Hewlett Packard
Enterprise

HPE Reference Architecture for Disaster Recovery solution with HPE Data Protector and HPE StoreOnce

Technical white paper

Contents

Executive summary

An effective business continuity and disaster recovery plan can be the difference between companies that survive, or not, when disaster strikes. Very careful planning and a technology deployment that supports the plan, are crucial in getting the business back to normal operations. Recovery objectives that were once measured in days are now measured in minutes, or seconds. Information Technology is the enabler of the business, and with increased requirements to enable the business, there is no tolerance for any downtime, whether it is a small or medium business, or a large enterprise.

Disasters can occur anywhere and anytime in your business environment. Natural disasters, such as earthquakes, tsunamis, hurricanes, fires, and floods can be devastating. Human errors, and hardware and software failures, are just as unpredictable and expensive. Proactive disaster recovery planning and implementation is the best way to prevent loss. In the U.S. alone the number of disaster declarations by the Federal Emergency Management Agency (FEMA) are increasing year by year. When disaster strikes your data needs to be accessible from an off-site location, and it needs to remain accessible for an unknown length of time. Staff may be limited and time is will be important, so you need to know you're protected, because every minute your system is down the financial implications grow. The Hewlett Packard Enterprise ConvergedSystem 700 and Converged Infrastructures with Disaster Recovery (DR) solution offerings can help you reduce business disruptions.

In the case of mission-critical transactional and analytical workloads, Oracle Databases are dominant among the applications driving these workloads. More importantly, many major applications rely on Oracle Database architectures within the application stack. Deploying a solid database architecture with a disaster recovery solution virtually or physically is a key success indicator that can mean the difference between leading and following your competition.

A solid database architecture can make the difference between competitive differentiation and simple comparative parity. Competitive organizations establish aggressive recovery point objectives (RPOs) and recovery time objectives (RTOs) to help minimize data loss and ensure application recovery and reliability. These organizations choose primary infrastructure and data protection strategies to deliver application-consistent backups, application-reliable recoveries, user-defined service levels, and high availability, along with the ability to maximize resource utilization. You can find these requirements in a non-integrated solution, but the long-term application lifecycle costs are often much greater in the end.

HPE ConvergedSystem 700 and HPE Converged Architecture 700 with HPE StoreOnce and HPE Data Protector provide an excellent DR solution for an Oracle Database. HPE ConvergedSystem 700 and HPE Converged Architecture 700, with their seamless computing, virtualization, and mobilization, are well suited for the Oracle Database application. Each system ships with factory-integrated compute, storage, networking, and management components, all preconfigured to address the most demanding workloads. Converged management and automation capabilities are built into the HPE ConvergedSystem 700 and HPE Converged Architecture 700 that allow customers to simplify everyday tasks through a single pane of glass management. Compute and storage can be scaled independently so that customers can easily adapt to new requirements and enable the solution to grow as business needs change. One-stop support from the HPE Center of Excellence (CoE) provides a single point of accountability with a faster problem resolution.

This paper will examine disaster recovery for both mission-critical applications and generalized virtual machine workloads. Oracle RAC managed with the Oracle Database application agent represents the mission-critical application. Recovery of generalized virtual machine workloads is accomplished with HPE Data Protector. The HPE Data Protector Cell Manager failover feature makes the Cell Manager available at the secondary site during disaster. The HPE Data Protector Object Copy feature enables HPE StoreOnce replication for secondary site data availability. All of these advanced data management features can be supported on HPE StoreOnce.

This configuration was tested and validated on HPE ConvergedSystem 700. The solution is also applicable on HPE Converged Architecture 700 and HPE Converged Infrastructure platforms as well.

Target audience: This document is intended for data center administrators, managers, and staff who want to learn more about disaster recovery solutions using an HPE ConvergedSystem 700 deployment. A working knowledge of Linux®, Oracle Databases, HPE Data Protector, and virtualization is recommended.

Document purpose: The purpose of this document is to describe a Reference Architecture, highlighting recognizable benefits to technical audiences.

Solution overview

A well-designed DR strategy does not use a one-size-fits-all approach. A combination of tactics, customized for the target applications, provides the best disaster protection and Service Level Agreement adherence. Following that principle, the environment described below takes different approaches for critical and generalized workloads. Critical applications that need deep integration with the data management system, such as Oracle Database, are protected using an Oracle Database application agent specific to the application. For generalized workloads, the Virtual Environment Protection Agent (VEPA) and the Media Agent provide a data management vehicle for all virtual machines.

A key consideration for DR is data availability. Several approaches are taken to ensure data is available at the DR site. The HPE Data Protector foreign Cell Manager[1] configuration ensures all the backup, restore and replication configuration is available at the standby Cell Manager. The HPE Data Protector Object Copy feature makes the application data available at the secondary site for DR, and the HPE Data Protector integration with VMware® is used for VM recovery at the secondary site.

Figure 1 shows the two HPE ConvergedSystem 700 sites that are used for testing along with HPE StoreOnce 4500 as the backup device. We have used two Cell Manager servers, one at the primary and another at the secondary site. The Cell Manager server communicates with all clients, Disk Agents and Media Agents and coordinates all operations, such as backup, restore, copy and media management. All the Oracle RAC VMs and the single instance VMs are configured to act as Cell Manager Clients and Media Agents. In the Primary Cell Manager we have imported a secondary Cell Manager as foreign Cell Manager and vice versa. Therefore, during replication, you can specify the foreign Cell Manager to be aware of the replicated copy.

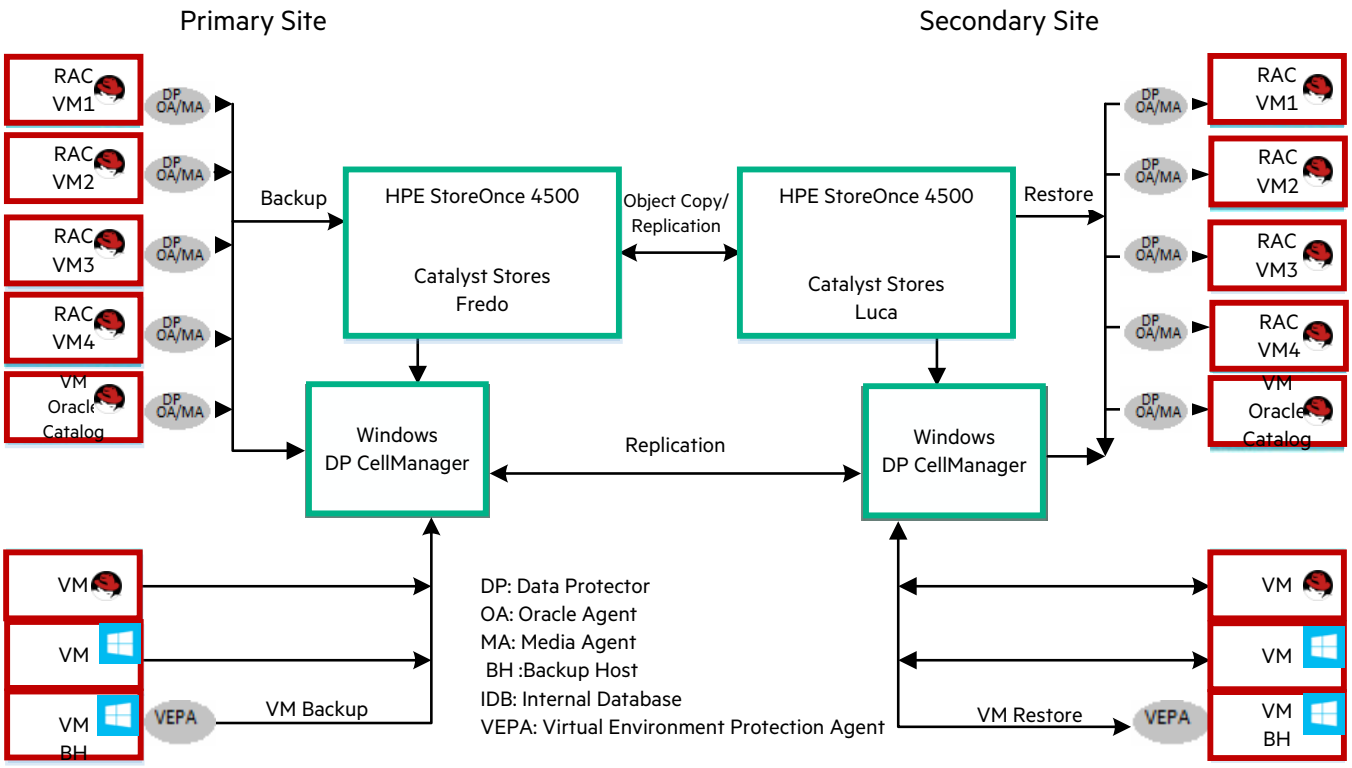

Figure 1. Overview of the solution

Using an Object Copy, we have replicated all of the backup images to the secondary site's HPE StoreOnce and then restored the data from either local or remote backup copies on HPE StoreOnce.

[1] Foreign Cell Manager is a cell manager configured at the secondary site for Disaster Recovery purpose

To protect VM images:

- **HPE Data Protector** uses VMware vStorage API for Data Protection (VADP) to automatically discover VMs
- **HPE StoreOnce Catalyst** software deduplicates VM backup images at the backup host
- **HPE Data Protector** copies the VM backup images to an HPE StoreOnce Catalyst store
- **Object Copy** is used to backup images that are replicated to the secondary site, and the Cell Manager at the secondary site is aware of the replicated data

Solution components

This solution includes the following key components.

Hardware

HPE ConvergedSystem 700

New workloads and business demands are forcing customers to reevaluate the way they purchase and manage infrastructure. Do-it-yourself and integrated systems assembled from components provided by multiple vendors often cause increases in cost, time, and effort as well as create complexity with interoperability challenges. HPE ConvergedSystem 700 addresses these new requirements with a tightly integrated, workload optimized system architected to support on-demand IT infrastructure, private cloud, and mixed workloads. Preconfigured to meet a range of business needs, HPE ConvergedSystem 700 offerings can be deployed easily and rapidly to support a variety of virtualized application environments. HPE engineering has done all the testing and validation so that IT professionals can focus their time on critical activities. Specific technology has been included in the system that addresses critical business objectives such as lower costs, increased performance, and high availability.

HPE ConvergedSystem 700 is extremely agile, with a broad range of configurations available. Based on a common, scale-out architecture, it starts with as few as two workload servers and can scale up to 128 servers with multiple storage options. Compute and storage can be customized and scaled individually, providing the ability to deliver the right mix of performance for unique customer requirements.

For more details, see http://h20195.www2.hpe.com/V2/GetDocument.aspx?docname=4AA5-6612ENW.

HPE Converged Architecture 700

The HPE Converged Architecture 700 is delivered as a flexible and validated Reference Architecture. Hewlett Packard Enterprise has done all the engineering, testing, and design to save your time and effort. It provides the configuration, sizing, bill of materials, and deployment details that reduce your risk and TCO. Based on a full HPE stack of HPE BladeSystem, HPE 3PAR storage, HPE Networking, and HPE OneView infrastructure management software, the HPE Converged Architecture 700 can be easily modified to fit within your existing networking requirements. It will accommodate the tools you use and support multigenerational HPE ProLiant blade servers, including HPE ProLiant Gen9. This architecture delivers interoperability with a choice of third party top-of-rack network switches and hypervisors from leading industry vendors. The HPE Converged Architecture 700 provides an easy upgrade path to HPE Helion CloudSystem. You can deploy at your site with your trusted channel partner who can provide value added services tailored to fit your needs. This solution is complete with a variety of lifecycle support options.

For more details, see: http://h20195.www2.hpe.com/V2/GetDocument.aspx?docname=4AA6-1963ENW.

HPE 3PAR StoreServ

HPE 3PAR StoreServ storage offers the performance and flexibility that you need to accelerate new application deployments. Including server virtualization, cloud, IT as a Service (ITaaS). With HPE 3PAR StoreServ storage, you can serve unpredictable and mixed workloads, support unstructured and structured data growth, and meet both file and block storage needs.

The modular HPE 3PAR StoreServ architecture can be scaled from 1.2 TB to 1.2 PB, making the system deployable as a small or very large centralized system. Until now, enterprise customers were often required to purchase and manage at least two distinct architectures to span their range of cost and scalability requirements. HPE 3PAR StoreServ storage is the ideal platform for virtualization and cloud computing environments. The high performance and scalability of the HPE 3PAR StoreServ architecture is well-suited for large or high-growth projects and consolidation of mission-critical information, demanding performance-based applications and data lifecycle management.

For more details on product ranges, see hpe.com/storage.

HPE StoreOnce

HPE StoreOnce is designed to cater to the needs of all types of customers from entry-level to large-scale enterprises. HPE StoreOnce backup systems deliver scale-out capacity and performance to keep pace with shrinking backup windows, reliable DR, simplified protection of remote offices, and rapid file restore to meet today's SLAs. The models vary by capacity and connectivity protocol. Customers can start out by purchasing a single HPE StoreOnce base unit/couplet, and then expand with additional couplets and expansion shelves. These units deliver cost-effective, scalable, disk-based backup with deduplication for long-term, on-site data retention for the HPE ConvergedSystem 700 single-rack configuration.

For more details, see https://www.hpe.com/h20195/v2/GetDocument.aspx?docname=c04328820.

Software

HPE Data Protector

HPE Data Protector is a backup solution that provides reliable data protection and high availability for rapid growth. HPE Data Protector offers comprehensive backup and restore functionality specifically tailored for enterprise-wide and distributed environments. HPE Data Protector is a scalable and highly flexible architecture, offering central administration and a high-performance backup.

The HPE Data Protector cell is a network environment that has a cell manager, client systems, and devices. The HPE Data Protector Cell Manager is the central control point where the HPE Data Protector software is installed. After installing the HPE Data Protector software, you can add systems to be backed up. These systems become HPE Data Protector client systems that are part of the cell. When the HPE Data Protector backs up files, it saves them to media in the backup devices. Figure 2 explains about the HPE Data Protector architecture.

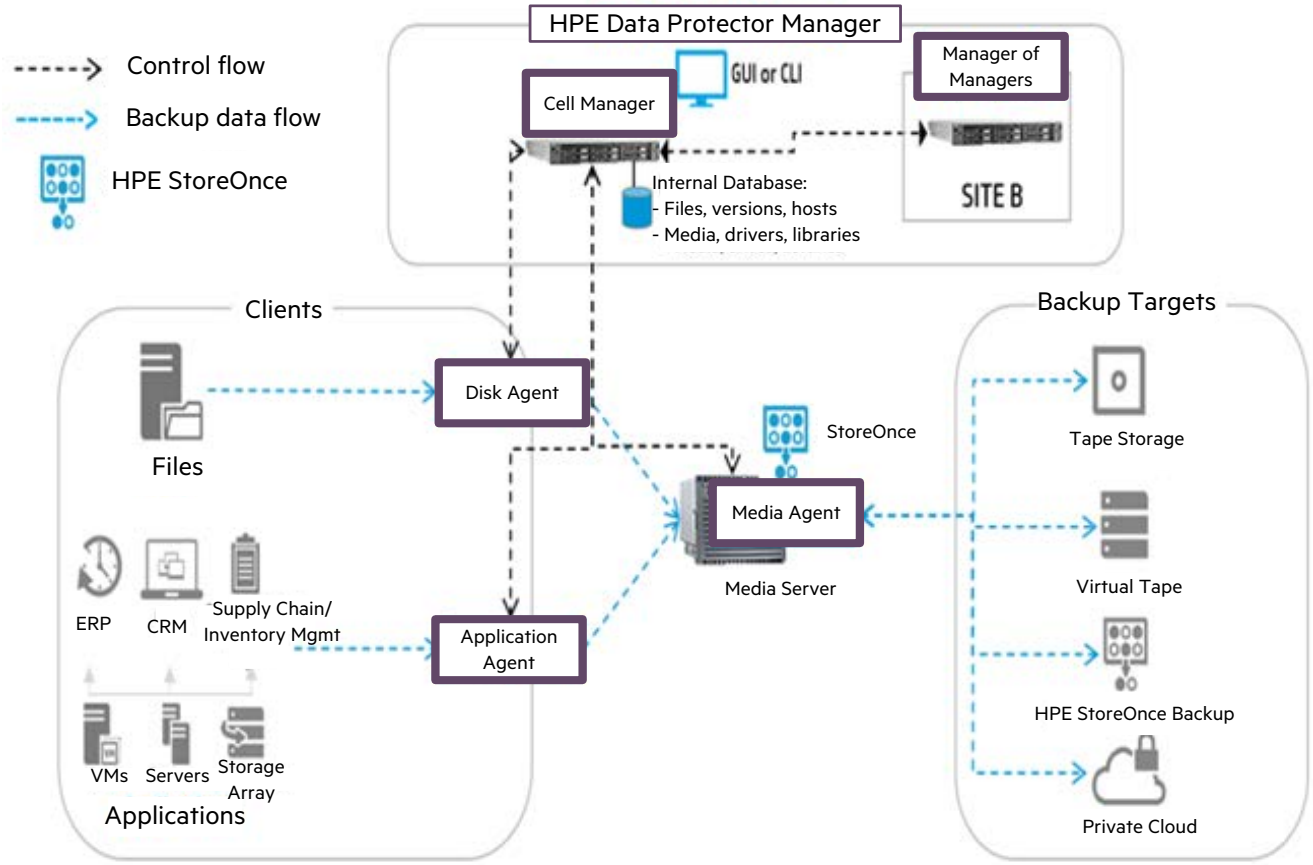

Figure 2. HPE Data Protector Architecture

For more details about HPE Data Protector, see http://www8.hp.com/us/en/software-solutions/data-protector-backup-recovery-software.

Test bed architecture overview

The test bed infrastructure consists of an HPE ConvergedSystem 700 and an HPE StoreOnce 4500 at each site (primary and secondary). Both sites are connected via a 20GbE network. At the primary site, eight VMs are created: four for Oracle RAC, one for Oracle Single Instance database, one for Oracle/catalog database, and two for HammerDB application. Each Oracle VM has the capacity of 16 virtual CPUs (vCPUs) and memory of 96 GB. An additional VM with 16 vCPUs and 32 GB capacity is created for the Cell Manager database server. The storage for all the VMs has been allocated from HPE 3PAR StoreServ 7400c and a datastore was created. We have created a similar virtual machine infrastructure at the secondary site also for the disaster recovery purpose.

In the secondary site, we have created eight virtual machines (RAC VMs, Single Instance, Catalog DB, HammerDB, and Cell Manager database server) and installed the necessary operating systems. Oracle Database binaries have been installed on the RAC setup to speed up the recovery process, thereby reducing the RTO. To better utilize the infrastructure in the secondary site, it can be used to run tests and development databases or any other applications.

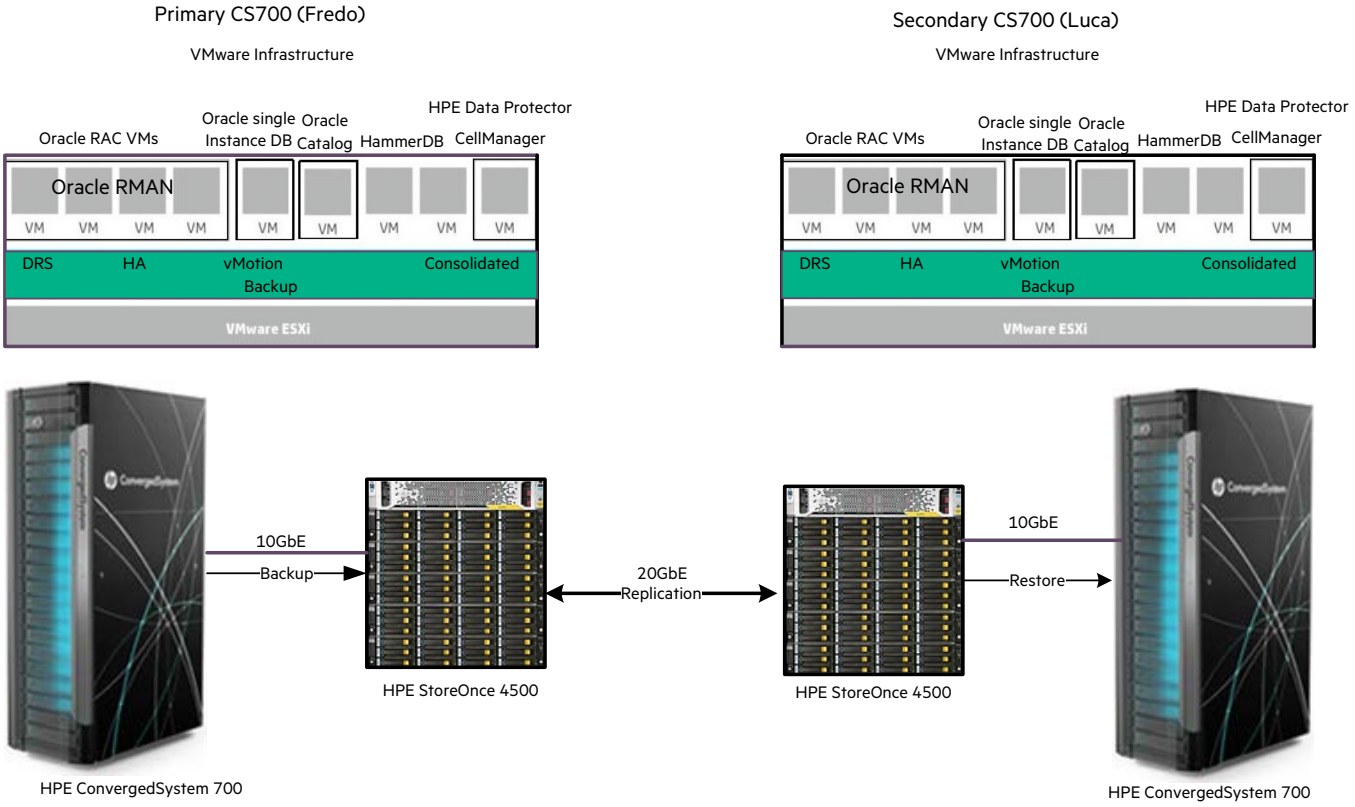

Figure 3. HPE ConvergedSystem 700, Oracle RAC, and HPE Data Protector DR Test Bed architecture overview

Table 1 discusses the test bed configuration of the Disaster recovery for Oracle Database for a single site.

Table 1. HPE ConvergedSystem 700 hardware for the test bed

One HPE BladeSystem c7000 enclosure	Enclosure to host blades and Virtual Connect modules
Two HPE Virtual Connect FlexFabric 10Gb/24-Port modules	Virtual Connect module for Ethernet and SAN connectivity
Eight HPE ProLiant BL460c Gen9 server blades	Server blades to host one Oracle RAC 12c 4-node cluster, two single Instance Oracle Database, HammerDB application running as VMs and two VMs for Live Sync testing
One HPE 3PAR StoreServ 7400c storage	Storage for Oracle Database and data stores for VM
Two HPE 5900AF-48XG-4QSFP+ network switches Two HPE 5900AF-48G-4XG-2QSFP+ network switches	Network switches for LAN traffic. Network switches for management traffic
Two HPE StoreFabric SN6000B 24-port SAN switches	Fibre Channel switches for SAN connectivity between servers and HPE 3PAR StoreServ storage
Software configurations	Red Hat® Enterprise Linux (RHEL) 6.5, Microsoft® Windows Server® 2012 R2, Oracle 12c, HPE Data Protector 9.05, HammerDB 2.16
One HPE StoreOnce 4500	Target for Oracle Database backup

Storage configuration for the test bed

The HPE 3PAR StoreServ 7400c hosted the two Oracle Databases of 500 GB each. To support this layout, the HPE 3PAR StoreServ primary storage array of the HPE ConvergedSystem 700 is used to allocate and present storage pools . After allocation, the storage objects are presented to the VMware vSphere environment. The storage is allocated and provisioned using a thick format to enhance I/O performance during the workload-testing phase of this solution.

Table 2. Storage configuration for the test bed

Contents	Size	Provisioning	Disk type	Notes
Oracle Data Files	500 GB	Thick	15k SAS	RAID 5
Oracle Temp Files	500 GB	Thick	15k SAS	RAID 5
Oracle Redo logs	500 GB	Thick	15k SAS	RAID 5
Oracle Control Files	500 GB	Thick	15k SAS	RAID 5
Virtual Machine OS, Binaries specifications	150 GB	Thin	15k SAS	RHEL 6.5, Oracle 12c (16 core and 96 GB memory)
Cell Manager Specifications	200 GB	Thin	15k SAS	Windows Server 2012 R2, HPE Data Protector 9.05
HammerDB	50 GB	Thin	15k SAS	RHEL 6.5, HammerDB 2.16

HammerDB

HammerDB is an open source database load testing and benchmarking tool for Oracle and other databases.

HammerDB is used to load the data into the Oracle Databases in the primary site. Transactions were run against the databases during backup to simulate a customer production environment. The load on the system was qualified as medium and consumed 45% of the CPU. The daily data change rate was estimated between 1.8% and 2%. After the DR procedure was completed at the secondary site, HammerDB was used again to verify the database consistency by running the transactions against the DB.

- For more details, see hammerdb.com.

Disaster Recovery for Oracle Database using HPE Data Protector Object Copy

The following steps are required to implement disaster recovery for Oracle Databases:

1. Import the secondary site Cell Manager as a client in the primary Cell Manager.
2. Backup the catalog database and restore it at the secondary site.

Note

The catalog database is used by rman[2] to store metadata information for the backup. Metadata information like the number of backups, time of the backup, and other detail can be stored. In our setup, the catalog database is running on a separate VM.

---

3. Backup the production Oracle Database and restore it at the secondary site. This is the actual database which contains the production data to be protected.
4. Take a file level backup of all the configuration files of HPE Data Protector from primary site and restore it on the secondary site. HPE Data Protector expects the configuration in the primary and secondary cell to be identical to detect the backup jobs on the secondary site cell manager.
5. Take an Object Copy of the backup (catalog, production and config file backup) to the secondary site store.
6. Restore the config files, catalog database and Oracle Database in order.

The next section explains the detailed procedure for the disaster recovery for Oracle Database using HPE Data Protector Object Copy method.

1. To make the secondary site disaster tolerant, first import the secondary Cell Manager to the Primary Cell Manager as a Client as shown in Figure 4.

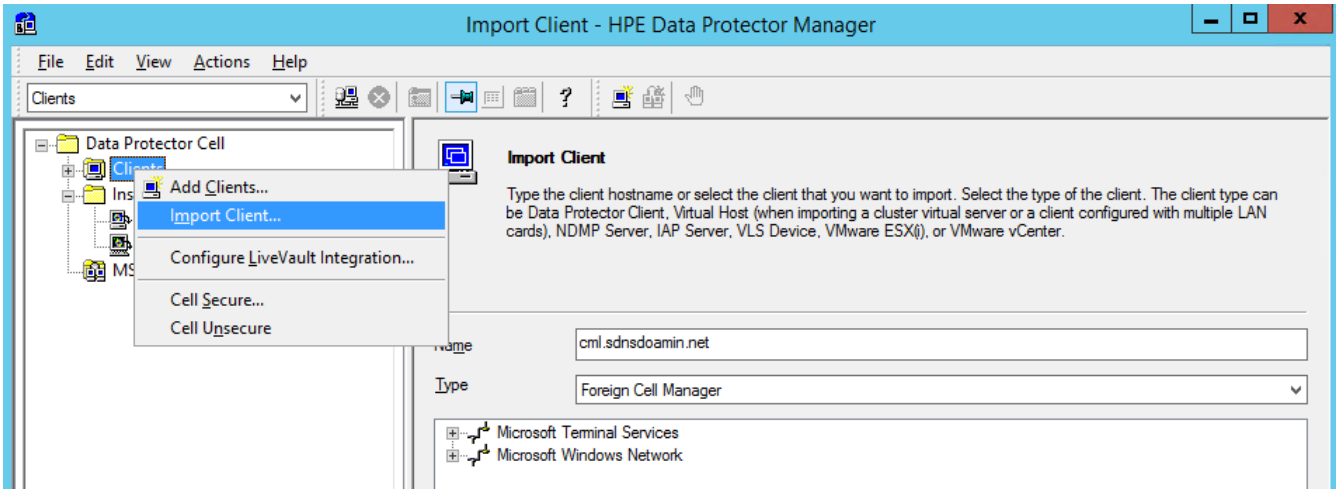

Figure 4. Importing Foreign Cell Manager as a client.

2. Enter the Name of the foreign/secondary Cell Manager.
3. Select Foreign Cell Manager as the client Type, and click Finish. The Foreign Cell Manager will be added as the Client.

[2] Oracle Recovery Manager used for Oracle Backup, restore and Recovery purpose

Configuring HPE Data Protector with Oracle/catalog database

To take backup of Oracle Database using HPE Data Protector, we need to configure the Oracle/catalog database with HPE Data Protector.

Below are the steps to configure Oracle database with HPE Data Protector

1. In the HPE Data Protector Manager GUI, select **Backup → Oracle Server → Add Backup** to open the Create New Backup window.
2. In this window, select **Blank Oracle Backup**.
3. In the Backup type drop-down list, select **Local or network backup**.
4. By default, this will be a load-balanced backup. Click **OK**.

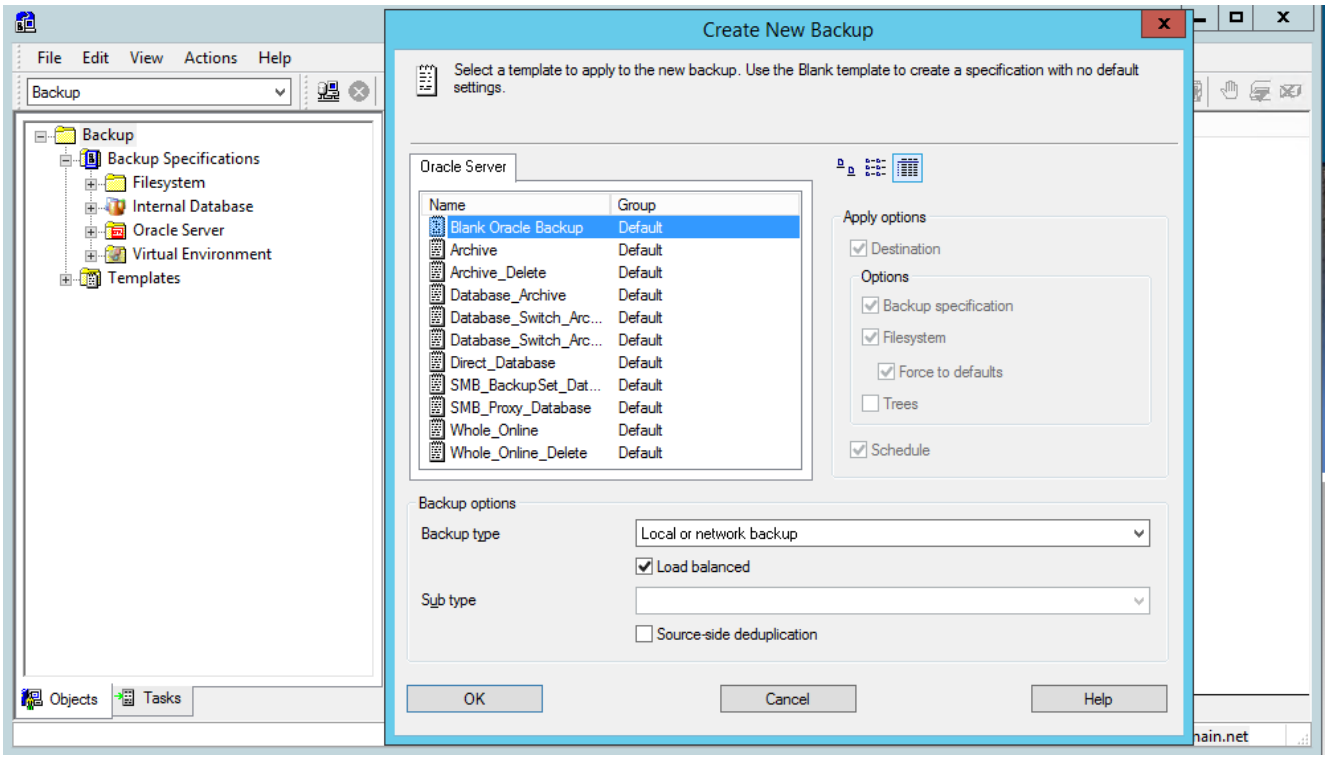

Figure 5. HPE Data Protector creating Oracle Backup specification.

5. In the Client drop-down list, select the scan name of the database to load balance the backup operation.
6. Choose the Application database name, and provide the user and group/domain credentials. Click **Next**.

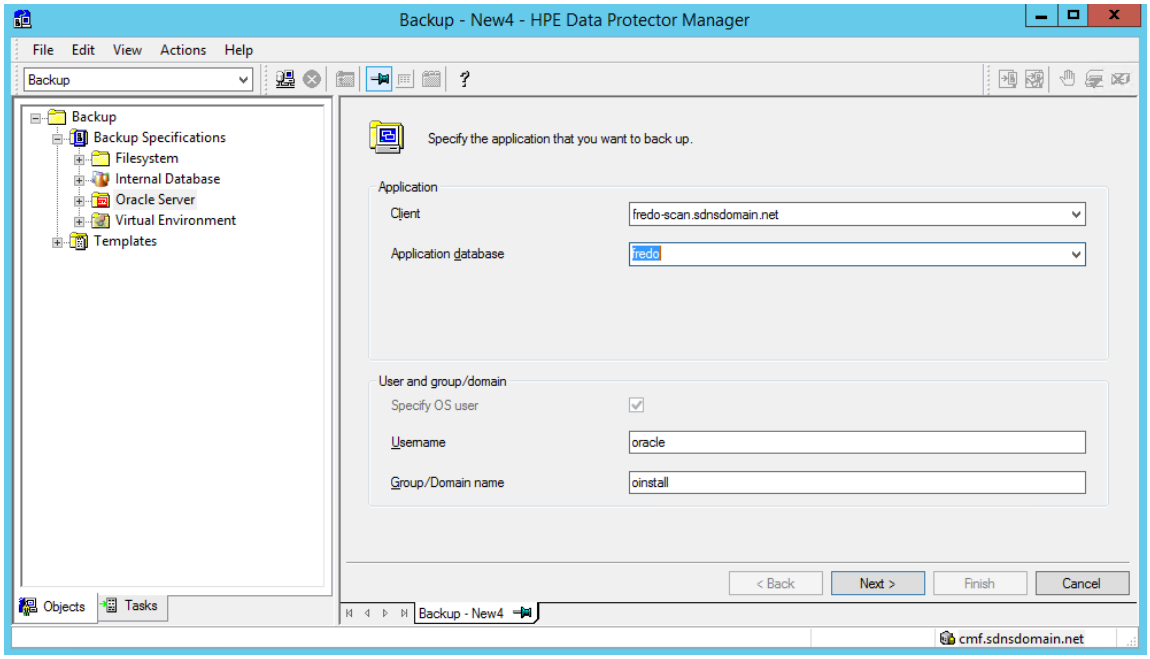

Figure 6. HPE Data Protector – specifying the application details

7. Next right-click the database name to provide the Oracle General, Primary, Catalog, and Standby parameters.

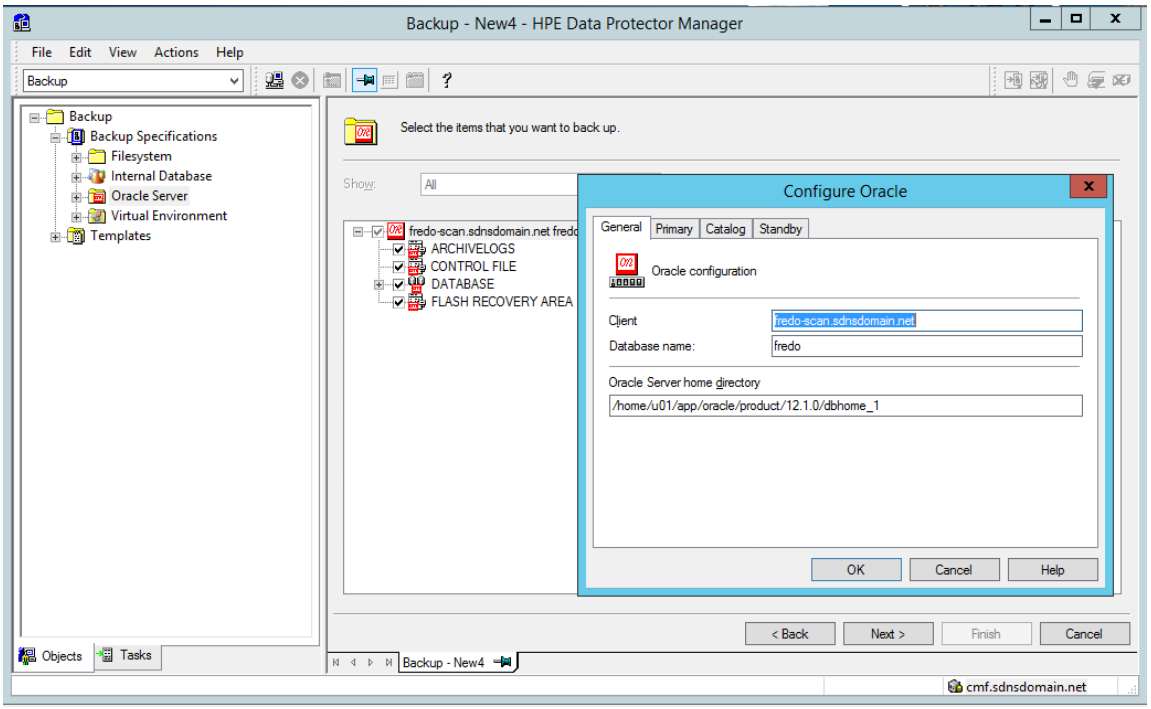

Figure 7. HPE Data Protector – Configuring Oracle home directory.

8. In the next screen, select the destination device to take the backup, and also provide the application-specific configuration parameters.
9. Save or execute the created backup specification.

Oracle/catalog database backup procedure

For the test configuration, full and incremental backups were scheduled to run every 24 hours.

After the HPE Data Protector and Oracle integration are completed as explained in the previous sections, we can use the procedure below to back up the Oracle Database. Make sure that both the Oracle Database and the catalog database are backed up using this procedure.

1. Go to the **HPE Data Protector Manager**, as seen in figure 8.
2. Select **Backup** from the drop-down box.
3. Under Oracle Server, right-click on the respective backup specification, and click **Start Backup** as shown in Figure 8.

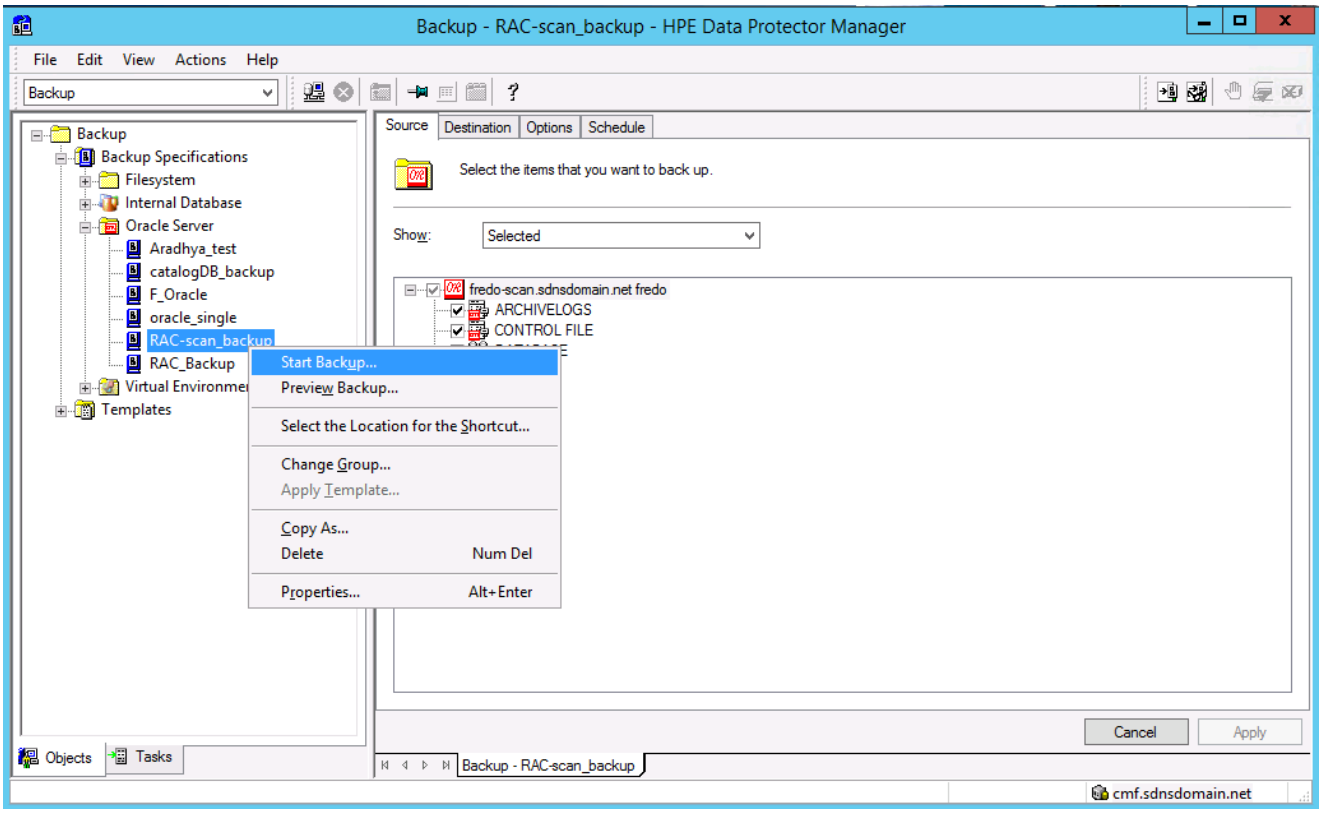

Figure 8. Oracle RAC Subclient backup option for Fredo Instance

The backup will be initiated, and you can monitor it on the Backup tab in the DP Cell Manager window. You need to backup both catalog and Oracle Database using above procedure.

HPE Data Protector Configuration files backup procedure

Once the database backup has completed, configure and run a file level backup for the configuration files of HPE Data Protector as shown in Figure 9.

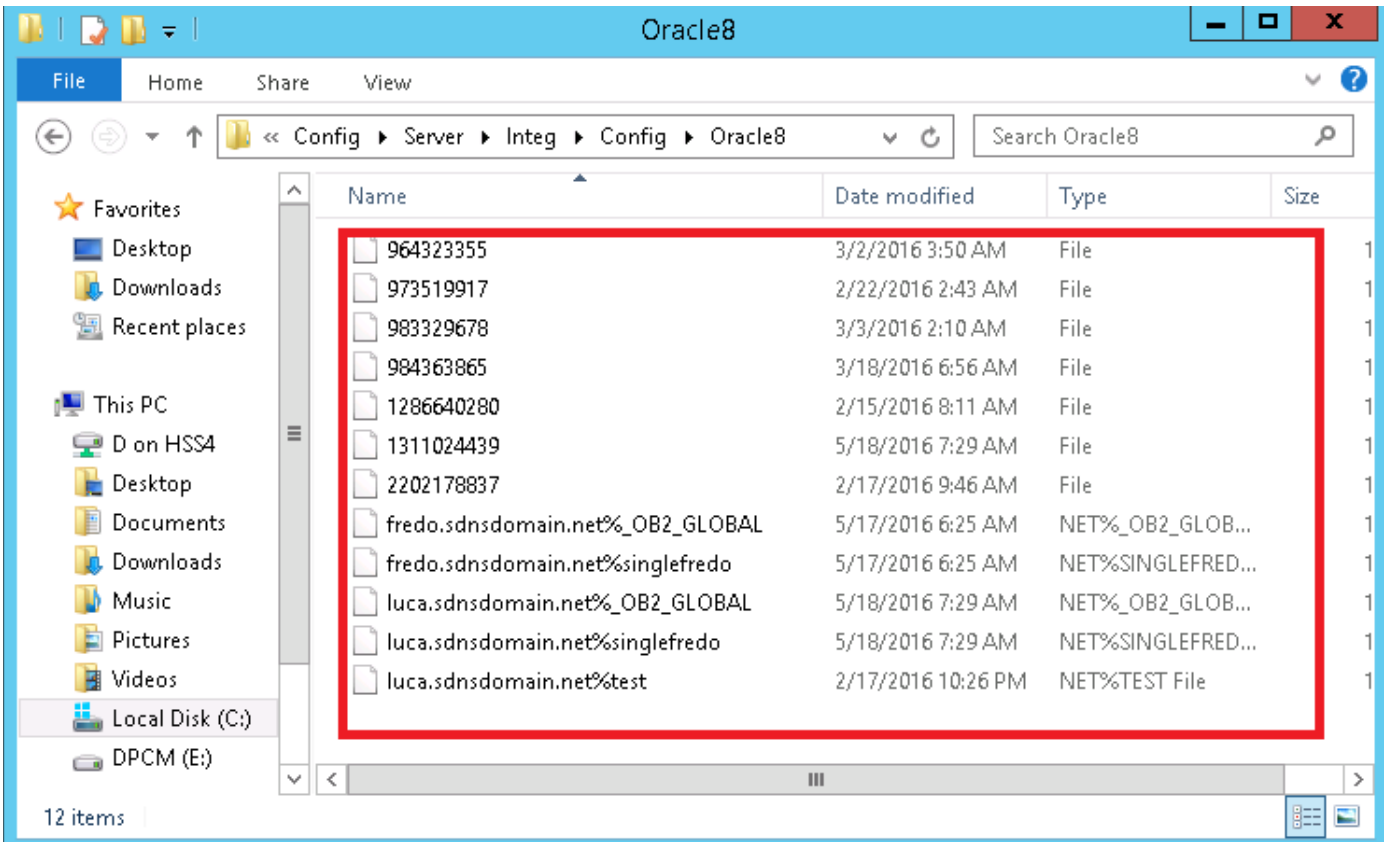

Figure 9. Copying the config files from primary Cell Manager to secondary

As soon as the backup operation is complete, we schedule the post-backup Object Copy operation to replicate the backup images of all the three (catalog, oracle and HPE DP config files) as explained in below section.

Creating the Object Copy specification for backup jobs

HPE Data Protector Object Copy

The HPE Data Protector's Object Copy functionality enables you to copy selected object versions to a specific media set. You can select object versions from one or several backups, Object Copy, or object consolidation sessions. During the Object Copy session, HPE Data Protector reads the backed-up data from the source media, transfers the data, and writes it to the target media.

The result of an Object Copy session is a media set that contains copies of the object versions you specified. You can start an Object Copy session interactively or specify an automated start of the session.

The Object Copy is load balanced by default. HPE Data Protector makes optimum use of the available devices by utilizing as many devices as possible.

1. Go to the **HPE Data Protector Manager**, as seen in figure 10.
2. Select **Object Operations** from the drop-down box in the top left corner.
3. Click **Object Copy → Post Backup**.
4. Select the backup specification that needs to be replicated to the secondary site post backup.

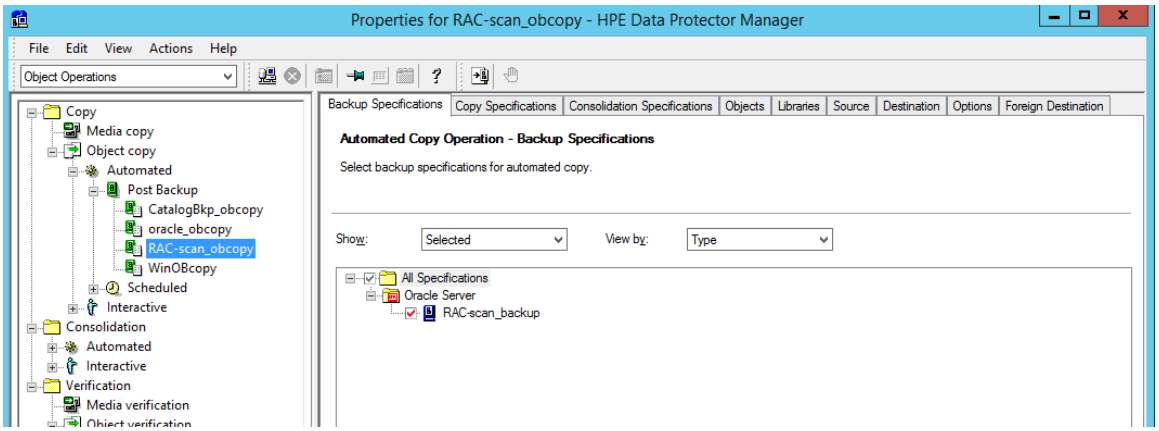

Figure 10. HPE Data Protector Object copy operation- Backup Specifications tab

Provide the necessary information on all the tabs.

1. On the Destination tab, select the destination HPE StoreOnce Catalyst store for the copy operation as shown in Figure 11.

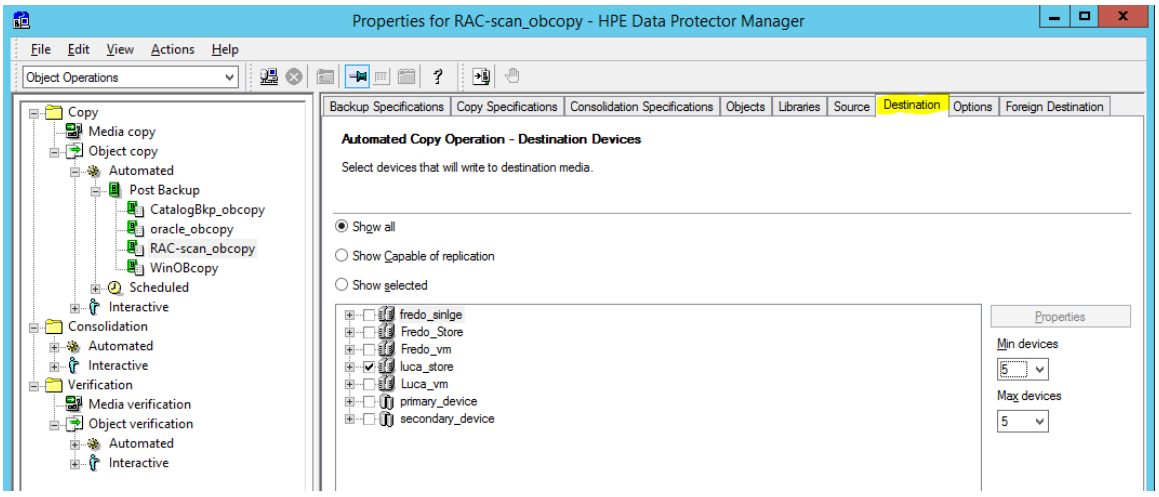

Figure 11. HPE Data Protector Object Copy operation - Destination tab

2. On the Options tab, select the **Use replication** and **Replicate to a foreign cell** checkboxes to move the data.

As you can see in figure 12, the foreign Cell Manager is made aware that the replicated image is available for restore/recovery at the secondary site.

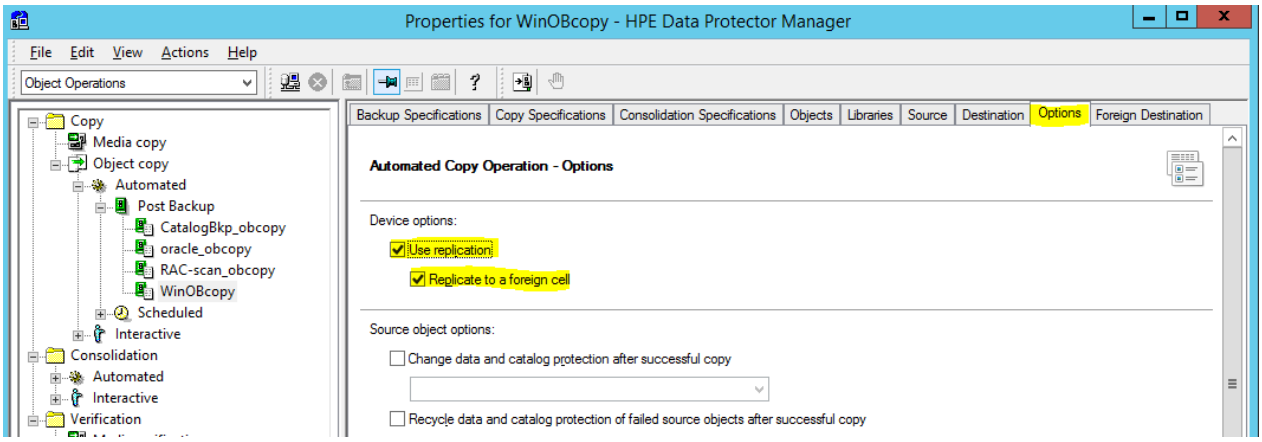

Figure 12. HPE Data Protector Object Copy operation - Options tab

3. On the Foreign Destination tab, select the **Foreign Cell Destination**.
4. Here you choose the foreign Cell Manager that you imported to the primary cell earlier.
5. Select the schedule to run the object specification and save the specification.

Note

Object Copy details are stored in the internal database (IDB) located on both the primary and secondary Cell Managers. It is vital to select the **Replicate to a foreign cell** option so that in the event of a disaster, the data/database can be restored to a new server from the secondary or foreign Cell Manager.

---

6. After the successful execution of the backup and the Object Copy operation at the primary site, go to the secondary site HPE Data Protector GUI and navigate to Restore, from the drop-down menu (Figure 13).
7. You can see the database backup image to restore.
8. Select the database.
9. On the Options tab, select the target host, enter the Oracle user and group name, and click Restore.

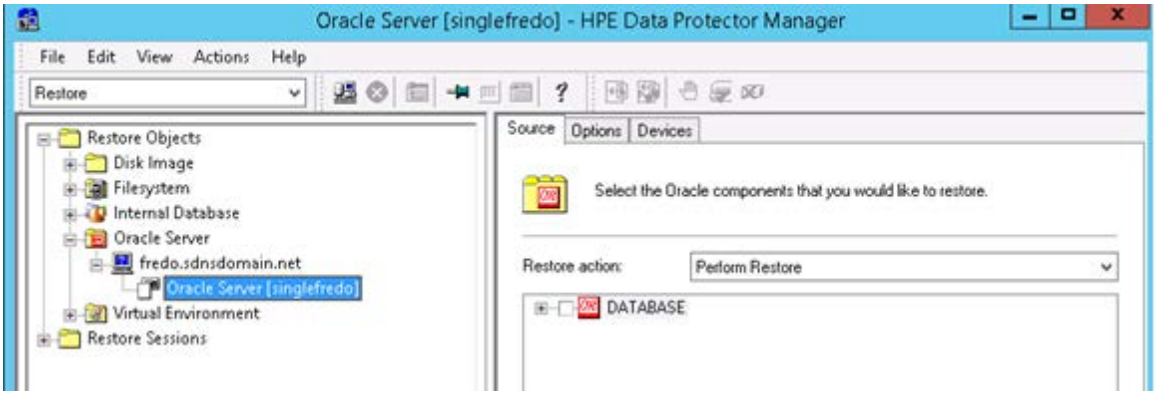

Figure 13. Replicated backup images on the Oracle RAC clients

DR on a secondary site

We have imported the Secondary site Data Protector Cell Manager as a foreign Cell Manager in the primary Cell Manager. Also while configuring Object Copies for the backup specification, we have mentioned to make the secondary Cell Manager aware of the backup images so that the secondary site Cell Manager will automatically detect the backup images for the restore.

Oracle failover scenario

To restore the Oracle Database on secondary site we need to follow below steps.

1. Restore the HPE Data Protector Configuration files.
2. Restore the catalog database.
3. Restore the production database using the catalog database.
4. Register the production database with the Oracle RAC cluster.

Restore the HPE Data Protector Configuration files

As mentioned in the backup section, we have taken the file level backup of HPE Data Protector Configuration files. Restore those files to same location at secondary site and restart the cell manager.

Restore the catalog database

Before restoring the actual database, we need to restore the catalog database. Catalog database is used by rman to store metadata information about the backup. Information like number of backups, time of the backup, and other details has been stored in it. In our setup, the catalog database is running on a separate VM, and we need to make a backup of it using the backup procedure mentioned in Oracle backup procedure section.

1. To restore the catalog database on secondary site, install the Oracle Database binary on the catalog VM, and start the dummy Oracle instance with dpcat name using “startup nomount force “from the SQL command prompt.
2. Then follow below steps to restore spfile using below rman script.

```
run{
allocate channel 'dev_0' type 'sbt_tape'
parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8_64bit.so,ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<instance
name>)';
restore spfile from autobackup;
restore controlfile from autobackup;
RELEASE CHANNEL dev_0;
}
```

3. Once spfile has been restored, restart the database using the spfile. Then restore the controlfile using below rman script.

```
run{
allocate channel 'dev_0' type 'sbt_tape'
parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8_64bit.so,ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<instance
name>)';
restore controlfile from autobackup;
RELEASE CHANNEL dev_0;
}
```

4. After that, restore the datafiles and open the database using below rman script.

```
run{
allocate channel 'dev_0' type 'sbt_tape'
parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8_64bit.so,ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<instance
name>)';
restore database;
recover database;
alter database open resetlogs;
RELEASE CHANNEL dev_0;
}
```

Alternatively you can also restore the datafiles from the GUI of HPE Data Protector as explained in below section.

Oracle Database restore

After the catalog database is restored to the secondary site, we are ready to restore the actual database using the following procedure.

ASM configurations

Because we have used Oracle ASM disks for the database, we need to create the same disk groups (RACDATA,ARCHIVELOGS) of the primary site in the secondary site. We also need to create the directory structure for the Oracle instance using the mkdir command inside the ASM diskgroup as shown in Figure 14.

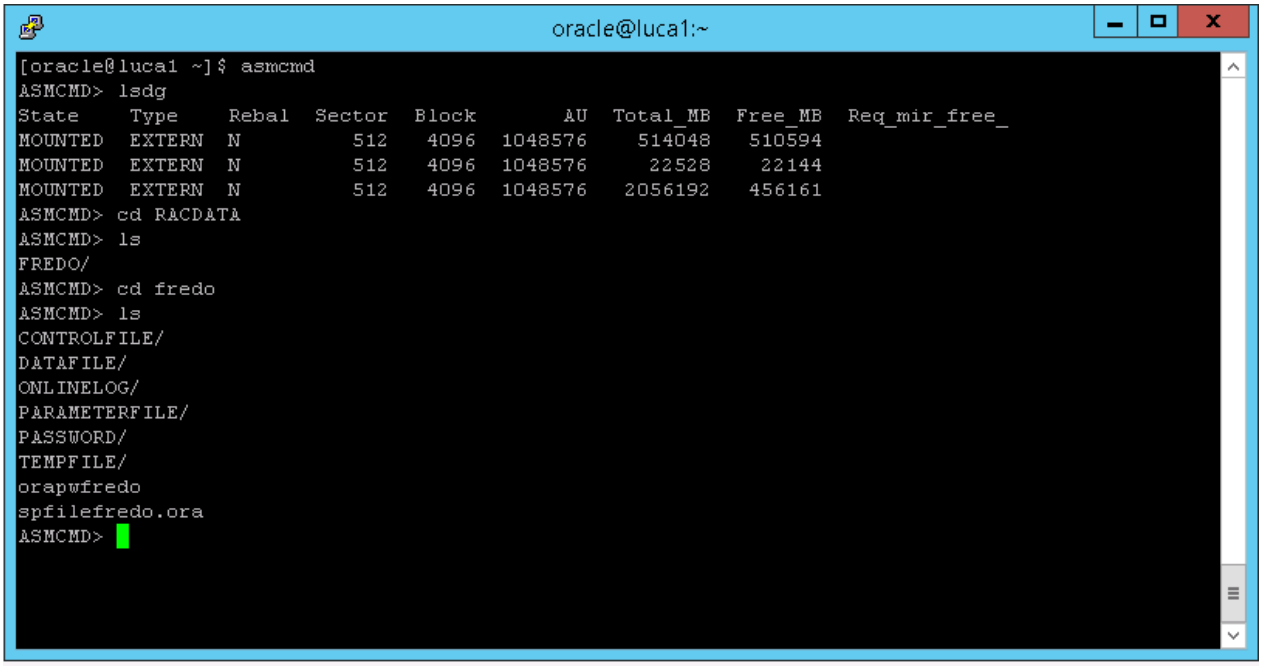

Figure 14. ASM diskgroup directory structure.

How to start the dummy instance on RAC node

1. Restore the oracle init file from filesytem agent to the oracle home directory.
2. Create the destinations and dump files directories based on the init file information.
3. Make appropriate entries in the tnsnames.ora file of the secondary site to resolve the instances names.
4. Create an Oracle password file using the orapwd command, and provide the sys password (which should be the same as in the primary site).
5. Log in at the SQL prompt, and issue the startup `nomount` command.

After the instance is started, the command line output will look similar to Figure 15.

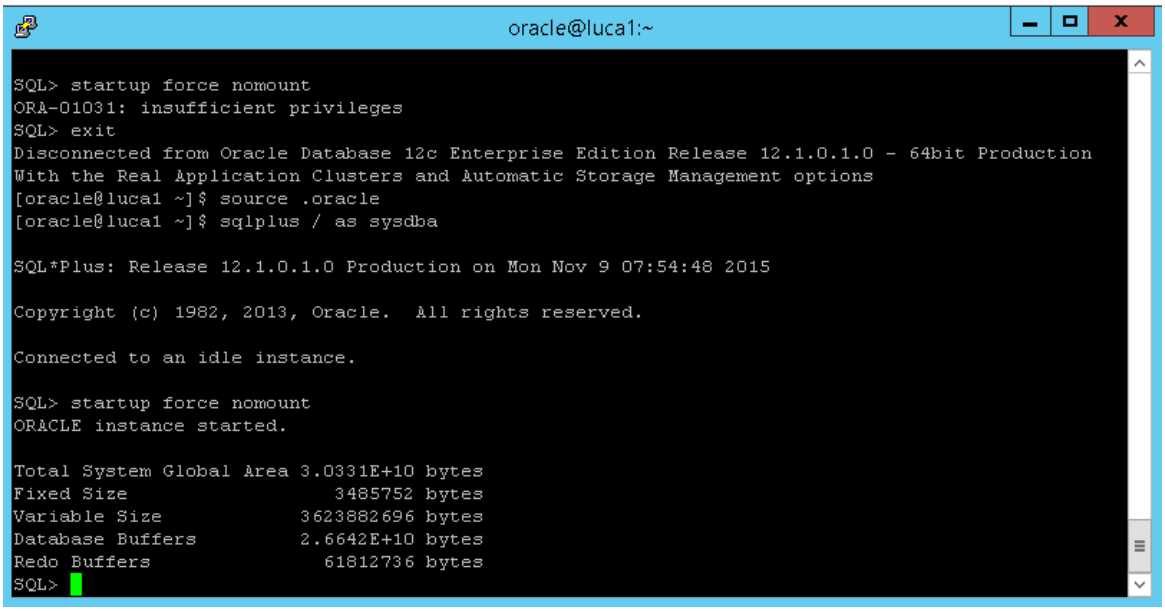

Figure 15. Starting of Dummy instance on RAC node

After the dummy instance is started, use the rman script provided below to restore the spfile and control file to the secondary site.

1. Connect to rman using "rman target /" for the catalog database restore and "rman target / catalog rman/cat@dpcat" for the fredo database.
2. Then restore the spfile and control file using the scripts below.

Note
Change the instance names accordingly in the script.

```
run{
allocate channel 'dev_0' type 'sbt_tape'
parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8_64bit.so,ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<instance
name>)';
restore spfile from autobackup;
RELEASE CHANNEL dev_0;
}
shutdown immediate;
startup force nomount;

run{
allocate channel 'dev_0' type 'sbt_tape'
parms
'SBT_LIBRARY=/opt/omni/lib/libob2oracle8_64bit.so,ENV=(OB2BARTYPE=Oracle8,OB2APPNAME=<instance
name>)';
restore controlfile from autobackup;
RELEASE CHANNEL dev_0;
}
```

3. After the control file is restored, shut down the database and start with the mount option.
4. Go to the secondary Cell Manager, and configure the backup specification for the newly started instance on the secondary site by providing the instance name and the catalog database info as shown in Figure 16.

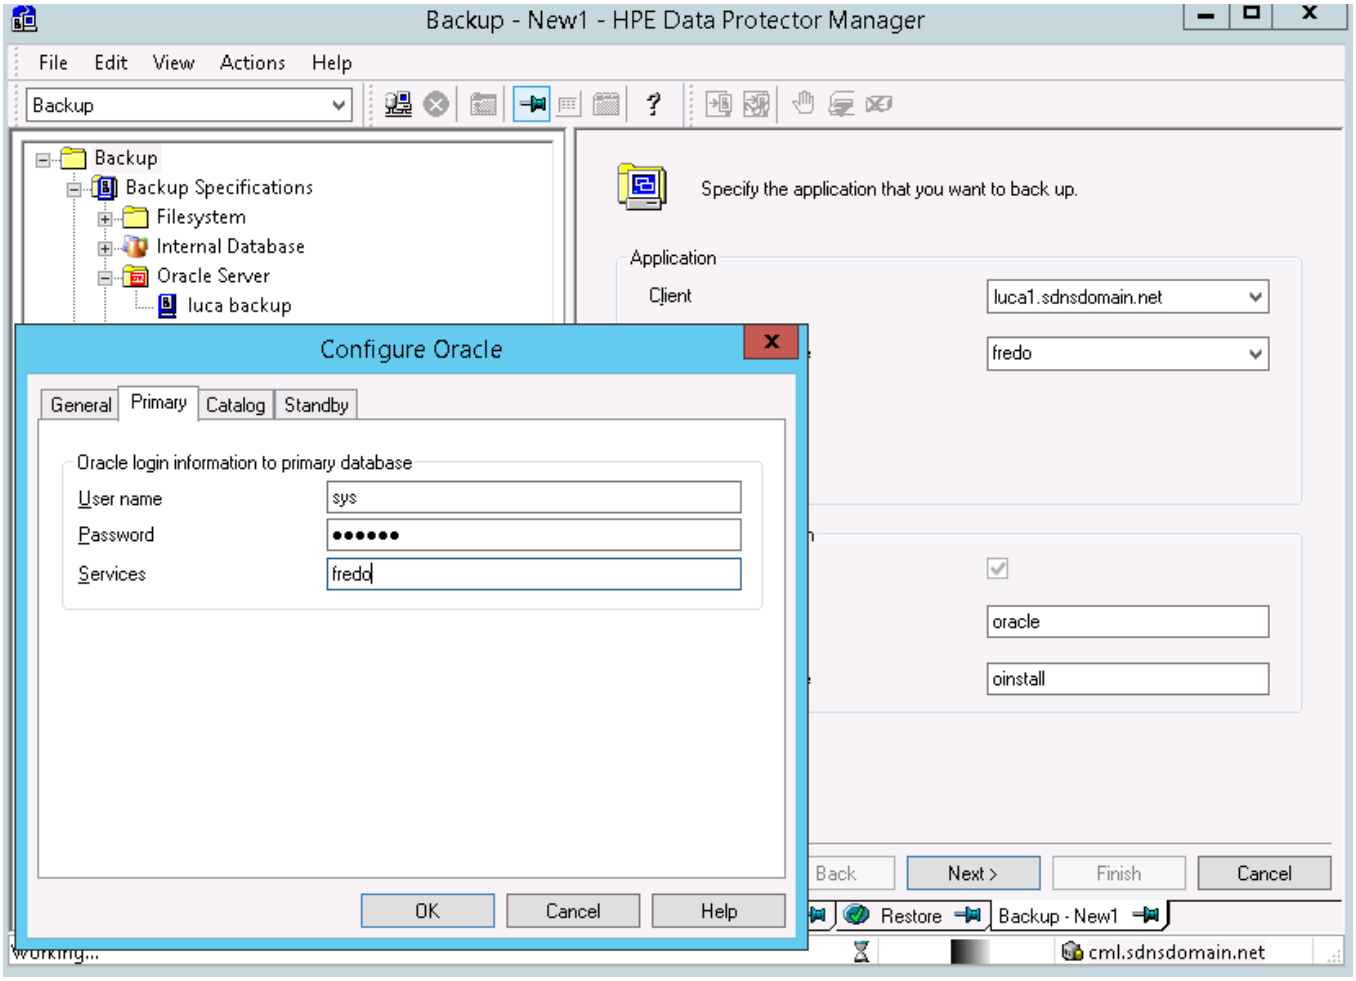

Figure 16. Configuring the backup specification for dummy instance on luca1

5. Go to the Restore option, and click fredo1.sdnsdoman.net.
6. Under the Source tab, select the database files to be restored.

7. Under the Options tab, in the Client drop-down list, select luca1.sdnsdomain.net.
8. In the User name text box, enter “oracle”.
9. In the User group text box, enter “oinstall”, and click the Restore button to restore the database files as shown in Figure 17.

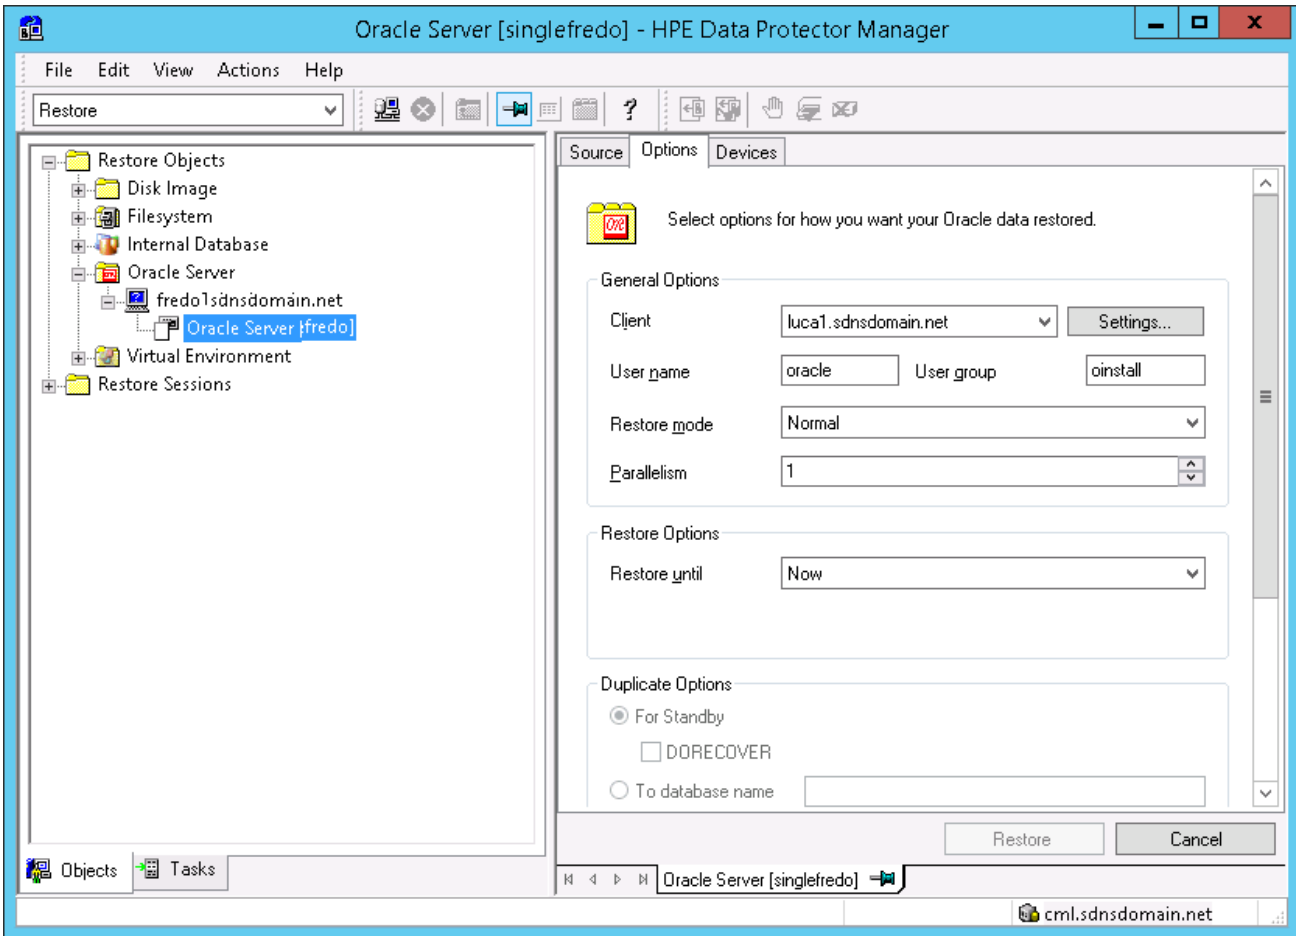

Figure 17. Data files restoration procedure for secondary site

After the database files are restored, we need to restore the archive log files.

1. Go to luca1.sdnsdomain.net and connect to rman.
2. To get the latest backup sets for the database, execute “list backup of database completed between 'SYSDATE -1' and 'sysdate';.”.
3. Make a note of the latest backup set.
4. Using “List of Archived Logs in backup set <Backup set>”, get the latest log sequence number for the backup set.
5. Then go to the HPE Data Protector Cell Manager of the secondary site.
6. Under the Restore drop-down box, click fredo1.sdnsdomain.net. Select the instance, and go to Options.
7. Under Restore Options, in the Restore until drop-down list, click Selected logseq/thread number.

8. Select the log sequence number that has been obtained using rman commands, and click Restore as shown in Figure 18.

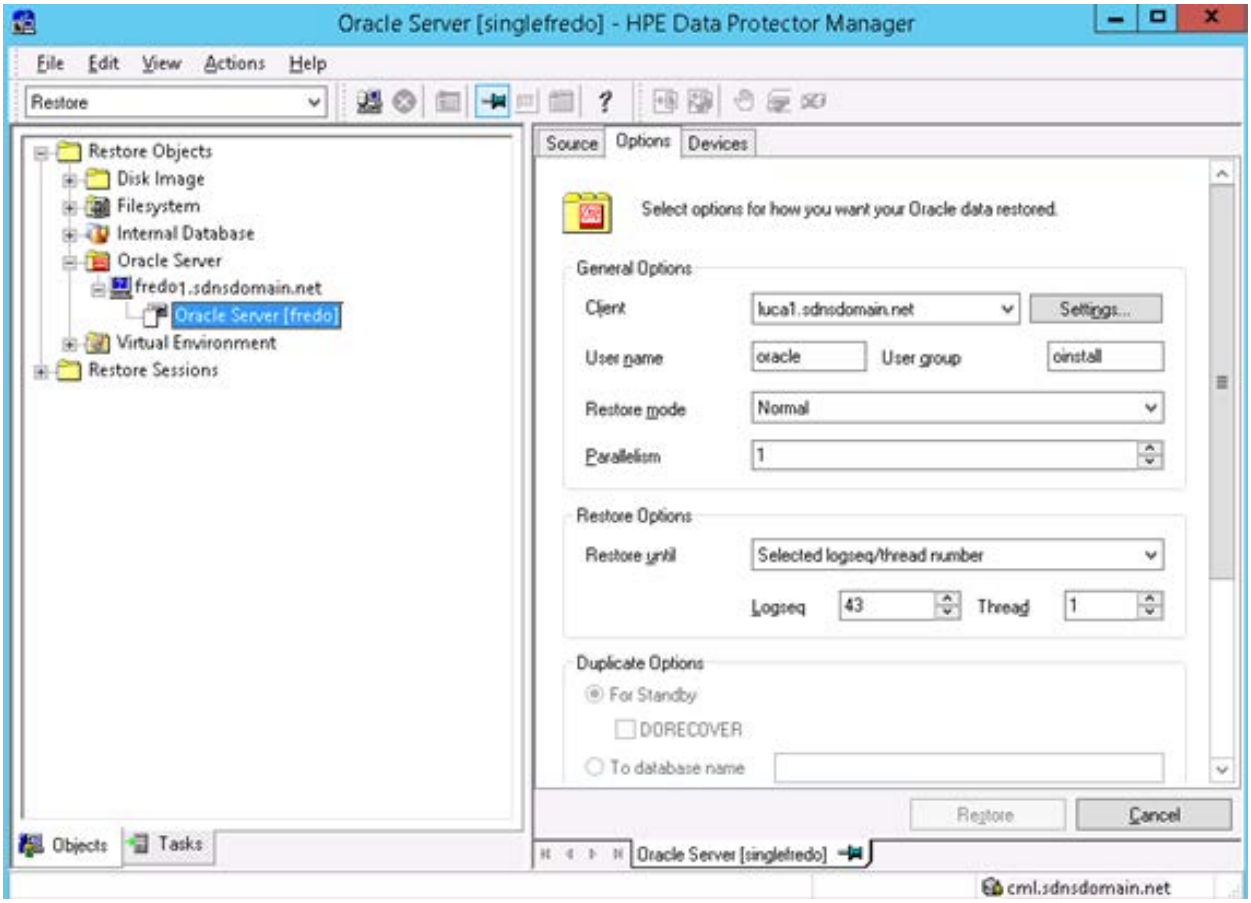

Figure 18. Archive log restoration procedure for luca1.

9. After the database and archive log files are restored successfully, use the following commands to register the database and to add the instance names.

```
srvctl add database -d fredo -o /home/u01/app/oracle/product/12.1.0/dbhome_1/
srvctl add instance -d fredo -i fredo2 -n luca4
srvctl add instance -d fredo -i fredo1 -n luca1
srvctl add instance -d fredo -i fredo4 -n luca2
srvctl add instance -d fredo -i fredo3 -n luca3
```

10. After these commands are successfully executed, start the database using the command below.

```
Srvctl start database -d fredo
```

11. The backup will be taken in the replicated disk created locally at the secondary site, and the data will be replicated to the replicated disk at the primary site.

With this database recovery on the secondary site, the secondary site is complete and ready to become the primary site.

HPE Data Protector and VMware integration: Protecting your VMs from disaster

Using the HPE Data Protector Virtual Environment integration, you can back up the virtual machines, virtual machine disks, and virtual machine templates. In a vCenter environment, HPE Data Protector communicates with VMware vSphere through the vCenter Server system. All backup and restore requests are sent to vCenter Server. In one session, you can back up virtual machines from one or multiple data centers.

HPE Data Protector identifies data centers and virtual machines by their VMware vSphere inventory path**.** When you back up a virtual machine, you actually back up virtual machine files of the .vmx and .vmdk types.

HPE Data Protector Virtual Environment Integration component

The HPE Data Protector Virtual Environment Integration component (VEAgent) must be installed on at least one HPE Data Protector client in the cell. This client is called the backup system. It can be a virtual machine, or HPE Data Protector Cell Manager, or a dedicated physical backup system.

vStorage Image backup method

The vStorage Image backup method provided by the HPE Data Protector Virtual Environment integration is based on the VMware vStorage technology. For this method, a single central backup host is used to back up all virtual machines hosted by ESX(i) Server systems in an HPE Data Protector cell. This backup host can be a dedicated physical host, a virtual machine, or the Cell Manager. The important point is that it has the HPE Data Protector Virtual Environment Integration component (VEAgent) installed.

During a vStorage Image backup, VEAgent first establishes connections between the backup host and the virtualization host. This connection can be either through a vCenter Server system or direct in the case of a standalone ESXi Server environment. It then requests a snapshot of the virtual machine that is to be backed up via the vStorage API for Data Protection (VADP). This snapshot is used during the period of the backup to keep the virtual machine in a consistent state.

VEAgent then opens the virtual machine disks across LAN or SAN, initializes the Media Agent client, and controls the transfer to it of the virtual machine and all its associated data.

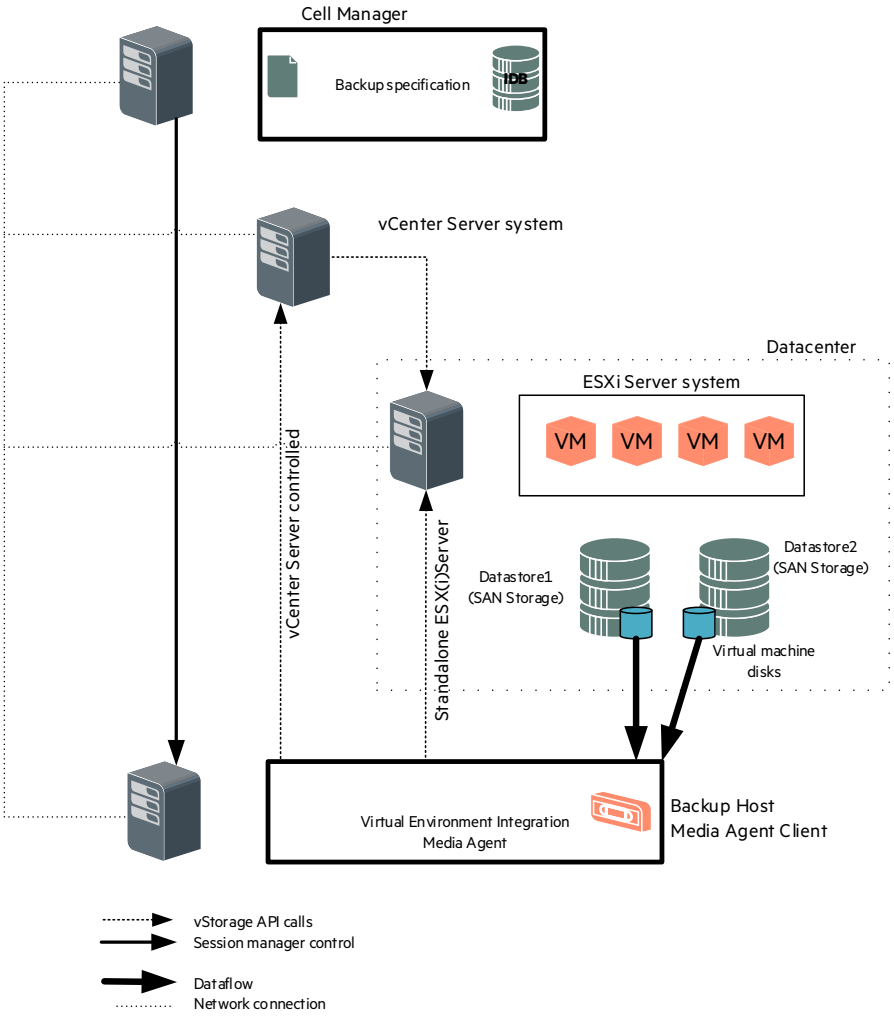

Figure 19. vStorage image method

Backup flow

When the virtual machine backup is initiated, HPE Data Protector triggers a snapshot of the Virtual machine and replicas of the source volumes are created. HPE Data Protector records the Change ID and initiates a backup (Full, Incremental or Differential) and when the backup is complete the snapshot is deleted.

Restore of VMware objects

Virtual machines, virtual machine disks, and virtual machines templates backed up with the vStorage Image method can be restored to a datacenter or to a directory on a backup host. By default, virtual machines are restored to the original datacenter and the original datastore, but you can select a different datacenter if needed. In this DR use case, we have used a different datacenter scenario, which is explained in the session below.

Backup and disaster recovery configuration of VMware VM

Import VMware clients into the HPE Data Protector Cell. For importing a VMware vCenter server, see the integration guide in the HPE Data Protector Manager GUI: HPE Data Protector Manager GUI → Help → Guides → integration guide.

1. To take the backup of a virtual machine, you need to add new backup specifications under Virtual Environment; and then specify the application that you want to backup.
2. Provide the necessary information, and click **Next**. Figure 20 shows how to specify the backup specification for a virtual machine.

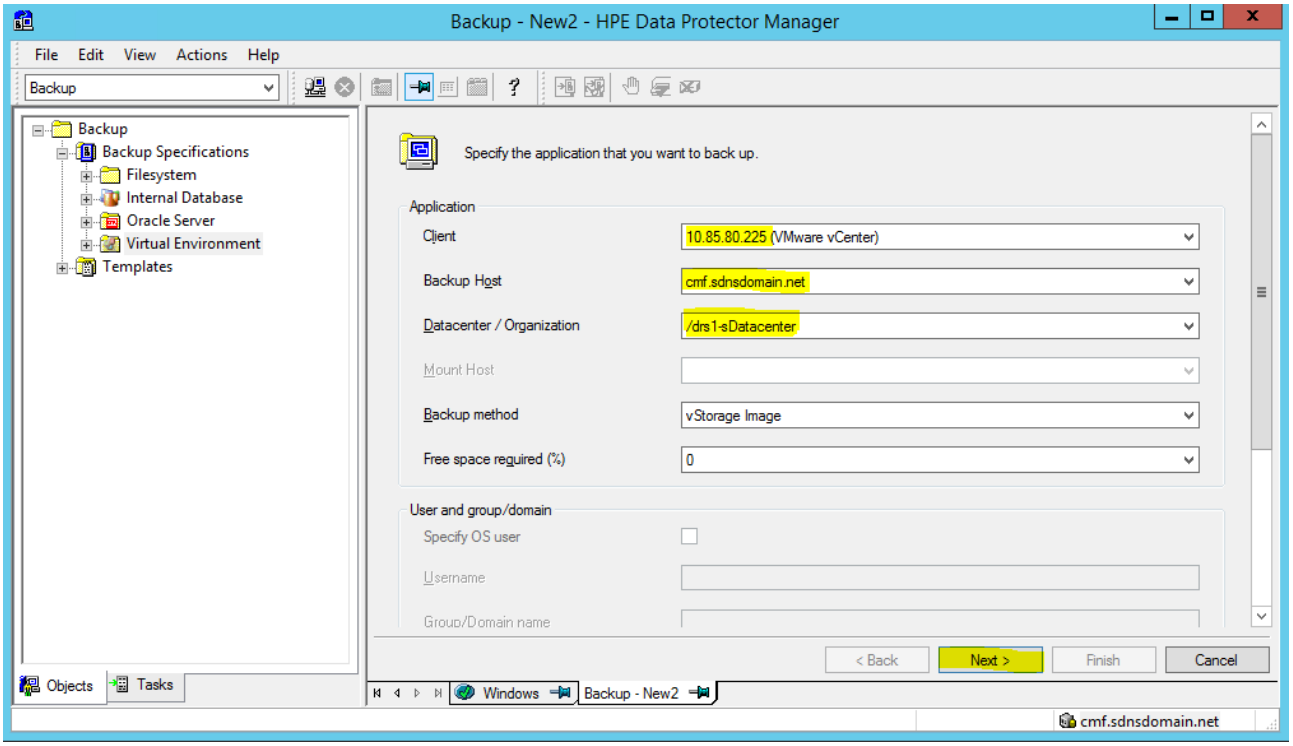

Figure 20. Shows the Backup specification for Virtual machine - select the application

3. Select the virtual machine that you want to back up, and click **Next**.

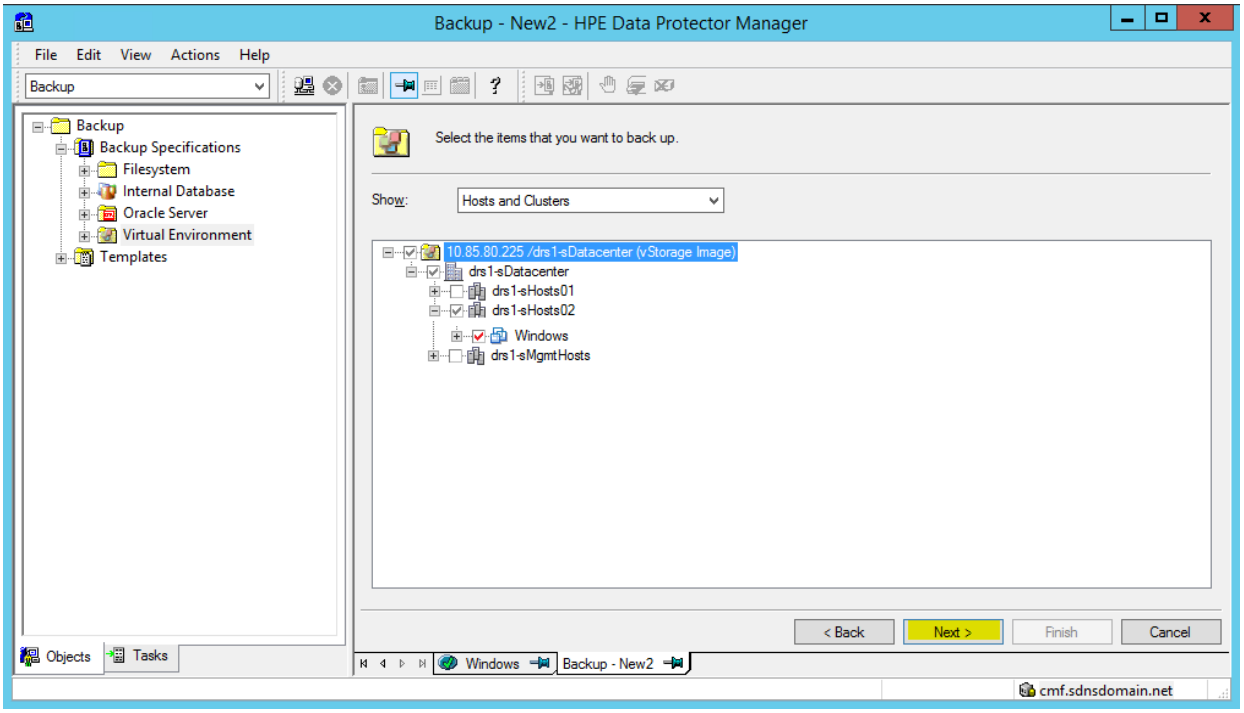

Figure 21. Shows the Backup specification for Virtual machine - select the virtual machine

4. In the subsequent windows, first select the Backup options for all the objects in the backup specification.
5. Then schedule the date and time for the backup specification. This will initiate the backup operation.

Note

We have imported the secondary site Cell Manager as the foreign Cell Manager in the primary Cell Manager. Also, while configuring Object Copies for the backup specification, we have mentioned to make the secondary Cell Manager aware of the backup images. The secondary site Cell Manager will automatically detect the backup images for the restore.

---

6. After the successful execution of the backup and the Object Copy operation at the primary site, go to the secondary site HPE Data Protector GUI, and navigate to **Restore** from the drop-down menu.
7. Go to the Virtual Environment, where you can find the Backup that was replicated to the secondary site as shown in Figure 22.

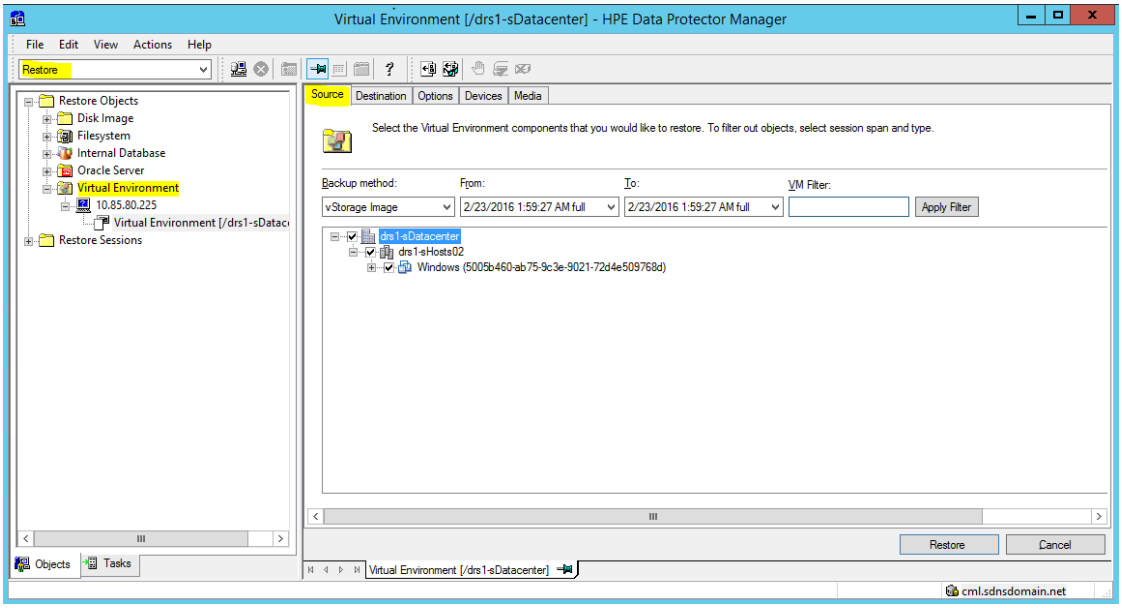

Figure 22. Shows the Source tab of the restore virtual machine window

8. Go to the Virtual Environment Destination tab, and provide the necessary information for the restore operation.

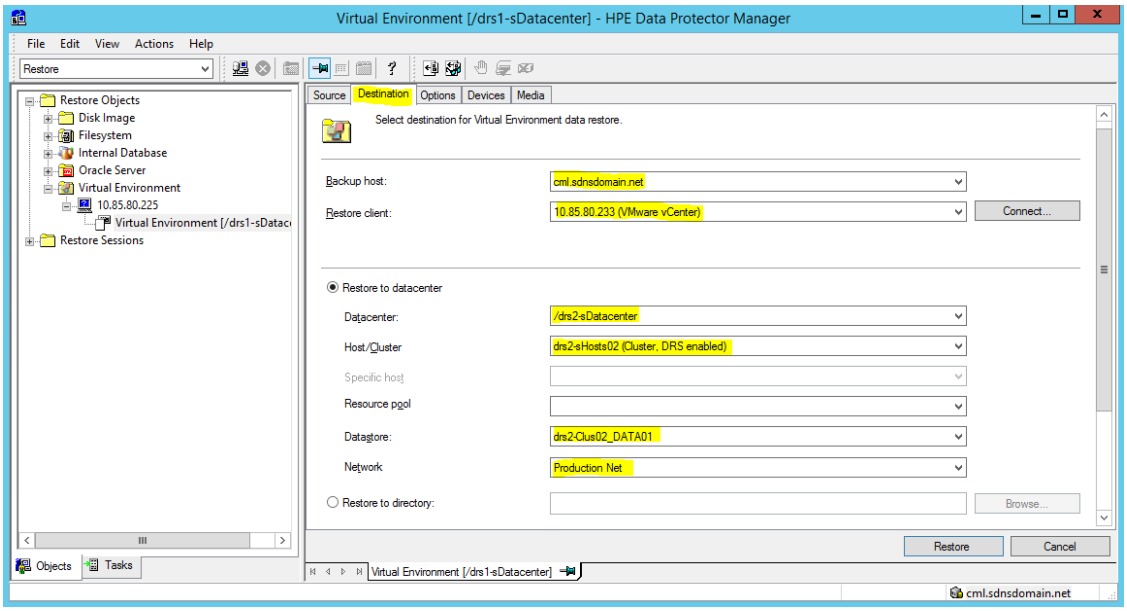

Figure 23. Shows the Destination tab of the restore virtual machine window

9. In all tabs, provide the necessary information for the restore operation and click **Restore**.
10. After successful completion of the restore operation, go to the VMware vSphere Web Client to check the VM status.

For the failback scenario, we have imported the primary site Cell Manager as the foreign Cell Manager in the secondary site Cell Manager. Also while configuring Object Copies for the backup specification, we have mentioned to make the primary Cell Manager aware of the backup images. Once the Object operation completes, the secondary site Cell Manager will automatically detect the backup images for the restore.

For more information on HPE Data Protector VMware integration, see the integration guide in the HPE Data Protector Manager GUI: HPE Data Protector Manager GUI → Help → Guides → integration guide.

Best practices

Object Copy best practices

Object Copy performance can be impacted by factors, such as device block sizes and the connection of devices. If the devices used in the Object Copy session have different block sizes, the data will be repackaged during the session, which takes additional time and resources. If the data is transferred over the network, there will be additional network load and time consumption. This impact can be minimized if the operation is load balanced.

HPE StoreOnce Best Practices

Always ensure that the appliance software in your HPE StoreOnce Backup system is fully up-to-date. Software upgrades also contain all the necessary component firmware upgrades. Where possible, group backups of similar data types to the same destination device (Share/VTL/Store). This grouping can help optimize deduplication ratios.

Use separate StoreOnce Catalyst stores and backup specifications for different data types. Run multiple backups simultaneously to improve the aggregate throughput for an HPE StoreOnce appliance.

Use blackout windows and replication windows to ensure that the appliance is not concurrently performing backup, replication, housekeeping, and offload to tape operations. This practice can keep system performance consistent throughout the backup period.

Configure multiple Ethernet ports in a network bond to achieve increased available network throughput. Identify and resolve other performance bottlenecks in your backup environment, such as slow clients and Media Agents.

Make use of the HPE Storage Sizing tool to size your HPE StoreOnce solution. This tool is available at hpe.com/storage/sizer

Networking Best Practices

The recommendation is to configure the HPE StoreOnce with bonded 10GbE connections where possible to allow increased throughput. Multiple Clients and Media Agents can write to HPE StoreOnce simultaneously, reducing the total backup window required.

- Adhere to the suggested maximum number of concurrent operations per share/appliance. Be aware that there is a limit of 48 concurrent writes to NAS shares on HPE StoreOnce. If more streams are required, use multiple shares.
- For NFS shares, use the correct mount options to help ensure in-order delivery and improve deduplication ratios. For more details, see the HPE StoreOnce Backup System Linux and UNIX Configuration Guide at, hpe.com/info/storeonce/docs.

Licensing

HPE Data Protector supports two licensing schemes:

- Traditional licensing based on features and backup targets – available for all versions of HPE Data Protector software
- Capacity Based Licensing – available with HPE Data Protector 7.01 and above

See HPE Data Protector software QuickSpecs for more details on licensing, at,
http://h20195.www2.hpe.com/V2/GetDocument.aspx?docname=c04109270

Summary

Organizations invest in converged technology stacks to reap the benefits of a single pretested and validated solution platform that can help reduce risk, increase operational efficiencies, and deploy workloads upon a predictable infrastructure set. With a converged approach, you get a modular solution design that can scale with the demands and growth of your workloads. This converged technology provides you with a simplified and consistent approach to key processes for Oracle Databases: data protection, recovery, DR, provisioning, cloning, replication and migration. The HPE solution built upon the HPE ConvergedSystem 700, HPE StoreOnce, and HPE Data Protector software removes the burdens of a build-it-yourself model. Following are the key takeaways of this project.

- HPE ConvergedSystem 700 is an optimized infrastructure platform for protecting virtualized mission-critical applications against disasters.
- HPE StoreOnce deduplication capability provides good RPOs and RTOs for virtualized mission-critical applications as well as backup storage efficiencies.
- HPE Data Protector provides a comprehensive DR solution for HPE ConvergedSystem 700 utilizing HPE StoreOnce Catalyst and replication technologies that reduce risk and costs for disaster recovery.

Appendix A: Bill of materials

The HPE ConvergedSystem 700 is a predefined solution with the ability to vary key components to meet a customer's needs. The product ordering information is defined in detail in the HPE ConvergedSystem 700 QuickSpecs in the How to order section of that document. This list notes the key customizations used for testing this solution.

(Per Site) HPE ConvergedSystem 700 consists of:

- 8 x HPE ProLiant BL460c Gen9 server blades, each with
  - 2 x Intel® Xeon® E5-2690 v3 12-core CPUs @ 2.60 GHz
  - 256GB memory
  - HPE Virtual Connect 10GbE FlexFabric
- 2 x HPE ProLiant DL360 Gen9 (management servers), each with
  - 2 x Intel Xeon E5-2680 v3 12-core CPUs @ 2.50 GHz
  - 256GB memory
- HPE BladeSystem c7000 enclosure
  - 2 x HPE Virtual Connect FlexFabric-20/40 F8 module
  - 2 x HPE Onboard Administrator
- HPE 3PAR StoreServ 7400c 4-node
  - 120 x 300GB SFF Fast Class (FC) in 10 SFF Drive enclosures
- Switches
  - 2 x HPE SN6000B Fibre Channel switches
  - 2 x HPE 5900AF-48XG-4QSFP+ network switches
  - 2 x HPE 5900AF-48G-4XG-2QSFP+ network switches
  - Optional Cisco Nexus series network switches
    - 2 x Cisco Nexus 56128P network switches
    - 2 x Cisco Nexus 3048TP-1GE network switches
    - 2 x Cisco Nexus 2248PQ network switches
- HPE StoreOnce 4500 24TB Backup

The HPE ConvergedSystem 700 QuickSpecs can be viewed as HTML at, http://h20195.www2.hpe.com/V2/GetHtml.aspx?docname=c04505387.

The HPE StoreOnce QuickSpecs can be viewed as HTML at, http://h20195.www2.hpe.com/V2/GetHTML.aspx?docname=c04328820.

Appendix B: Solution hardware configuration and components

Table B-1 outlines the hardware configurations deployed in HPE ConvergedSystem 700 based on the number of workload servers and the particular storage solution selected.

Table B-1. Hardware and software configurations

Component	Quantity	Description
Workload server	2–16 per c7000 enclosure	HPE ProLiant BL460c Gen9 server blade each with: • 2 x Intel Xeon processor E5-2690 v3 (2.66 GHz/12-core), E5-2680 v3 (2.5 GHz/12-core), E5-2670 v3 (2.3 GHz/12-core), E5-2660 v3 (2.6 GHz/10 core), E5-2650 v3 (2.3 GHz/10-core), E5-2640 v3 (2.6 GHz/8 core), E5-2683 v3 (2 GHz/14 core), E5-2630 v3 (2.4 GHz/8 core), E5-2620 v3 (2.4 GHz/6 core), E5-2623 v3 (3 GHz/4 core), E5-2609 v3 (1.9 GHz/6 core), E5-2685 v3 (2.6 GHz/12 core), E5-2603 v3 (1.6 GHz/6 core), E5-2650L v3 (1.8 GHz/12 core), E5-2698 v3 (2.3 GHz/16 core), E5-2630L v3 (1.8 GHz/8 core), E5-2695 v3 (2.3 GHz/14 core), E5-2637 v3 (3.5 GHz/4 core), E5-2697 v3 (2.6 GHz/14 core), E5-2667 v3 (3.2 GHz/8 core), E5-2643 v3 (3.4 GHz/6 core), E5-2699 v3 (2.3 GHz/18 core) • 64 GB–512 GB • HPE Smart Array P244br controller and 1G FBWC • Optional 2 x SAS hard drives (300 GB/15,000 rpm, 450 GB/10,000 rpm, or 600 GB/10,000 rpm) • 1 x HPE FlexFabric 20Gb 2-port 650FLB FIO Adapter
	2-8 per c7000 enclosure	HPE ProLiant WS460c Gen9 server blade • NVIDIA GRID K1 or K2 Graphics accelerator options • Mixed blade environment supported
Management server	2 or 4 (See Note)	HPE ProLiant DL360 Gen9 server, each with: • 2 x Intel Xeon Processor E5-2680 v3 (2.5 GHz/12-core) • 256 GB RAM • HPE Smart Array P440ar controller with 2 GB FBWC • 8 x SAS hard drives (600 GB/10,000 rpm) • HPE FlexFabric 10Gb 2-port 556FLR-SFP+ Adapter Note: CloudSystem SW option will add an additional 10Gb NIC to each DL360 Gen9 management server on the order. • HPE 82Q 8Gb Dual Port PCIe Fibre Channel Host Bus Adapter • Optional PCI Card: HPE StoreFabric CN1200E 10Gb CNA • SATA DVD-RW optical drive • Note: CloudSystem SW install option will include DL360 Gen9 server and additional 10Gb NICs for each management server.
HPE 3PAR StoreServ Storage	7200c 2N	HPE 3PAR StoreServ 7200c storage in a 42U HPE i–Series rack
	7400c 2N or 4N	HPE 3PAR StoreServ 7400c storage in a 42U HPE i–Series rack
	7440c 2N or 4N	HPE 3PAR StoreServ 7440c storage in a 42U HPE i–Series rack
	7450c 2N or 4N	HPE 3PAR StoreServ 7450c storage in a 42U HPE i–Series rack
Network switches (HPE option)	2 (4 when greater than 4 enclosures)	HPE 5900AF-48XG-4QSFP+ switches
	2	HPE 5900AF-48G-4XG-2QSFP+ switches (multi-rack only)
Network switches (Cisco option)	2 (4 when greater than 4 enclosures)	Cisco Nexus 2248PQ and 3048TP
	2	Cisco Nexus 56128P (multi-rack only)

Component	Quantity	Description
SAN switches	2 (4 when greater than 4 enclosures)	HPE SN6000B 16Gb 48/24 Fibre Channel switches (with an optional upgrade to activate 24 additional ports) or HPE SN6000B 16Gb 48/48 Fibre Channel switches (multi-rack only)
	Flat SAN	Single-rack only
Infrastructure	1–8	HPE BladeSystem c7000 Platinum enclosure, each with: • 2 x HPE Virtual Connect FlexFabric 20/40 F8 modules • 2 x HPE Onboard Administrator for the HPE BladeSystem c7000 enclosure
	1 (1 or 2 c7000 enclosures) or 2 (3, 4, or 5 c7000 enclosures) or 3 (6, 7, or 8 c7000 enclosures)	HPE i–Series rack—42U compute rack
	1–2 (Array 1 rack + Optional Array 1 extension rack) 3–4 (Optional Array 2 Rack + Array 2 extension rack)	HPE i–Series rack—42U storage rack

Note

You can customize the configuration components of HPE ConvergedSystem 700. Refer the document at this link http://h20195.www2.hpe.com/V2/GetDocument.aspx?docname=4AA5-6612ENW for more details.

Resources and additional links

HPE StoreOnce Backup System User Guide, http://h20564.www2.hpe.com/portal/site/hpsc/public/kb/docDisplay/?docId=c02295179

HPE StoreOnce Backup System Linux and UNIX Configuration Guide, http://h20564.www2.hpe.com/portal/site/hpsc/public/kb/docDisplay/?docId=c02299831

HPE Data Protector, http://h20195.www2.hpe.com/V2/GetHTML.aspx?docname=c04109270

HPE Storage, hpe.com/storage

HPE Storage Sizer, hpe.com/storage/sizer

HPE ConvergedSystem 700, hpe.com/info/cs700

HPE Reference Architectures, hpe.com/info/ra

To help us improve our documents, please provide feedback at hpe.com/contact/feedback.

Sign up for updates

4AA6-8131ENW, September 2016