



The science behind the report:

Get through the migration maze with less downtime with help from Dell Data Migration Services

This document describes what we tested, how we tested, and what we found. To learn how these facts translate into real-world benefits, read the report [Get through the migration maze with less downtime with help from Dell Data Migration Services](#).

We concluded our hands-on testing on April 11, 2022. During testing, we determined the appropriate hardware and software configurations and applied updates as they became available. The results in this report reflect configurations that we finalized on February 11, 2022 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

To learn more about how we have calculated the wins in this report, go to <http://facts.pt/calculating-and-highlighting-wins>. Unless we state otherwise, we have followed the rules and principles we outline in that document.

Table 1: Timed results for our migration testing. Source: Principled Technologies.

Migrating data manually		Block or file data transfer time	Downtime	Migrating data using Dell Data Migration Services	
Activity	Time (hh:mm)			Time (hh:mm)	Activity
Pre-work (research, documentation, prerequisite installs)	02:53			1:15	Pre-work (Datadobi infrastructure prerequisites)
Configuring the PowerStore array	00:23			0:23	Configuring the PowerStore array
Configuring storage components on the servers	00:23			0:23	Configuring storage components on the servers
Setup block data migration	NA**			2:00	Setup block data migration
Migration of block data	01:02			2:30	Migration of block data
Block migration downtime	01:02			0:00	Block migration downtime
Setup file data migration	NA*			3:45	Setup file data migration
Migration of file data	17:00			38:15	Migration of file data
File migration downtime	17:00			0:05	File migration downtime
Validating block and file data	02:56			0:25	Validating file data
				0:20	Validating block data
Post-migration cleanup (decommissioning storage)	00:24			0:05	Post-migration cleanup (decommissioning storage)
Total downtime				0:05	Total downtime
Total elapsed time	25:02			49:21	Total elapsed time

*Completed as part of Configuring storage components on the servers (above)

System configuration information

Table 2: Detailed information on the servers we tested.

Server configuration information	4x Dell™ PowerEdge™ R740 servers
BIOS name and version	Dell Inc. 2.12.2
Non-default BIOS settings	All defaults used
Operating system name and version/build number	Red Hat® Enterprise Linux® release 8.5
Date of last OS updates/patches applied	03/01/16
Power management policy	Balanced
Processor	
Number of processors	2
Vendor and model	Intel® Xeon® E5-2670 v3
Core count (per processor)	15
Core frequency (GHz)	2.60
Stepping	1
Memory module(s)	
Total memory in system (GB)	128
Number of memory modules	16
Vendor and model	Hynix HG4894829485757320X
Size (GB)	8
Type	PC4-17000
Speed (MHz)	2,133
Speed running in the server (MHz)	2,133
Storage controller	
Vendor and model	Dell PERC H730p
Cache size (GB)	2
Firmware version	1.0
Driver version	1.0.1
Local storage	
Number of drives	8
Drive vendor and model	Seagate SST900587
Drive size (GB)	900
Drive information (speed, interface, type)	15K, 12Gb SAS, HDD
Network adapter	
Vendor and model	Intel X540
Number and type of ports	2 x 10GbE
Driver version	1.0.1

Server configuration information		4x Dell™ PowerEdge™ R740 servers
Cooling fans		
Vendor and model	NMB 06038DA-12S-EWH	
Number of cooling fans	6	
Power supplies		
Vendor and model	Delta 0Y26KX	
Number of power supplies	2	
Wattage of each (W)	1,100	

Table 3: Detailed information on the storage we tested.

Storage configuration information	HPE 3PAR StoreServ 8450	Dell PowerStore 5000T
Controller firmware version	3.3.1.460 (MU3)	1.0.2.1.3.273
Number of storage controllers	2	2
Number of storage shelves	1	1
Number of drives per shelf	48	21
Drive vendor and model number	SanDisk® DOPE1920S5xnNMRI	Samsung® MZ-WLL1T9C
Drive size (TB)	1.92	1.92
Drive information	SAS SSD	NVMe™ SSD
NVRAM count	0	4
NVRAM size (GB)	N/A	8.5

How we tested

Installing the operating system

Installing Red Hat Enterprise Linux (version 8.0 on the database server and 8.3 on the web server)

1. Insert a bootable USB key.
2. Start the server, and press F11 to enter the BIOS Boot Manager.
3. In the settings interface BIOS Boot Menu, choose One-shot UEFI Boot Menu or One-shot BIOS Boot Menu.
4. Choose to boot from USB Disk Or Virtual CD (depending on the installation media).
5. From the installation interface, select the first item to start the installation. Note: The installation interface will appear automatically.
6. Press Enter, and the installer loads the pre-install order.
7. After the pre-install process is complete, the graphical interface appears. Select the installation language.
8. In the configuration interface, configure the following parameters:
 - Select the time and date.
 - Keep the default language.
 - Keep the default keyboard settings.
9. On the Installation Summary screen, click Installation Destination.
10. On the Installation Destination screen, accept the default, and click Done.
11. On the Installation Summary screen, click Software Selection.
12. From the Base Environment pane, select Server. Leave the additional software selections at default, and click Done.
13. On the Installation Summary screen, click Network & Host Name.
14. In Network & Host Name, choose your hostname, turn your Ethernet connection on, and click Configure.
15. In Editing [your Ethernet connection], click Connect automatically with priority: 0, and click Save.
16. Click Done.
17. At the Installation Summary screen, click Account.
18. Enter the Red Hat Customer Portal username and password for later repository access, leaving everything else at default.
19. Under Connect to Red Hat, a Registered message appears.
20. On the Installation Summary screen, click Begin Installation.
21. On the Configuration screen, click Root Password, and set a password.
22. After installation completes, select Reboot.

Installing Red Hat Enterprise Linux (version 7.9 on the file client and 7.5 on the mail server)

1. Boot the system, and press F10.
2. In the left pane, click OS Deployment.
3. In the right pane, click Deploy OS.
4. From the list of operating systems, select Red Hat Enterprise Linux 7x86_64 bit.
5. LifeCycle Controller extracts the driver update disk to an internal USB drive labeled OEMDRV.
6. After the drivers are extracted, LifeCycle Controller prompts you to insert the operating system installation media.
7. Click Next.
8. Select BIOS or UEFI, and click Next.
9. Insert the Red Hat Enterprise Linux installation media, and click Next.
10. Click Finish. Continue installing the operating system by booting to the operating system media.

Configuring the Fibre Channel switches

We created Fibre Channel zones according to best practices. For the server port connections, we created two peer zones; one containing the A ports from each bare-metal server and storage, and the other containing the B ports from each bare-metal server and storage.

Creating the Fibre Channel zones

1. Log into the Brocade switch GUI.
2. Click Configure→Zone Admin.
3. Click the Zone tab.
4. Click New Peer Zone, and enter a zone name.
5. Select the desired single server port WWNs as principal members and the desired storage WWNs as peer members.
6. To add the WWNs to the zone, click the right arrow.
7. Repeat steps 4 through 6 for all remaining peer zones.
8. Click the Zone Config tab.
9. Click New Zone Config, and enter a zone configuration name.
10. Select all newly created zones.
11. To add the peer zones to the zone config, click the right arrow.
12. To save the changes, click Zoning Actions→Save Config.
13. Click Zoning Actions→Enable Config.

Configuring the storage

Once we established all FC connectivity, we configured the 3PAR storage for our test environment. First, we enabled FilePersona on the 3PAR array (following the direction in <https://support.hpe.com/hpesc/public/docDisplay?docId=c04617546>) so we could use file storage in addition to block storage. Then, we configured volumes and file shares for each bare-metal server as we detail in Table 4.

Table 4: Storage breakdown for each bare-metal server in our tests.

Name	Allocated capacity (TB)	Used capacity (TB)	Notes
Database server			
Vol0 (data)	2	2	
Vol1 (logs)	2	2	
File client			
NFS mounted to /mnt/shares/test via 10.222.2.120	10	4.4	Shared across three NFS and CIFS mounts
File Share SMB-RO mounted to /mnt/cifsro 10.222.2.120	10	4.4	Shared across three NFS and CIFS mounts
File Share SMB-CIFS mounted to /mnt/cifs via 10.222.2.220	5	1.4	
Web server			
Vol2	2	2	
Vol3	2	1.1	
Mail client			
File Share NFS2 mounted to /mnt/shares/nfs via 10.222.2.120	10	4.4	Shared across three NFS and CIFS mounts

Configuring the 3PAR storage array

We followed the steps in the following sections of the [Document - HPE 3PAR Red Hat Enterprise Linux, CentOS Linux, Oracle Linux, and Oracle VM Server Implementation Guide | HPE Support](#) documentation for the 3PAR setup: Configuring the HPE 3PAR StoreServ Storage and host—FC, Setting up multipathing software on the host, Provisioning storage and host LUN discovery, Using HPE 3PAR devices on the host, Booting the host from the HPE 3PAR StoreServ Storage, and Configuring additional HPE 3PAR software features.

Configuring the Dell PowerStore array

For initial PowerStore setup, we followed all applicable steps in the documentation from <https://dl.dell.com/content/manual37523884-dell-emc-powerstore-host-configuration-guide.pdf> and <https://www.dell.com/support/manuals/en-us/powerstore-500t/pwrstr-ntwkg/initial-configuration-wizard?guid=guid-7c24263b-bc9a-4c43-b523-215f86215b6b&lang=en-us>. We configured LUNS, shares, and IPs as part of the manual install.

Installing the software

To simulate a small-scale data center environment, we set up each server to fulfill a certain role within the data center.

Table 5: The Linux versions and other software we installed on each server.

Host role	OS	Software
Database server	RHEL 8.0	SQL Server, DVDStore
File client	RHEL 7.9	Kept Default Software from RHEL installation
Web server	RHEL 8.3	Apache, MySQL™, PHP
Mail server	RHEL 7.5	Postfix

Installing Microsoft SQL Server 2019 on RHEL 8.0

1. On the server that will run SQL Server 2019, log in as root and navigate to the command-line interface.
2. Install and configure Python version 2 as the default for the server:

```
yum install -y python2
alternatives --config python # select option for python2
```

3. Install Microsoft SQL dependencies:

```
yum install -y bzip2 gdb cyrus-sasl cyrus-sasl-gssapi
```

4. Download the repository file from <https://packages.microsoft.com/config/rhel/8/mssql-server-2019.repo>.
5. Copy the following file to the yum repo directory:

```
cp mssql-server-2019.repo /etc/yum/repos.d/
```

6. Download MS SQL Server RPM, and install it without checking dependencies:

```
yum download mssql-server
rpm -i --nodeps mssql-server-15.*.x86_64.rpm
```

7. Configure and start Microsoft SQL Server 2019:

```
/opt/mssql/bin/mssql-conf setup
```

8. To select Evaluation, type 1.
9. To accept the license, type Yes.
10. Create the administrator password for SQL Server by typing it twice.

Configuring the DVD Store 2 benchmark for database generation

We generated the data using the Install.pl script included with DVD Store version 2.1 (DS2), providing the parameters for our 100GB database size and the database platform we used. We ran the Install.pl script to generate the database schema.

After processing the data generation, we built the 500GB database in SQL Server.

The only modification we made to the schema creation scripts were the specified file sizes for our database. We explicitly set the file sizes higher than necessary. Other than this file size modification, we created and loaded the database in accordance to the DVD Store documentation. Specifically, we followed these steps:

1. Generate the data, and create the database and file structure using database creation scripts in the DS2 download. Make size modifications specific to the 500GB database, making the appropriate changes to drive letters.
2. Create database tables, stored procedures, and objects using the provided DVD Store scripts.
3. Set the database recovery model to bulk-logged to prevent excess logging.
4. Load the generated data into the database using the provided DVD Store scripts.
5. Create indices, full-text catalogs, primary keys, and foreign keys using the database-creation scripts.

Installing Apache HTTP Server 2.4 Lifestream and httpd package on RHEL 8.3

1. On the RHEL server designated as the web server, log in as root and navigate to the command-line interface.
2. Install the httpd package, including its default dependencies, by entering:

```
sudo yum -y install httpd
```

3. To start this newly installed web server and make sure it will automatically be started after a reboot, enter:

```
sudo systemctl enable --now httpd
```

4. Open up ports 80 and 443 by entering:

```
sudo firewall-cmd --add-service=http --add-service=https  
sudo firewall-cmd --add-service=http --add-service=https --permanent
```

Installing MySQL and PHP on RHEL 8.3

1. On the web server, log in as root and navigate to the command-line interface.
2. Add the MySQL Yum repository to your system's repository list. Follow these steps:
 - a. Go to the download page for MySQL Yum repository at <https://dev.mysql.com/downloads/repo/yum/>.
 - b. Select and download the release package for your platform.
3. Install the downloaded release package with the following command, replacing platform-and-version-specific-package-name with the name of the downloaded package:

```
sudo rpm -Uvh mysql80-community-release-el8-3.noarch.rpm
```

4. Install the MySQL package and other required packages:

```
sudo yum install mysql-community-server  
Start the MySQL server with the following command:  
systemctl start mysqld
```

5. Install PHP:

```
sudo yum -y module install php
```

6. Enable connecting to the MySQL database from PHP:

```
sudo yum -y install php-mysqlnd
```

7. Restart Apache to pick up PHP and the PHP MySQL library:

```
sudo systemctl restart httpd
```

Installing Postfix on RHEL 7.5

1. On the web server, log in as root and navigate to the command-line interface.
2. Check to make sure that sendmail is not already running:

```
systemctl status sendmail
```

3. If sendmail is running, you must stop it before installing and configuring Postfix:

```
systemctl stop sendmail
```

4. Disable sendmail so that it does not get restart automatically when the system is rebooted:

```
systemctl disable sendmail
```

5. Remove sendmail from the system:

```
dnf remove sendmail
```

6. Add the smtp service to the firewall as follows:

```
firewall-cmd --permanent --add-service=smtp
```

7. Install Postfix, and configure a special Postfix user in the /etc/passwd file:

```
dnf install postfix
```

8. Edit and save the /etc/postfix/main.cf file in the format as follows:

```
myhostname = mtal.domain.com
mydomain = domain.com
myorigin = $mydomain
mydestination = mydestination = $myhostname, localhost.$mydomain, localhost, $mydomain
inet_interfaces = all
mynetworks = subnet
```

9. Start Postfix:

```
systemctl start postfix
```

10. Configure Postfix to start automatically at system startup:

```
systemctl enable postfix
```

Configuring the Dell PowerStore 5000T array

This section contains the steps we took to set up and configure the Dell PowerStore 5000T storage array. The first step to configure the array is to define the hosts that will attach to it.

Configuring the hosts

1. Log into the Dell EMC PowerStore web management application, and navigate to the Compute drop-down list. Select Hosts & Host Groups.
2. Click Add Host.
3. Provide a Friendly name and Linux Operating System, and click Next.
4. Leave the Fibre Channel radio button selected, and click Next.
5. Choose the appropriate Initiator(s), and click Next.
6. Click Create.
7. Repeat the process for each host under test.

Creating the Fibre Channel SCSI LUNs

1. At the top of the screen, navigate to the Storage drop-down.
2. From the Storage drop-down list, choose Volumes.
3. Click on Create
4. Choose an appropriate name for the volume, for Category, choose Other, and skip Application.
5. For Quantity, enter 1, enter the desired volume size, and click Next.
6. Leave the SCSI radio button selected, and select the check box for the desired host.
7. Click Next, and click Create.
8. Repeat the above steps for the required number of LUNs.

Creating the network-attached volumes

1. At the top of the screen, navigate to the Storage drop-down.
2. Select NAS Servers, and click Create.
3. Provide a name for the NAS server, enter an IP address and Subnet mask, and click Next.
4. For CIFS-based connections:
 - a. Select the SMB box, and click Next
 - b. Select Standalone for Windows Server Type, provide a Netbios name, and fill in appropriate WorkGroup.
 - c. Provide an Administrator password, and click Next.
 - d. On the DNS page, click Next, and click Create NAS Server.
5. For NFS-based connections:
 - a. Click the NFSv4 box, and click Next.
 - b. Click Use Local Files, and click Next.
 - c. On the DNS page, click Next, and click Create NAS server.
6. At the top of the screen, navigate to the Storage drop-down.
7. Click File Systems.
8. Click Create.
9. Choose the appropriate NAS server check box, and click Next.
10. Provide a name, enter a size for the File System, and click Next.
11. Provide a name for SMB and NFS shares, and click Next.
12. For NFS shares, select Read/Write, under Default Access, allow Root, and click Add Host.
13. Provide the required host IP address, and click Save.
14. Click Next, and click Create File System.
15. For NFS shares: After the file system has been created, navigate to the NAS Servers page, and click the NAS Server name.
16. On the next screen, click the Naming Services tab, and click Local Files.
17. Click Upload Local Files, and upload the passwd, group, & hosts file to the PowerStore array.

Configuring storage components on the servers

Attaching and configuring Fibre Channel LUNs to hosts

Configuring multipath manager and map LUNs

1. Log into the appropriate server, and issue the `multipath -ll` command to ensure that the DM software is present.
2. Save the `/etc/multipath.conf` file to a backup file.
3. Using a Linux editor (e.g., VI), edit the `/etc/multipath.conf` file, and add the following content to the end of the file:

```
devices {
  device {
    vendor DellEMC
    product PowerStore
    path_selector "queue-length 0"
    path_grouping_policy group_by_prio
    path_checker tur
    detect_prio yes
    failback immediate
    no_path_retry 3
    rr_min_io_rq 1
    fast_io_fail_tmo 15
    max_sectors_kb 1024
  }
}
```

4. Save the `multipath.conf` file, and restart the service:

```
systemctl restart multipathd
```

5. Ensure the `multipath` service restarted properly, and shut down the server.
6. Log into the Dell EMC PowerStore web management application, and from the Storage drop-down list, select Volumes.
7. Select the check box for the desired FC volume, click Provision, and choose Map.
8. Click the check box for the desired host, and click Apply.

Configuring LVM2 (Logical Volume Manager) to create file systems and mount LUNs

1. Power on the server.
2. Log into the server, and issue the `multipath -ll` command to ensure that the selected PowerStore LUNs are visible to the host.
3. Make note of the device name for the required LUNs (i.e., `mpathx`).
4. Issue the `pvcreate` command to create a physical volume for each LUN being added to the host:

```
pvcreate /dev/mapper/<value for mpathx>
```

5. Issue the `pvs` to confirm creation.
6. Issue the `vgcreate` command to create a volume group to hold the physical volumes:

```
vgcreate <vggroup name> /dev/mapper/<mpathx> /dev/mapper/<mpathy>
```

7. Issue the `vgs` or `vgdisplay` commands to confirm the creation of the volume group:

```
vgs
vgdisplay <vggroup name>
```

8. Issue the `lvcreate` command to create the required number of logical volumes:

```
lvcreate -L <LUN size> -n <logical volume name> <volume group name>
```

- Execute the `lvs` or `lvdisplay` command to verify create of logical volumes:

```
lvs
lvdisplay
```

- For each new volume, create a file system on each:

```
mkfs -t xfs /dev/<vg name>/<lv name>
```

- Make a directory for each volume to be mounted. For example:

```
mkdir /mnt/vol4
mkdir /mnt/vol5
```

- Mount the new volumes to their mount points. For example:

```
mount /dev/mapper/vgroup1-lvol4 /mnt/vol4
mount /dev/mapper/vgroup1-lvol5 /mnt/vol5
```

- Verify the volumes are mounted:

```
df -h
```

- Edit the `/etc/fstab` file to include the newly created PowerStore volumes. For example:

```
/dev/mapper/vgroup1-lvol4 /mnt/vol4          xfs     defaults    0 0
/dev/mapper/vgroup1-lvol5 /mnt/vol5          xfs     defaults    0 0
```

Attaching and configuring network-attached volumes to hosts

- Log into the Dell EMC PowerStore web management application, and from the Storage drop-down, select File Systems.
- Click the SMB SHARES tab, and reference the SMB Share Path (IPv4) address for the hosts with CIFS connectivity.
- Click the NFS EXPORTS tab, and reference the NFS Export Path (IPv4) address for the hosts with NFS connectivity.
- For the host(s) with SMB connectivity, create a file in `/etc` to hold the host's login credentials for the PowerStore's SMB share(s).
- For example, a file named `/etc/pssmb_credentials.txt` would contain the following data:

```
username=Administrator
password=<password>
domain=<host's domain name>
```

- Log into the appropriate server, and edit the `/etc/fstab` file to add a mount point for the shared IP address from the PowerStore.

- For example, the host with both SMB & NFS shares would look something like:

```
10.222.200.161:/PS_FS1_NFS2 /mnt/shares/test nfs defaults 0 0
//10.222.200.162/PS_SMBCIFs /mnt/cifs cifs credentials=/etc/pssmb_credentials.txt,rw,uid=1000,gid=1000,nounix,iocharset=utf8,file_mode=0777,dir_mode=0777 0 0
//10.222.200.161/PS_SMB-RO /mnt/cifsro cifs credentials=/etc/pssmb_credentials.txt,rw,uid=1000,gid=1000,nounix,iocharset=utf8,file_mode=0777,dir_mode=0777 0 0
```

- Issue the `mount -a` command to mount the new network shares.
- Verify the volumes are mounted:

```
df -h
```

Migrating the data (manual migration)

Once we mapped, formatted, and made the PowerStore volumes available to the hosts, we copied all the data from the 3PAR storage array over to the new PowerStore array:

From each host, issue the appropriate (recursive) copy commands to copy all the data from the 3PAR array to the PowerStore array. For example:

```
cp -r nfs/images/ pctest/ &
cp -r /mnt/shares/nfs /mnt/shares/pctest &
cp -r /mnt/vol13 /mnt/vol15 &
```

Validating the data (manual migration)

Once all the migration to the PowerStore array was complete, we ensured that all data had moved successfully. We used the Linux `diff` command:

```
diff /mnt/shares/nfs /mnt/shares/pctest
diff /mnt/vol12 /mnt/vol14
diff /mnt/vol13 /mnt/vol15
```

Removing legacy storage

Once we completed data validation, we removed the old storage array from the environment.

1. Create new mount points to serve as temporary mount points to remove the old storage array. For example:

```
mkdir /mnt/vol2old
mkdir /mnt/shares/nfs_old
```

2. Edit /etc/fstab on all hosts to mount previous storage locations to the new mount points.
3. Reboot hosts to ensure new mount points contain the proper locations and data from old storage unit.
4. Edit /etc/fstab on all hosts to use the original mount points for the new PowerStore array locations.
5. Reboot hosts, and ensure the locations and data from the PowerStore array are mounted on the original mount points.
6. Edit /etc/fstab on all hosts to remove references to 3PAR, and reboot.
7. Unmap hosts from the 3PAR array.

Read the report at <https://facts.pt/CbS3bgW> ▶

This project was commissioned by Dell.



Facts matter.®

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

DISCLAIMER OF WARRANTIES; LIMITATION OF LIABILITY:

Principled Technologies, Inc. has made reasonable efforts to ensure the accuracy and validity of its testing, however, Principled Technologies, Inc. specifically disclaims any warranty, expressed or implied, relating to the test results and analysis, their accuracy, completeness or quality, including any implied warranty of fitness for any particular purpose. All persons or entities relying on the results of any testing do so at their own risk, and agree that Principled Technologies, Inc., its employees and its subcontractors shall have no liability whatsoever from any claim of loss or damage on account of any alleged error or defect in any testing procedure or result.

In no event shall Principled Technologies, Inc. be liable for indirect, special, incidental, or consequential damages in connection with its testing, even if advised of the possibility of such damages. In no event shall Principled Technologies, Inc.'s liability, including for direct damages, exceed the amounts paid in connection with Principled Technologies, Inc.'s testing. Customer's sole and exclusive remedies are as set forth herein.