensure that the DC can be written in a timely manner and prevent disruption to normal Core activities.

There are several factors that affect DC sizing, including the type of data and amount of unique data. In most situations, 1GB of DC is suitable for proper deduplication of about 1TB of raw protected data, assuming a 50% deduplication ratio. This ratio should be adjusted if the data protected is expected to yield a lower ratio. It is extremely important to follow this guideline to prevent the dedupe cache from filling up. Once the cache is full, deduplication rates of new data may be reduced and the chances of prematurely filling the Rapid Recovery repository will increase. For example, a 10TB repository should have a DC 10GB in size. The DC can be set to

a maximum of 50% of the installed system RAM. Ensure that the Core server contains enough RAM to satisfy the DC requirements, Rapid Recovery Core service, and RAM required by Windows. Properly sizing your DC will help prevent performance problems and prevent unduplicated data from being stored in your repository.

Additionally, the location where the DC resides should be a minimum of two times the size of the configured DC. This will ensure that there is enough space for both the primary and secondary DC to be stored on disk. For the utmost data protection, both the primary and secondary DC should be stored on a separate disk when possible, or on disk in a RAID configuration. This ensures that only one copy of the DC is lost in the event of a disk failure.

Rapid Recovery Repository Server Architecture

The DVM repository should reside on disks that are dedicated solely to the repository and to no other applications. The only exception to this is when the DVM is run as virtual machine. When the DVM repository is run on shared storage with other applications, ensure that the storage has enough I/O capacity to meet the needs of the DVM and the shared application(s). For guidelines on I/O requirements see DVM Repository Spindle recommendations.

When using multiple storage arrays, such as Dell MD arrays, make sure that they are chained as one single volume. This will help to ensure that I/O load is evenly distributed across all disk and not just a small subset of disk.

RAID Considerations

Each RAID level has pros and cons, and each provides a mixture of I/O performance and data protection suitable for different environments. In most installations, including the factory default with the DL Series appliances, RAID 6, is used. Review the below table to help decide what RAID level is right for you.

Features	RAID 6	RAID 10	RAID 60	RAID 50
Minimum Number of Drives	4	4	6	8
Data Protection	Up to two disk failure	Up to one disk failure in each sub-array	Up to one disk failure in each sub-array	Up to two disk failure in each sub-array
Write Performance	Low	Medium	Medium	Medium
Read Performance	High	High	High	High
Utilization Capacity	50%-88%	50%	67%-94%	50%-88%

Table 4. RAID level descriptions

Rapid Recovery Volumes

When designing a Rapid Recovery Core, it is important to ensure the DVM repository is placed on fast storage that is able to maintain high levels of I/O. When choosing storage for the Rapid Recovery repository the below guidelines should be followed (For more detailed guidelines see DVM Repository Spindle Recommendations):

- Disk used for the DVM repository should be 7200rpm or faster Near-Line SAS or SAS.
- Primary storage used for the repository should be local and not located at a remote site or the cloud. The ideal storage for the DVM repository is Direct Attached Storage (DAS).
- Network Attached Storage (NAS) should only be used in environments with up to 10 clients, up to 10GB CR/h (Change rate per hour), and a repository size of up to 6TB.

- Repository volumes should not be placed on tiered storage unless they are all on the same tier.

When possible, strive to place the DVM repository on dedicated disk. This will ensure that the DVM does not have to fight other applications for I/O resources. If this is not possible and the DVM repository must be placed on shared storage, ensure that the storage has enough I/O to support all applications that reside on it. Never place a repository on the operating system volume.

Note: Use Dell Performance Analysis Collection Kit to validate that you have enough I/O resources for your applications and the Rapid Recovery DVM.

In situations where Rapid Recovery must be run as a VM see the Rapid Recovery as a Virtual Machine section of this document for more detailed volume setup information.

DVM Configurations

This section presents several examples of how to configure a Rapid Recovery Core and DVM repository. While there are other configurations possible, only the ones listed in this guide have been tested to determine sizing requirements. Deviation from these configurations will require additional work to determine the required resources.

For all of the configurations below use the Sizing Recommendations section to determine the recommended resources for the Rapid Recovery Core and Rapid Data Service Server.

Rapid Recovery DVM configuration 1:

This is the Rapid Recovery default configuration, which will require two volumes and RAID groups. Volume 1 contains the OS, primary and secondary dedupe cache, and uses RAID 1 for redundancy. Volume 2 is dedicated solely to the Rapid Recovery repository and should be in RAID configuration when possible. This is the typical configuration used in small environments and when the Core is run as a VM.

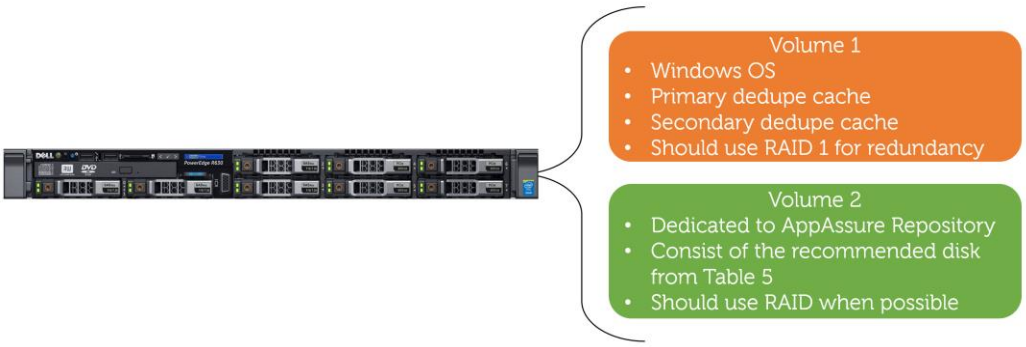

Figure 2: DVM Configuration 1

Rapid Recovery DVM configuration 2:

This configuration is similar to configuration 1, but it provides additional performance for those who have large deduplication caches. In this configuration the DC is installed on SSD drives that allow the cache to be written to disk quickly, speeding up backups.

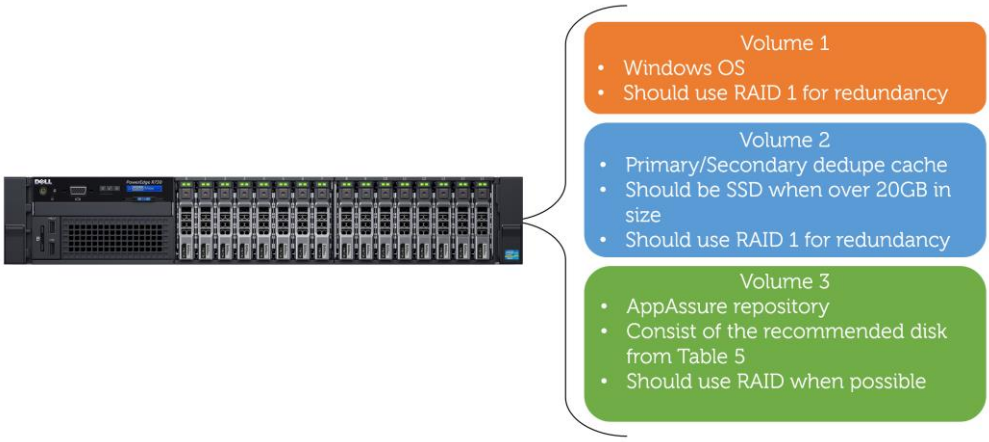

Figure 3: DVM Configuration 2

Rapid Recovery DVM configuration 3:

This configuration is the most advanced in that each Rapid Recovery component resides on a dedicated volume. The deduplication cache and repository metadata reside on SSD storage to provide the best application performance possible. This configuration is highly advised in large environments with high change rates.

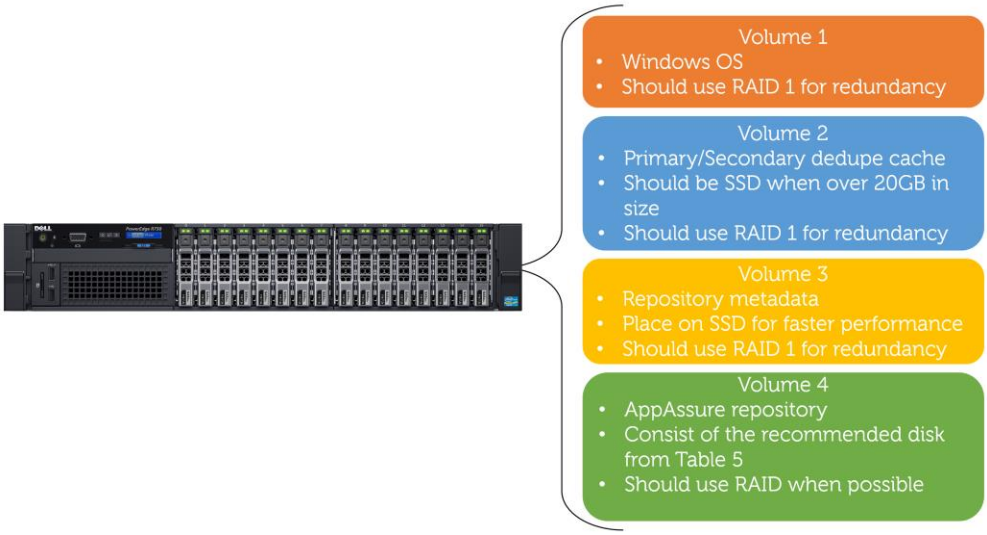

Figure 4: DVM Configuration 3

Rapid Recovery Sizing Recommendations

This section is designed to provide hardware recommendations for the Rapid Recovery software when installing on new or existing hardware. For DL appliance recommendations please refer to DL Appliance/Quick Sizing Reference for guidance.

Note: Recommendations are subject to change as product enhancements are made and new features are added. Make sure to always download the latest version of this guide from the Rapid Recovery technical documentation page.

DVM Repository Spindle Recommendations

Rapid Recovery generates significant random and sequential I/O, and thus requires fast-performing storage. When assessing the speed of the storage to be used with Rapid Recovery, it is extremely important to ensure that the selected storage can support the high amounts of random I/O that Rapid Recovery can generate. When planning your storage needs, make sure to focus on drive quantity to meet your disk I/O performance, rather than capacity requirements. This is because it can take more disk to satisfy your I/O needs than it can for your capacity needs. Failure to size storage in this order can result in having the proper storage capacity but not enough I/O capability. This may result in multiple jobs being queued up, sluggish Core performance, slow backups and slow task execution.

Table 5 represents the typical number of 7200 rpm NL-SAS/SAS drives or spindles that Rapid Recovery requires to perform successfully in a given environment. These numbers are based on multiple factors including the average I/O performance of 7200 rpm NL-SAS/SAS drives and the assumption that multiple backup, restore, mount, Verified Recovery, virtual standby, and replication tasks are being performed on the Core server. These requirements represent the typical usage for a Rapid Recovery administrator. The ability of your hardware to meet these requirements should allow Rapid Recovery to perform adequately. Please note that some environments may exceed the expected usage tested by Quest and require additional resources on their Rapid Recovery Core server.

Total Change Rate/hour	DVM Repository Up to 100 Clients
Up to 10GB	NAS/SAN/DAS 4 spindles RAID 6
Up to 20GB	SAN/DAS 4 spindles RAID 6
Up to 60GB	SAN/DAS 6 spindles RAID 6
Up to 100GB	SAN/DAS 8 spindles RAID 6
Up to 200GB	SAN/DAS 12 spindles RAID 6
Up to 300GB	SAN/DAS 16 spindles RAID 6
Up to 400GB	SAN/DAS 20 spindles RAID 6
UP to 500GB	SAN/DAS 24 spindles RAID 6
Other	Contact a Quest representative

Table 5: Spindle recommendations for Rapid Recovery and DL appliances

Note: This chart represents the minimum number of spindles required to obtain optimal Rapid Recovery repository performance. This table assumes that each 7200pm drive is capable 75 IOPS, measured at 32KB with 75% reads in RAID 6 with 60% random I/O. You may need additional drives to satisfy capacity requirements as outlined in previous sections.

Rapid Recovery Core Server RAM Recommendations

Rapid Recovery clients can be broken down into four protection classes based on their targeted protection level and protection type. The more aggressive the protection the more resources Rapid Recovery Core will need. Below are sizing estimates for the Rapid Recovery protection tiers based on average RAM consumption observed in testing and various customer environments.

Clients	Agent Protection	Agentless Protection	RAM
Class 1	Snapshot only	Snapshot only	.3 GB
Class 2	Snapshot + Virtual Standby + Verified Recovery Checks	Snapshot + Virtual Standby + Volume Integrity Check	.5 GB
Class 3	Snapshot + Replication + Verified Recovery Checks	Snapshot + Replication + Volume Integrity Check	.5 GB
Class 4	Snapshot+ Virtual Standby + Replication + Verified Recovery Checks	Snapshot+ Virtual Standby + Replication + Volume Integrity Check	1 GB

Table 6: Rapid Recovery client RAM consumption

Note: When sizing a target Core used for replication use the same guidelines shown in table 5.

Calculating Core Server RAM

The following calculations should be applied to all Rapid Recovery configurations listed in the Rapid Recovery Repository Server Architecture section. This ensures that the appropriate resources are available in all scenarios.

DVM RAM calculation

Once the total number of Rapid Recovery clients for each class is known and the dedupe cache size has been calculated, the Core server RAM can be calculated with the below formula. 8GB is the minimum amount of required RAM for a Core with a DVM repository:

Core Server RAM = (Class 1 protection * .3 GB) + (Class 2 protection * .5 GB) + (Class 3 protection * .5 GB) + (Class 4 protection * 1 GB) + (Dedupe Cache Size > 1.5 GB) + 8 GB (Minimum required RAM)

Note: If the Rapid Recovery Core server is running on Windows Server 2008 or 2008 R2, add an additional 2GB of RAM per TB of repository space.

Rapid Recovery Core Server Processor Recommendations

Rapid Recovery processor consumption increases as the number of clients, tasks and total change rate per hour (CR/h) increase. The following table is meant to provide a general guideline as to what class of CPU will be required for a given number of clients and a given data consumption rate.

Total CR/h	10 Clients	20 Clients	50 Clients	80 Clients	100 Clients
10 GB	Single Quad-Core CPU	Single Quad-Core CPU	Dual Quad-Core CPU	Single Six-Core CPU	Single Six-Core CPU
20 GB	Single Quad-Core CPU	Single Six-Core CPU	Single Six-Core CPU	Dual Six-Core CPU	Dual Six-Core CPU
60 GB	Single Quad-Core CPU	Single Six-Core CPU	Dual Six-Core CPU	Dual Six-Core CPU	Dual Six-Core CPU
100 GB	Single Six-Core CPU	Dual Six-Core CPU	Dual Six-Core CPU	Dual Eight-Core CPU	Dual Eight-Core CPU
200 GB	Dual Six-Core CPU	Dual Six-Core CPU	Dual Eight- Core CPU	Dual Eight-Core CPU	Dual Eight-Core CPU
300 GB	Dual Six-Core CPU	Dual Six-Core CPU	Dual Eight- Core CPU	Dual Eight-Core CPU	Dual Eight-Core CPU
400 GB	Dual Six-Core CPU	Dual Eight-Core CPU	Dual Eight- Core CPU	Dual Eight-Core CPU	Dual Ten-Core CPU
500 GB	Dual Six-Core CPU	Dual Eight-Core CPU	Dual Ten- Core CPU	Dual Ten-Core CPU	Dual Ten- Core CPU

Table 7: Rapid Recovery Core Processor recommendations

Note: Processors should be 2.2 GHz or faster.

Rapid Recovery Core Server Networking Recommendations

Providing adequate bandwidth to the Rapid Recovery Core server is vital to maintaining ideal Core performance and preventing jobs from stacking up. Below in Table 8 are general recommendations as to what speed network cards to use and how many should be present in the Rapid Recovery Core. These recommendations are based on a given CR/h while assuming all clients are performing outgoing replication and Virtual Standby tasks. It is important to note that in some situations the Core may require more network bandwidth than recommended in the below table. Some situations that may cause this to occur are high numbers of virtual standbys, large amounts of incoming and/or outgoing replication tasks, ongoing Live Recovery or ongoing bare metal recovery tasks.

Network Card(s)	Max CR/h
1 GB NIC	10GB
2x 1 GB NIC	100GB
4x 1 GB NIC	200GB
6x 1 GB NIC or 1x 10GB NIC	500GB

Table 8: Network recommendations

Note: When using NIC teaming be sure to use LACP (Switch independent).

A quick method to tell if your Rapid Recovery Core server has adequate bandwidth is using the Windows Task Manager. Under the performance tab of the Windows Task Manager you can see the current Send/Receive rate of your Ethernet card(s). If you are not consistently sustaining the maximum bandwidth of your Ethernet connection during peak backup periods (See Figure 5), then your cards are sized appropriately. However, if you are constantly at the limits of what your network card can provide, you should consider adding extra network cards. See Using Windows Performance Analysis Collection to collect network usage stats throughout longer time periods.

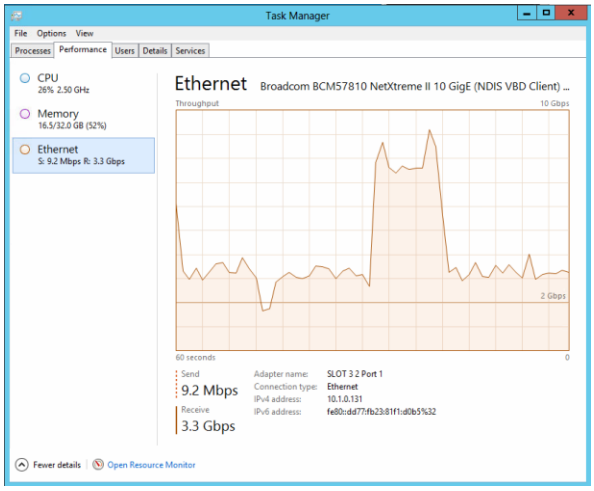

Figure 5: Network performance

Rapid Recovery Client Requirements

While a properly sized Rapid Recovery Core is vital to backup performance, a well performing client is just as important. Failure to have protected clients meet the minimum system requirements listed in the Rapid Recovery Release Notes can result in slow backups, VSS errors and other quality issues.

In addition to meeting minimum systems requirements it is important to ensure the disk on protected systems is not over-taxed. Often, systems under heavy I/O load will manifest issues sending changed blocks of data to the Rapid Recovery Core server. Common symptoms of a client with too much I/O load are VSS errors during backup and excessively slow transfer speeds. Slow transfer speeds for one or multiple clients can have a large impact on your backups, as there is a limited number of backup jobs that can run at one time.

To avoid such situations, ensure that your protected systems are not already approaching their I/O capacity when Rapid Recovery backups are not active. Systems that are nearing capacity during periods when backups are not running typically incur a greater risk of having performance problems during Rapid Recovery snapshots. This is also the case with agentless VMware backups. Attempting to back up too many virtual machines at one time can place additional load on the ESX environment, and in extreme cases may cause issues with backups.

Ensure that systems running Microsoft Exchange and SQL, and also performing Verified Recovery tasks, have ample resources. Verified Recovery tasks will place additional load on SQL and Exchange servers if checks are done on the agent side. If Exchange and/or SQL

servers are resource constrained, consider running the Verified Recovery checks on the Rapid Recovery Core.

Note: When troubleshooting backup performance issues, look at the performance and resource utilization of the protected clients as well as the Core. Problems do not always stem from the Core even if errors and issues appear at the Core. To determine the current performance of your protected machines or troubleshoot suspected I/O issues please see Use Dell Performance Analysis Collection Kit.

Rapid Recovery Performance Optimization

Rapid Recovery as a Virtual Machine

When operating the Rapid Recovery Core on a virtual machine there are some best practices that should be followed to ensure Core performance. Failure to follow these best practices may result in poor Core performance. These best practices apply to both VMware vSphere and Microsoft Hyper-V.

- Ensure that all volumes are thick provisioned, Eager Zeroed for VMware, or configured as Fixed Disk in the case of Microsoft Hyper-V.
- Place the Rapid Recovery repository and dedupe cache on a dedicated datastore when possible. If not possible, then place repository and dedupe cache on the datastore with the most available I/O.
- Ensure you have the recommended number of dedicated spindles or I/O available for the Rapid Recovery repository. (Disk depth queue should not be greater than the number of spindles + 1).
- Ensure that Rapid Recovery VM resources are not balanced, reserved, shared or limited.

Verified Recovery Limits

The Rapid Recovery Core provides the ability to perform application checks, also known as Verified Recovery. With Verified Recovery, Rapid Recovery validates the integrity and constancy of Microsoft SQL and Exchange databases, and checks the file system on each volume. The checks performed on MS SQL and Exchange databases place additional load on the Rapid Recovery repository and the Core. As such, it is recommended to follow the below guidelines to minimize the load on the Core and the repository. Failure to adhere to these guidelines may result in adverse Core performance.

Additionally, the same principles apply to Virtual Standby tasks. Each Virtual Standby task places additional load on the Rapid Recovery repository that can lead to performance issues if not properly sized. It is highly advised to adhere to the Virtual Standby guidelines provided below, in conjunction with the sizing recommendations provided in this guide.

Virtual Standby tasks:

- Tasks should be limited to 10 Virtual Standby tasks on Cores sized to protect up 10 clients
- Tasks should be limited to 15 Virtual Standby tasks on Cores sized to protect 20 clients or more

Verified Recovery for Microsoft SQL Servers:

If class 2, 3 or 4: Total MS SQL servers should not exceed 20 databases or 3TB per Rapid Recovery Core

If class 1: No limitations when checks are disabled

Verified Recovery for Microsoft Exchange Servers:

If class 2, 3 or 4: Total MS Exchange servers should not exceed 10 databases or 2TB per Rapid Recovery Core

If class 1: No limitations if checks are disabled

Scheduling Rapid Recovery Snapshots

Rapid Recovery has the ability to perform snapshots of protected machines as frequently as every 5 minutes and as infrequently as once per day. Protection schedules can also be broken down into periods. This allows you to tune Rapid Recovery for a high snapshot frequency during peak hours, and to detune snapshot frequency during off-peak hours.

An important factor to consider when setting up off-peak backup schedules is that Rapid Recovery will use this time to perform I/O intensive tasks such as attachability checks, checksum checks, rollups and other nightly tasks. Depending on the number of machines that need to perform the aforementioned tasks and the amount of data, the time for completion can be anywhere from 1 hour to over 4 hours. To help mitigate how long these tasks take and prevent them from spilling over into production backup hours, it is a good idea to halt all backups while these tasks are running. By default, Rapid Recovery will perform nightly tasks at 12:00AM, leaving ample time for all tasks to complete for most environments. However, in environments with a large number of clients and/or data this time may not be sufficient. In these situations, configure a gap in your protection intervals starting at 12:00AM and ending when production starts (see Figure 6). This should enable Rapid Recovery to complete its nightly tasks considerably faster and within a reasonable time frame. If you continue to have problems with nightly tasks finishing before production hours start, consider starting the nightly tasks sooner in addition to creating a period where no backups are taking place.

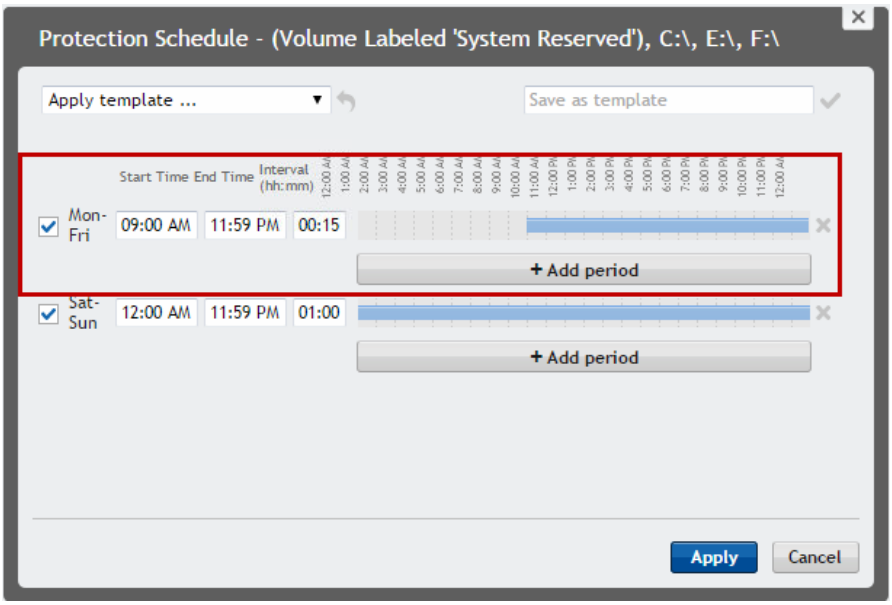

Figure 6: Rapid Recovery off-peak schedule

Number of Concurrent Tasks

Rapid Recovery allows users to define how many backup, Virtual Standby, replication and other tasks take place simultaneously. While it may be tempting to increase these tasks from their default values, do so with caution. In addition to causing extra load on the Rapid Recovery Core, increasing the number of tasks that can run simultaneously on the Core can cause backups and other tasks to perform more slowly. Too many tasks operating on the Core at one time can cause the Core to bump into CPU, RAM and I/O constraints. In most situations, I/O is the first constraint reached, which is much harder to detect than CPU and RAM constraints. Exhausted resources become bottlenecks that slow down the Core in multiple places, such as the UI, backup, replication, restore, and virtual standby tasks.

If you are suffering from any of the above issues, and have modified any of the default settings, start your trouble shooting by reverting to defaults. If reverting concurrent tasks back to their default settings does not help, contact Quest support.

Managing Large Amounts of Clients

In order to maintain a fast and responsive UI when protecting more than 20 clients it may be necessary to disable Toast alerts. These alerts can be disabled by using the following steps:

- Click the More icon and select **Notifications**.
- Click the drop down under Notification Groups and select **Edit**.
- Click **Notification Options** in the Edit Notification Group window.
- Deselect **Notify by Toast alerts** and click **OK**.

Rapid Recovery Sizing Troubleshooting

Dell Performance Analysis Collection Kit

If you are experiencing performance problems with your Rapid Recovery Core, please contact your Quest Support representative or partner to gain access to the Dell Performance Analysis Collection Kit (DPACK). Running DPACK on the Rapid Recovery Core will collect vital information about the Rapid Recovery Core such as, RAM, CPU and disk usage statics. This information will allow Quest representatives to assess the issues that you may having.

DL Appliance Sizing

DL series appliances are pre-configured servers that have AppAssure 5.4.3 pre-installed and optimized for best performance. Purchasing a DL series appliance simplifies the sizing and deployment process as a majority of the work is already done for you. To properly size a DL appliance, simply follow the Quick Sizing Reference table. The Quick Sizing Reference is color coded in **green**, **grey**, and **blue** with each color square indicating which appliance is best suited for a given client and CR/h combination. Each square provides baseline recommendations, but you will still need to use the DVM Repository capacity, Deduplication Cache, and Rapid Recovery Core Server RAM Recommendations sections to determine the appropriate repository capacity, deduplication cache size and RAM required for the DL appliance.

Although the DL appliance is mostly sized by the Quick Sizing Reference table to simplify deployment, it is still important to follow the recommendations presented in this document. Doing so will help ensure that the DL operates properly in your given environment. Also be sure to revisit this guide any time you add additional clients to be protected by Rapid Recovery that were not already accounted for.

Quick Sizing Reference

The table on page 21 is a quick reference guide that combines many of the tables presented in this document into one table. This table may be used to size AppAssure 5.x, Rapid Recovery 6.x, and DL series appliances.

Total Change Rate \hour	Up to 10 Clients	Up to 20 Clients	Up to 50 Clients	Up to 80 Clients	Up to 100 Clients
Up to 10GB	NAS\SAN\DAS 4 useable spindles min Single Quad Core CPU 1GB NIC Example: Dell R320\ Virtual Machine\ DL1000	SAN\DAS 4 useable spindles min\RAID 6 Single Quad Core CPU 1GB NIC Example: Dell R630\ DL1000	SAN\DAS 4 useable spindles min\RAID 6 Dual Quad Core CPU 1GB NIC Example: Dell R630\ DL4300	SAN\DAS 4 useable spindles min\RAID 6 Single Six Core CPU 1GB NIC Example: Dell R630\ DL4300	SAN\DAS 4 useable spindles min\RAID 6 Single Six Core CPU 1GB NIC Example: Dell R630\ DL4300
Up to 20 GB	SAN\DAS 4 useable spindles min\RAID 6 Single Quad Core CPU 2X 1GB NIC Example: Dell R630\DL4300	SAN\DAS 4 useable spindles min\RAID 6 Single Six Core CPU 2X 1GB NIC Example: Dell R630\DL4300	SAN\DAS 4 useable spindles min\RAID 6 Single Six Core CPU 2X 1GB NIC Example: Dell R630\ DL4300	SAN\DAS 4 useable spindles min\RAID 6 Dual Six Core CPU 2X 1GB NIC Example: Dell R630\ DL4300	SAN\DAS 4 useable spindles min\RAID 6 Dual Six Core CPU 2X 1GB NIC Example: Dell R630\ DL4300
Up to 60 GB	SAN\DAS 6 useable spindles min\RAID 6 Single Quad Core CPU 2X 1GB NIC Example: Dell R630\DL4300	SAN\DAS 6 useable spindles min\RAID 6 Single Six Core CPU 2X 1GB NIC Example: Dell R630\DL4300	SAN\DAS 6 useable spindles min\RAID 6 Dual Six Core CPU 2X 1GB NIC (LACP 802.3 is required) Example: Dell R630\ DL4300	SAN\DAS 6 useable spindles min\RAID 6 Dual Six Core CPU 2X 1GB NIC (LACP 802.3 is required) Example: Dell R630\DL4300	SAN\DAS 6 useable spindles min\RAID 6 Dual Six Core CPU 2X 1GB NIC (LACP 802.3 is required) Example: Dell R630\DL4300
Up to 100 GB	SAN\DAS 8 useable spindles min\RAID 6 Single Six Core CPU 2X 1GB NIC (LACP 802.3 is required) Example: Dell R630\ DL4300	SAN\DAS 8 useable spindles min\RAID 6 Dual Six Core CPU 4X 1GB NIC (LACP 802.3 is required) Example: Dell R630\ DL4300	SAN\DAS 8 useable spindles min\RAID 6 Dual Six Core CPU 4X 1GB NIC (LACP 802.3 is required) Example: Dell R630\ DL4300	SAN\DAS 8 useable spindles min\RAID 6 Dual Eight Core CPU 4X 1GB NIC (LACP 802.3 is required) Example: Dell R730xd\Dl4300	SAN\DAS 8 useable spindles min\RAID 6 Dual Eight Core CPU 4X 1GB NIC (LACP 802.3 is required) Example: Dell R730xd\Dl4300
Up to 200 GB	SAN\DAS 12 useable spindles min\RAID 6 Dual Six Core CPU 4X 1GB NIC (LACP 802.3 is required) Example: Dell R630 & MD1400\ DL4300	SAN\DAS 12 useable spindles min\RAID 6 Dual Six Core CPU 4X 1GB NIC (LACP 802.3 is required) Example: Dell R630 & MD1400\DL4300	SAN\DAS 12 useable spindles min\RAID 6 Dual Eight Core CPU 4X 1GB NIC (LACP 802.3 is required) Example: R730xd\ DL4300	SAN\DAS 12 useable spindles min\RAID 6 Dual Eight Core CPU 4X 1GB NIC (LACP 802.3 is required) Example: R730xd & MD1400\ DL4300	SAN\DAS 12 useable spindles min\RAID 6 Dual Eight Core CPU 4X 1GB NIC (LACP 802.3 is required) Example: R730xd & MD1400\ DL4300
Up to 300 GB	SAN\DAS 16 useable spindles min\RAID 6 Dual Six Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: Dell R630 & MD1400\ DL4300	SAN\DAS 16 useable spindles min\RAID 6 Dual Six Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: Dell R630 & MD1400\ DL4300	SAN\DAS 16 useable spindles min\RAID 6 Dual Eight Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: R730xd\ DL4300	SAN\DAS 16 useable spindles min\RAID 6 Dual Eight Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: R730xd\ DL4300	SAN\DAS 16 useable spindles min\RAID 6 Dual Eight Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: R730xd \ DL4300
Up to 400 GB	SAN\DAS 20 useable spindles min\RAID 6 Dual Six Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: Dell R630 & MD1400\ DL4300	SAN\DAS 20 useable spindles min\RAID 6 Dual Eight Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: Dell R730xd & MD1400\ DL4300	SAN\DAS 20 useable spindles min\RAID 6 Dual Eight Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: R730xd & MD1400\ DL4300	SAN\DAS 20 useable spindles min\RAID 6 Dual Eight Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: R730xd & MD1400\ DL4300	SAN\DAS 20 useable spindles min\RAID 6 Dual Ten Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: R730xd & MD1400\DL4300 HC
Up to 500 GB	SAN\DAS 24 useable spindles min\RAID 6 Dual Six Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: Dell R630 & MD1400\ DL4300	SAN\DAS 24 useable spindles min\RAID 6 Dual Eight Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: Dell R730xd & MD1400\ DL4300	SAN\DAS 24 useable spindles min\RAID 6 Dual Ten Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: R730xd & MD1400\ DL4300 HC	SAN\DAS 24 useable spindles min\RAID 6 Dual Ten Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: R730xd & MD1400\DL4300 HC	SAN\DAS 24 useable spindles min\RAID 6 Dual Ten Core CPU 6X 1GB NIC or 1x 10GB NIC (LACP 802.3 is required) Example: R730xd & MD1400\DL4300 HC
Other	**If you have a change rate per hour above 500GB or need to protect more than 100 agents please contact a Dell representative.**				

Best Practice Summary

Repository Sizing

- Size your repository for what you will grow into.
- Extending the repository – Extend in large sizes. The fewer the extents the better the repository performance.
- Maximum repository size – Maximum repository size depends on the HW the repository resides on. The maximum size we can recommend will be the largest size supported by the latest DL appliance as long as your HW is of matching or greater specifications.

Repository Hardware

- Logical or Virtual Disk:
  - Minimum number of spindles per logical disk is four. The higher the spindle count, the better the disk IO.
  - Minimum disk speed is 7200 RPM. The faster the disk, the better the disk IO.
- Order of preference for HW storage is DAS, SAN, and then NAS.
  - Presenting the repository volume to the OS as a local volume is the preferred connection method. Local volumes are presented through RAID controller or iSCSI connected targets and appear in Disk Management.
  - CIFS shares do not have the connection stability nor error tracking needed in case connection issues arise. CIFS do not appear in Disk Management.
  - NAS is supported, but only for repositories 6 TB or less, supporting 10 agents or less, with a change rate of 10 GB per hour or less.
  - Single USB drives are not considered DAS nor a HW option for a repository.
- Virtual Disk Policies:
  - RAID: Choose whichever RAID suits your performance needs
  - Read Policy: Read Ahead (options may vary by vendor, but it should be enabled)
  - Write Policy: Write Back (options may vary by vendor, but it should be enabled)
  - Stripe Element Size: Default (64 KB is default for most)
  - Disk Cache Policy: Enabled (options may vary by vendor, but it should be enabled)

  Note: Virtual disk policies are recommendations for non-appliance HW only. Virtual disk policies for DL appliances are preconfigured during provisioning and should not be changed for any reason.
- If a repository is connected over a network, connect using a dedicated connection, 1 Gbps or higher.
- It is not recommended to place a repository volume on tiered storage. You may experience performance issues.
- Cloud connected repositories are not supported. Repositories must be local to the core.

Repository Volumes

- Never place the repository on the same volume as the operating system.
- The repository should be placed on a dedicated volume on a dedicated logical disk.
- Do not compress or enable Windows dedupe on the repository volume.
- Do not run a schedule defragmentation program on the repository volume.
- Do not use any type of SW or HW to take a backup of the repository.

- For maximum disk IO you should have one repository with one extent per volume, one volume per logical disk.
    - Several extents added to the same volume may impact disk performance. The fewer the extents the better.
    - More than one repository per volume will impact disk performance.
    - A single repository per volume with several volumes on a single logical disk will impact performance.
- It is not recommended to place a repository volume on tiered storage. You may experience performance issues.
- Always thick provision the volume(s) a repository resides on. For example, disks attached to virtual machines such as vhd, vhdx, and vmdk or HW solution such as SANs that thin provision LUNs.

Dedupe Cache Sizing

- Know the amount of raw data you plan to back up. Basic rule of thumb is 1 GB of dedupe cache will dedupe 1 TB of raw data.
- Dedupe cache greater than 10 GB can reduce core performance when the dedupe cache is persisted to disk.
- If you size a dedupe cache greater than 20 GB, it is recommended the Primary cache, Secondary cache, and Cache metadata locations be moved to a dedicated SSD.
- Size your dedupe cache wisely, for example: If you have 32 TB of raw data this does not mean you need a 32 GB dedupe cache.
    - You will never obtain 100% dedupe.
    - Dedupe % varies according to data types.
    - Large dedupe cache will affect core performance.
    - You may be wasting memory for dedupe cache that the core can utilize for better performance.

      Note: The higher the dedupe and compression %'s the longer it will take to hydrate data during restore.
- Suggested steps in sizing dedupe cache:
    - Start your dedupe cache size at 25% to 50% of the amount of data to be deduped.
    - Let jobs run normally for seven days.
        - If you have a very low dedupe percentage, leave the dedupe cache size alone. If you wish to increase the dedupe cache size to see if the rate changes, do so in small amounts.
        - If you have a high dedupe percentage, increase the dedupe cache size in small increments.
        - Repeat until you no longer gain higher dedupe percentages.
        - When increasing the dedupe cache size, ensure you have enough free RAM, free disk space, and HW that can handle the increased disk IO before increasing the size.

      Note: Increasing the dedupe cache size is easy and you do not lose existing dedupe cache when doing so. However, decreasing the dedupe cache deletes the current dedupe cache files and creates new ones. You will lose all current dedupe cache data if you downsize.

DVM Repository Optimization Job

A new optimization job was added to Rapid Recovery 6.0 to allow repository data to be re-deduplicated.

The optimization job reviews all data in the repository, and reclaims duplicate space by effectively deduplicating all data contained in the existing recovery points. The optimization job is processor intensive, and the amount of time it takes to run this job depends on several factors. These factors include the size of your repository; the amount of data in your repository; available network bandwidth; and existing load on the input and output of your system. The more data in your repository, the longer this job runs.

For best results, the *repository optimization job* should be run only after first extending the Dedupe Cache. Note that this job must run to completion, otherwise, all progress will be lost if the running job is canceled. While reclamation of space in the repository is possible, it is not guaranteed.

For more information on this topic, see Optimizing a DVM Repository in the Rapid Recovery User Guide, found on the Rapid Recovery technical documentation page.