

# Zerto

a Hewlett Packard  
Enterprise company

## Protecting Oracle Databases on VMware with Zerto

Using Zerto continuous data protection  
for VMware single-instance Oracle  
database environments

---

# Contents

OVERVIEW .....	3
ZERTO FEATURES .....	3
ZERTO MODES OF OPERATION .....	4
Local continuous data protection .....	4
Remote disaster recovery and continuous data protection .....	4
Data mobility and migration .....	5
SAMPLE ZERTO SETUP: SOLUTION COMPONENTS .....	5
VMware .....	6
Zerto solution .....	6
Zerto Virtual Manager .....	6
Virtual Replication Appliance .....	7
Oracle application consistency or crash consistency .....	8
Creating a Zerto virtual protection group .....	9
HPE STORAGE .....	10
HPE Alletra 6000 .....	11
HPE Alletra 9000 .....	11
HPE Primera 600 .....	11
HPE StoreOnce VSA for long-term retention .....	11
Data Services Cloud Console .....	12
PUTTING ZERTO MODES INTO SPECIFIC WORKFLOWS .....	12
Failover .....	12
Restore .....	13
File restore .....	13
VM/VPG restore .....	13
LONG-TERM RETENTION ON ZERTO .....	13
Retention repository options .....	13
HPE StoreOnce Catalyst and HPE StoreOnce VSA .....	14
ZERTO ANALYTICS .....	14
Comprehensive analytics .....	14
SUMMARY .....	15

## OVERVIEW

Zerto, a Hewlett Packard Enterprise company, delivers everything you need to protect virtualized and containerized infrastructures in a single, simple, and scalable cloud data management and protection solution. Zerto also provides sizing tools to design the Zerto protection environment to meet the needs of the application and database workload characteristics through a unique data analytics feature.

Deploying Zerto in Oracle® database environments provides key benefits such as continuous data protection (CDP) to achieve the lowest recovery point objective (RPO) and recovery time objective (RTO) in the event of failures, ransomware attacks, or natural disasters. Protecting Oracle databases on VMware® with Zerto delivers CDP to an Oracle environment based on VMware. Zerto can easily provide CDP through replication of the entire virtual machine (VM) and associated storage for Oracle databases.

With Zerto Virtual Protection Groups (VPGs), the Oracle database environment can be protected along with associated applications as a single, user-defined, consistent entity. Multiple VPGs can host single or multiple database applications. With proper sizing, multiple VPGs can be defined for protection of other applications and databases.

Zerto integrates with most major cloud providers and hybrid cloud technology companies and is ideal for moving and protecting data in private (e.g. VMware vSphere) hybrid cloud (e.g. HPE GreenLake), and public cloud (e.g. Azure). The Zerto solution can integrate data to and from hybrid and public clouds and perform long-term data retention using HPE StoreOnce Catalyst or NFS storage. It can even provide immutable copies on Azure blobs, Amazon S3, or other S3-compatible storage in the cloud or on-premise. This paper specifically focuses on using Zerto for vSphere, but the concepts and benefits discussed here are widely applicable when using Zerto for Hyper-V, AWS, Kubernetes, or when using Zerto for VMware as a service on a managed service provider or public cloud (e.g. IBM Cloud, Azure VMware Solution, Oracle VMware Solution, etc.).

Although there are many options for replicating databases, some options have expensive licensing fees or high setup complexity and even then, might focus only on replicating one component of the entire solution (app, database, supporting data, and so forth). Oracle Data Guard or Oracle Golden Gate provide similar services with added requirements for licensing. Zerto enables you to integrate all the components (Oracle and non-Oracle) of a solution into one protection group and maintain crash and copy consistency.

Zerto supports single-instance Oracle databases and the use of Oracle Grid Infrastructure and Automatic Storage Manager (ASM) as well as traditional file systems. Zerto is not currently supported with Oracle Cluster File System (OCFS) or Oracle Real Application Clusters (RAC).

## ZERTO FEATURES

Features of Zerto protection for Oracle databases include dedicated, security-hardened Virtual Replication Appliances (VRAs) for hybrid and public cloud, disaster recovery to the cloud, integration with Kubernetes and container technologies, journal-based recovery, patch testing, and complete solution replication.

- **Disaster recovery**—Always on availability with orchestrated failovers, failback, and non-disruptive testing to deliver the RPOs in seconds and RTOs in minutes. Real-time monitoring of RPO times.
- **Continuous data protection**—Continuous replication of applications and associated single instance databases in virtualized and containerized environments. Continuous journaled replication with frequent checkpoints separated by 5-15 seconds.

- **Long-term data retention**—Long-term offsite copies of data using HPE StoreOnce Catalyst, NFS, or even the public cloud, includes options for immutable storage as of Zerto 9.0.
- **Data mobility and migrations**—Database environment movements to and from the cloud. The ability to move and migrate data swiftly as required in the cloud and hybrid cloud environments. Application consistency with an RPO of zero by taking a final powered-off checkpoint as part of the orchestration to migrate the data.
  - The migration happens with only minutes of downtime because all the requisite data has already been staged on the recovery site by the continuous replication.
  - Migrations can be tested using the Failover Test operation or a Move without a commit operation to ensure there are no surprises on migration day.
  - Migrations can be between dissimilar hardware or dissimilar on-premises hypervisors or public clouds.
- **Test and development**—Ability to quickly deploy and test solutions with on-demand copies of production that are only seconds behind the original VM.

## ZERTO MODES OF OPERATION

Zerto is extremely flexible for its use cases with Oracle databases in a hybrid cloud. Within VMware vCenter® environments, inter-vCenter data movements and migrations can be required as well as intra-vCenter requirements. Critical data can be protected in real-time with the ability to create Oracle application-consistent or crash-consistent copies. Whether local or across vCenters, Oracle application consistency can be maintained using pre- and post-scripts inserted into the recovery process to put the database into hot backup, take a checkpoint, and end hot backup mode.

### Local continuous data protection

A VPG can be set up just for backup purposes where the Oracle database VM and storage can be replicated to a local journal and configured to an off-site, long-term retention destination such as an HPE StoreOnce VSA system or a public cloud platform such as AWS or Microsoft Azure. The local replication enables instant recovery of the VM and storage. This can all be done within the same vCenter environment.

Possible use cases for Oracle include simple backup of the database, splitting off a production copy for test or development, or reporting or extracting data for data warehousing or decision support.

### Remote disaster recovery and continuous data protection

A VPG can be set up for disaster recovery purposes with the ability to replicate Oracle database or database objects to one or more vCenter sites. Long-term retention copies can also be configured in this mode. All databases and applications within the VPG will be protected with crash consistency. Oracle VPGs can be configured for application consistent copies as well.

The RTO is the time it would take to bring up the protected database on the disaster recovery side. Data availability and server availability on the disaster recovery side is immediate. Database availability would be the time it takes to bring the network connections to the disaster recovery server and the time it takes for the Oracle database to perform an application recovery or crash recovery. The RPO can be monitored continuously from the Zerto Virtual Manager (ZVM) or with Zerto Analytics for multi-site monitoring. You can monitor the RPO in real-time within Zerto by viewing the status of the VPG, as shown in Figure 1. The RPO is a function of the write rate of the database and application in the VPG as well as the available network bandwidth (although only a 5 Mbps minimum connection is required).

Sizing the systems and WAN resources is also key to creating a robust replication solution. Zerto provides a planning tool with Zerto Analytics, a SaaS offering included by default with Zerto.

⚠	VPG Name (#VMs) ↑	VPG Type	Direction	Peer Site	Priority	Protection Status	VPG State	Actual RPO
✓	MySQLVPG (1)	Remote	↩	Zerto-Oracle	● ● ○	Meeting SLA		10 sec
✓	Oracle (1)	Remote	↩	Zerto-Oracle	● ● ○	Meeting SLA		10 sec
✓	ZVM1 (1)	Remote	↩	Zerto-Oracle	● ○ ○	Meeting SLA		10 sec

FIGURE 1. VPG status and viewing RPO

## Data mobility and migration

A VPG can be set up to provide data mobility and migration using the move workflow within Zerto. Data can be moved within the vCenter cluster or across clusters. The defined repository for Oracle data can be a storage pool, hybrid cloud storage, or public cloud storage.

Oracle ASM or filesystem data can be migrated as well as a copy of the entire database solution. This can be useful in situations where test environments need to be a representation of another Oracle database environment, such as production. Figure 2 shows the HPE StoreOnce as the repository in this Oracle configuration.

FIGURE 2. Data mobility and migration for Oracle data

⚠	VPG Name (#VMs) ↑	Last Run Result	Success Rate	Date of Last Run	Last Successful Run	Repository
✓	MySQLVPG (1)	Success	100% (1/1)	January 28, 2022		StoreOnce (Zerto-Oracle)
✓	Oracle (1)	Success	100% (1/1)	January 28, 2022	January 28, 2022 4:00:24 A...	StoreOnce (Zerto-Oracle)

## SAMPLE ZERTO SETUP: SOLUTION COMPONENTS

The diagram in Figure 3 shows a sample setup for a disaster recovery solution using Zerto and protection of an Oracle database and a MySQL database. The replication is between two VMware vSphere environments. Each environment uses VMs for the databases and the Zerto VRAs. There is another VM that runs the Zerto Virtual Manager (ZVM), the management and orchestration component of Zerto. Each site must have a ZVM since there is a 1:1 relation between vCenters and ZVMs. Zerto is hardware and storage agnostic to allow deep infrastructure flexibility, so in this example, the storage is HPE Alletra and HPE Primera arrays. All the datastores for the databases and recovery are on HPE external storage.

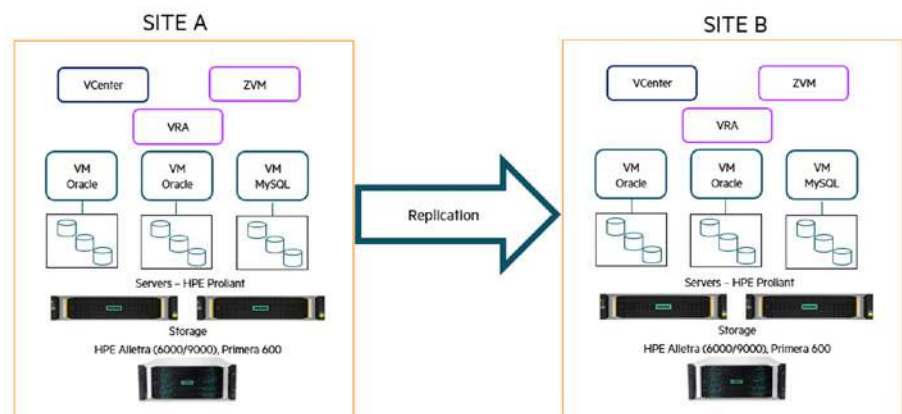


FIGURE 3. A sample Zerto disaster recovery environment using Oracle and MySQL databases

## VMware

The VMware environments would typically be similar with each having its own vCenter server between each site. Each site can have several VMware ESXi™ servers in which some may or may not be configured into the Zerto replication environment. In VMware, the Oracle database would be in a datastore that is attached to the VM running the Oracle database. Within the datastore are the VMDK files that represent the Oracle volumes. These volumes may be multiple VMDK volumes for a single ASM group and there may be multiple ASM groups depending on the specific solution configuration. If these volumes are defined within a single VPG, then all the data will be crash-consistent.

The VMware VM and the defined storage is replicated through the Zerto VRA, which is a virtual appliance VM. There is at least one VRA at each site, and ideally one on each ESXi host in both sites. The recovery VM and storage for the VMDKs are also defined.

<input type="checkbox"/>	VM Name	Protected Volume Location	Virtual machine OS and Oracle binaries	Provisioned	Thin
<input type="checkbox"/>	OL8.3_Ora19	[ds-a630-zerto-DS]:OL8.3(2)(1)/OL8.3-000001.vmdk	ds-a630-zerto-DS	300.0 GB	<input checked="" type="checkbox"/>
<input type="checkbox"/>	OL8.3_Ora19	[ds-a630-zerto-DS]:OL8.3(2)/OL8.3-000001.vmdk		200.0 GB	<input checked="" type="checkbox"/>
<input type="checkbox"/>	OL8.3_Ora19	[ds-a630-zerto-DS]:OL8.3(2)/OL8.3(2)-000001.vmdk		200.0 GB	<input checked="" type="checkbox"/>

FIGURE 4. An example of volume requirements for Zerto to recover a database

## Zerto solution

The main components of Zerto solution are: the Zerto Virtual Manager (ZVM), the Virtual Replication Appliance (VRA) created by Zerto, and the Virtual Protection Group (VPG), created on the ZVM to manage the replication of applications and databases. The VPG is the application group created to define the VMs and storage to be replicated as well as information needed for recovery. In the case of Oracle databases, the VPG would include the VMs running the database and the associated ASM or filesystem volumes needed. A single Zerto site can have many VPGs running that represent multiple databases, applications, or data servers.

## Zerto Virtual Manager

The ZVM is the main management tool for setting up and maintaining continuous replication of applications and databases. In this document, it is an Oracle database. This tool can be used to manage and monitor the activity of the defined replication processes. The main dashboard on the ZVM allows for the monitoring of the health status in terms of meeting the solution service level agreements (SLAs) and the RPO of the components in the VPGs.

Figure 5 shows performance panes for monitoring the replication activity for applications and databases inside of VPGs. Also available is the ability to understand the general health of the solution as well as its ability for the RPO to meeting SLAs. The IOPS, throughput, and WAN throughput are metrics for all VPGs configured in the Zerto solution. In the example shown in Figure 5, there is only one Oracle VPG configured but there could be several VPGs. The ZVM dashboard provides information on numbers of VPGs and VMs as well as average RPOs, alerts, and events.

The VPG status pane is an important pane that indicates the ability of the solution (in VPGs) to meet SLAs. This is a function of the overall write rate between sites or repositories and the defined journal size as well as the throughput capabilities within VMware and the WAN connection. Alerts can be monitored here as well as general events.

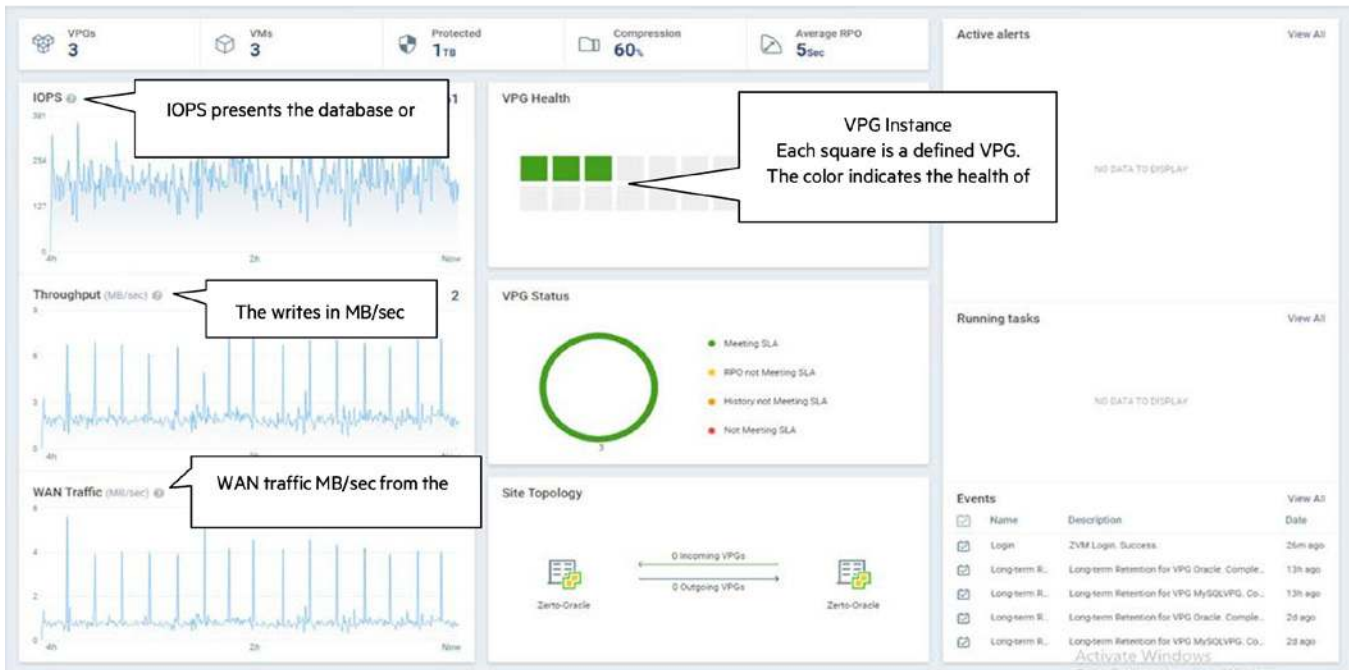


FIGURE 5. Zerto Virtual Manager dashboard

## Virtual Replication Appliance

The journals are attached to the VRA on the chosen recovery host for each VM. The checkpoints are stored in the journal for the duration of the defined journal history SLA within the VPG settings, with a user-defined history of one hour up to 30 days. The checkpoints are created by a continuous synchronization process between the source and target VRAs as well as the ZVM.

As the replication is occurring, Zerto is continuously monitoring and maintaining the defined RPO. The ability of Zerto to maintain the defined RPO, providing the solution has been sized properly, is based on the sizing of the systems and database workload write activity. Zerto is taking a checkpoint as fast as possible, typically every 5-15 seconds, and you can make specific checkpoints through the GUI interface or by scripting. The replicated data is written to a journaling system and eventually updated to disk. The VRAs in Figure 6 represent an example of two separate sites that contain the Oracle VPGs.

### NOTE

For an Oracle database to maintain consistency of its data, all database groups and objects for a specific database instance must be a part of the same VPG. Zerto ensures write-order fidelity and consistent checkpoints between all VMs in the same VPG.

Cluster / Host	Host Version	VRA Name	VRA Status	VRA Version	VRA Address	# VPGs	# VMs
FTC-lab							
bm8.hpelab.local	7.0	Z-VRA-bm8.hpelab...	Installed	Latest	192.168.20.100	2	2
bm9.hpelab.local	7.0	Z-VRA-bm9.hpelab...	Installed	Latest	192.168.20.101	3	4

Callouts in the image point to:

- ESXi host location for the VRA (points to the Cluster/Host column)
- VRA name (points to the VRA Name column)
- VRA status (points to the VRA Status column)

FIGURE 6. Zerto VRAs

## Oracle application consistency or crash consistency

Figure 7 illustrates that the Oracle database must be placed into hot backup mode to achieve database application consistency. Hot backup mode freezes the system change number (SCN) in the redo logs and it gives Oracle a specific application-consistent location to recover. After Oracle is in hot backup mode, the Oracle specific checkpoint for Zerto can be taken. After the checkpoint is complete, hot backup can be disabled.

In the case of crash consistency, because all the database objects in a specific database instance are part of the same VP, the data is consistent and proper write ordering will take place. In the event of a crash recovery, the database will recover from a specific checkpoint and behave like it is recovering from a power failure.



FIGURE 7. Workflow for creating an application-consistent checkpoint on an Oracle database

### Example Implementation of application consistency checkpoints

The example in Figure 8 shows the basic commands used to create an application consistent checkpoint for Oracle recovery. These basic commands can be integrated into another script that fits into your specific environment.

For this instance, the script is running on a Windows 2019 server running the ZVM using PowerShell 7.0 and the version 1.03 Zerto PowerShell Commandlet API. Communication to the Oracle server from the ZVM is via SSH and executing commands on SQLPlus.

```

Disable SSL (if no certificate is being used)
Remove-ZvmSslCheck

Connect to the ZVM through PowerShell
Connect-Zvm -Host ZVMHost -Credential administrator@vsphere.local

Perform Log Switch from ZVM with PowerShell
ssh oracle@oracleserver ". .bash_profile;sqlplus -s system/manager << EOF
>> ALTER SYSTEM SWITCH LOGFILE;
>> exit;
>> EOF"

Put database into hot backup
ssh oracle@oracleserver ". .bash_profile;sqlplus -s system/manager << EOF
>> alter database begin backup;
>> exit;
>> EOF"

Create a Zerto checkpoint for recovery
Start-ZvmVpgTaggedCheckpointInsert -VpgID OracleVPG -CheckPointName OracleAppCheckpoint

Take the database out of hot backup
ssh oracle@oracleserver ". .bash_profile;sqlplus -s system/oradata1 << EOF
>> ALTER DATABASE END BACKUP;;
>> exit;
>> EOF"

Perform log switch from ZVM
ssh oracle@oracleserver ". .bash_profile;sqlplus -s system/oradata1 << EOF
>> ALTER SYSTEM SWITCH LOGFILE;
>> exit;
  
```

FIGURE 8. Commands needed to script an application consistent checkpoint for an Oracle database



After you create the checkpoint, you can choose the checkpoint name during the Zerto recovery process. For example, from Figure 8 you would find a name of “OracleAppCheckpoint” in the list. This is the checkpoint to use during recovery.

For more information about Zerto best practices with Oracle databases, see [Best Practices When Protecting an Oracle Database](#).

These commands are the basic commands used to create an application consistent backup and are intended to be integrated into a customer script written by an administrator for making the logic fit the environment and provide proper error checking. For more information on the tools, see Zerto Commandlets in the PowerShell Gallery. Other useful Oracle SQLPlus queries are listed in Table 1.

Description	Command
Check that archive log mode is enabled	<code>select log_mode from v\$database;</code>
View tablespaces	<code>select tablespace_name from dba_tablespaces;</code>
Check tablespace readiness for hot backup	<code>select * from v\$backup;</code>
List archive log status	<code>archive log list;</code>

TABLE 1. Other useful SQLPlus queries

## Creating a Zerto virtual protection group

Figure 9 shows a list of all VMs as part of the associated vCenter in the solution. The unselected VMs are whatever applications, databases, or data sets might be available to replicate. If a VM is to be part of the new VPG, move that VM to the “Selected VMs” section and it will be included in the replication set.

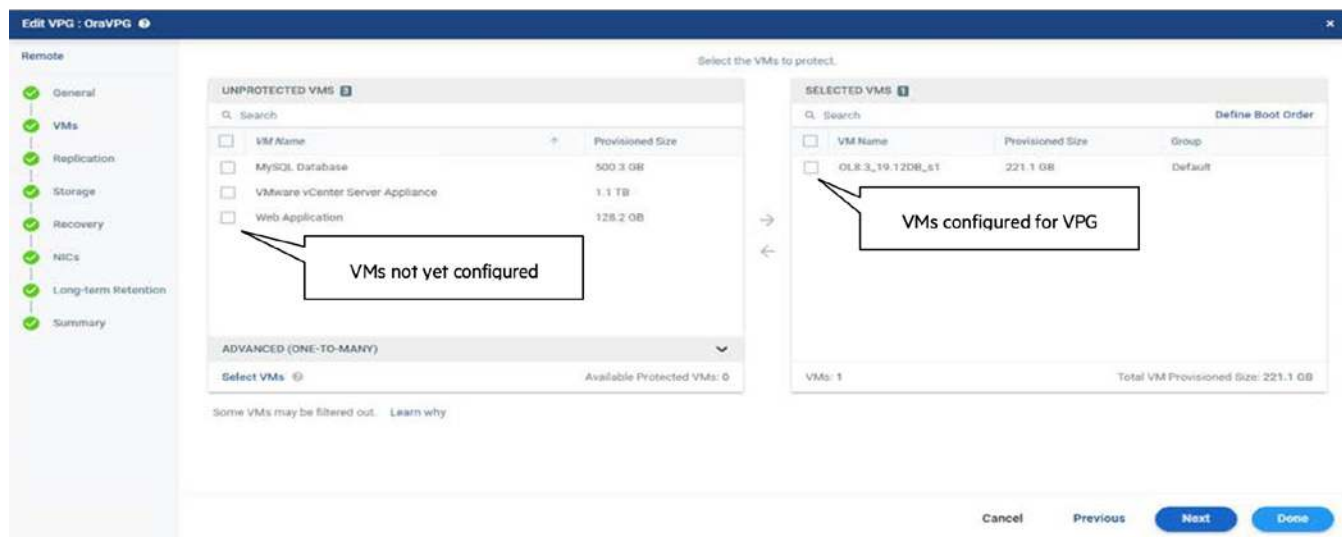


FIGURE 9. VPG configuration—VM selection

As shown in Figure 10, other important parameters needed to create a VPG are the hostname or IP address of the recovery VMware ESXi® server in the target side of replication along with the associated datastore for the VMs in the VPG.

The journal history is selected as part of the VPG definition. Data is replicated to the journal as well as to the recovery volume datastore. Within the journal are the I/O data being replicated along with checkpoints. You can recover to a specific checkpoint. A scratch journal is created to be used for test and live failovers. The purpose for the journal history is for holding the latest checkpoints. Journals are always stored on the recovery site in order to allow recovery in case the source site becomes completely unavailable.

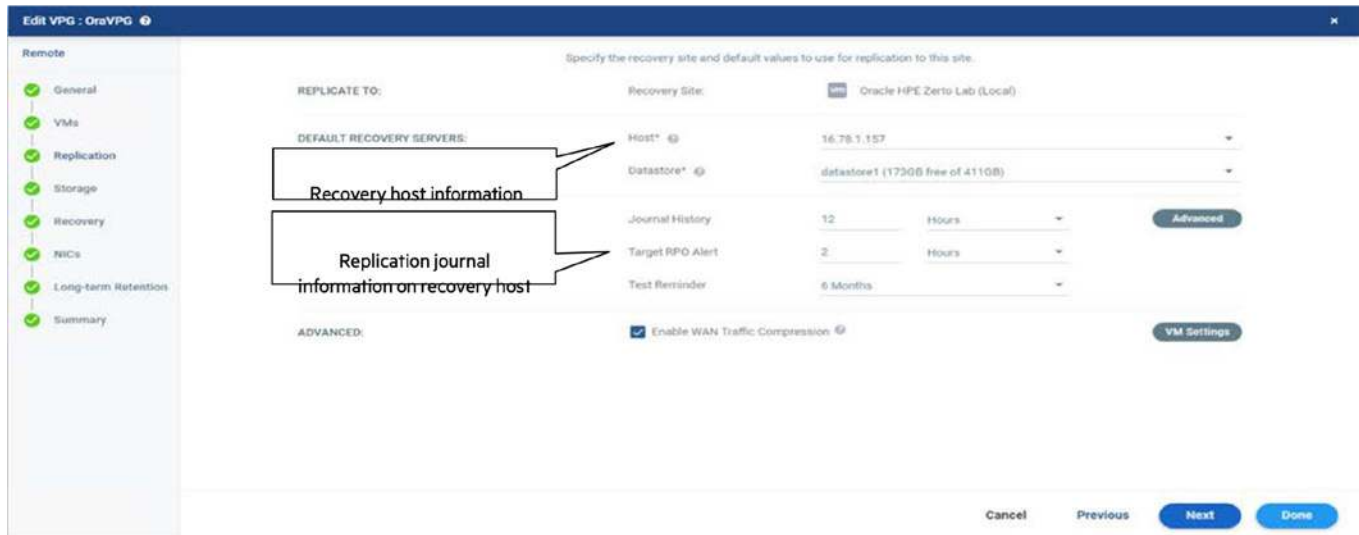


FIGURE 10. VPG recovery server information

## HPE STORAGE

For Oracle databases, the most popular HPE storage systems are HPE RAID protected storage arrays such as HPE Alletra 9000, HPE Alletra 6000, HPE Primera 600, HPE Nimble Storage, and HPE XP8 storage. Figure 11 shows a wide variety of storage options within Oracle for storage platforms, depending on the application.

The HPE Alletra 9000, HPE Alletra 6000, and HPE Primera 600 platforms are the latest in the portfolio. These three platforms can be used for both traditional hardware infrastructures as well hybrid cloud implementations.

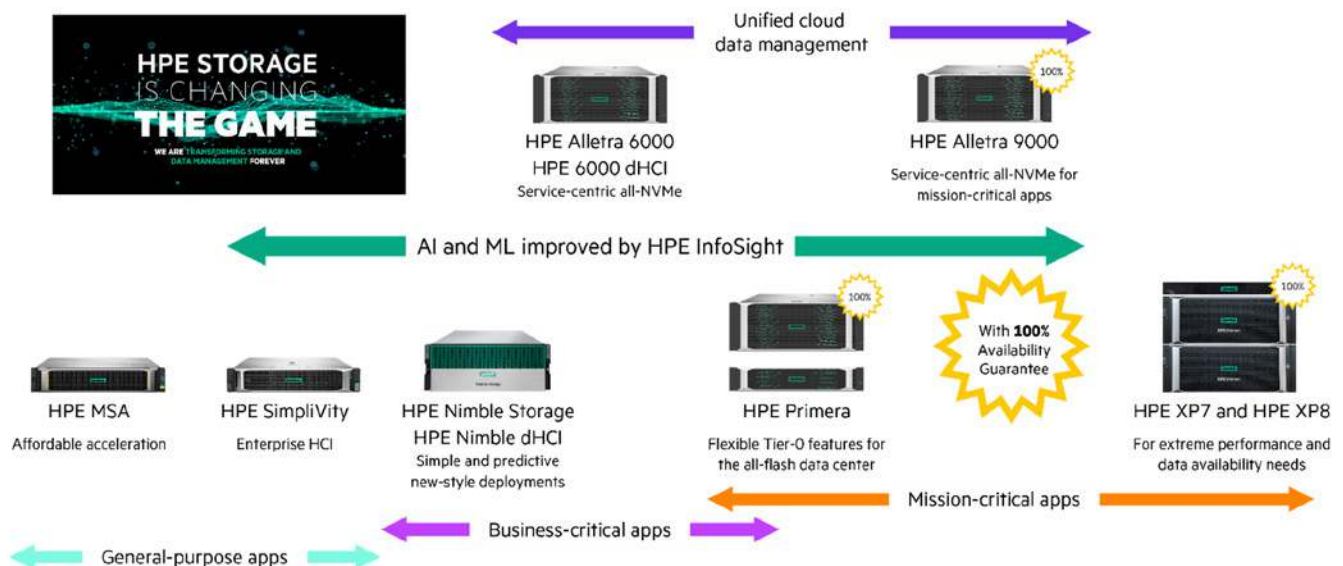


FIGURE 11. Feature comparison matrix of HPE Storage

## HPE Alletra 6000

The HPE Alletra 6000 is an excellent choice for environments running business critical content. HPE Alletra 6000 arrays are an excellent choice for containerized environments and deployment in small to medium-scale production deployments.

## HPE Alletra 9000

The HPE Alletra 9000 is the best fit for mission critical, high IOPS and data transfer performance capabilities with minimized latencies and a high processing density of extreme power and availability in a single 4U rack space. It is the storage system most capable for handling a mixture of block sizes and delivering high throughputs, positioning the storage into a private cloud or traditional data center implementation.

HPE engineering experimented with multiple workloads such as Oracle Database, MongoDB, and a complex synthetic workload that demonstrated seven times greater latency benefits of a traditional array without adverse effects to the Oracle database. You can configure HPE Alletra 9000 provisioning by using the Data Services Cloud Console (DSCC), which is based on a cloud management process Hewlett Packard Enterprise is now implementing with storage.

## HPE Primera 600

HPE Primera can run seamlessly in a clustered or single instance environment with extremely good production performance and mission critical availability where an HPE Alletra array is not required. HPE Primera is currently used for Oracle databases on a worldwide basis and is a proven solution for large enterprise Oracle production databases.

## HPE StoreOnce VSA for long-term retention

HPE StoreOnce products are secondary storage systems for data protection and copy data management. Using data deduplication and compression, backup data is efficiently and economically stored. HPE StoreOnce Integrity Plus is industry-leading patented technology that ensures the backup data is available for restore and recovery whenever it is needed. HPE StoreOnce systems are available as purpose-built hardware appliances and as virtual appliances hosted by a hypervisor. Figure 12 shows a backup of Oracle data to a Zerto long-term retention repository hosted on HPE StoreOnce. Catalyst shows possible dedupe benefits; many capacity details and configuration details have to do with Catalyst stores created through Zerto.

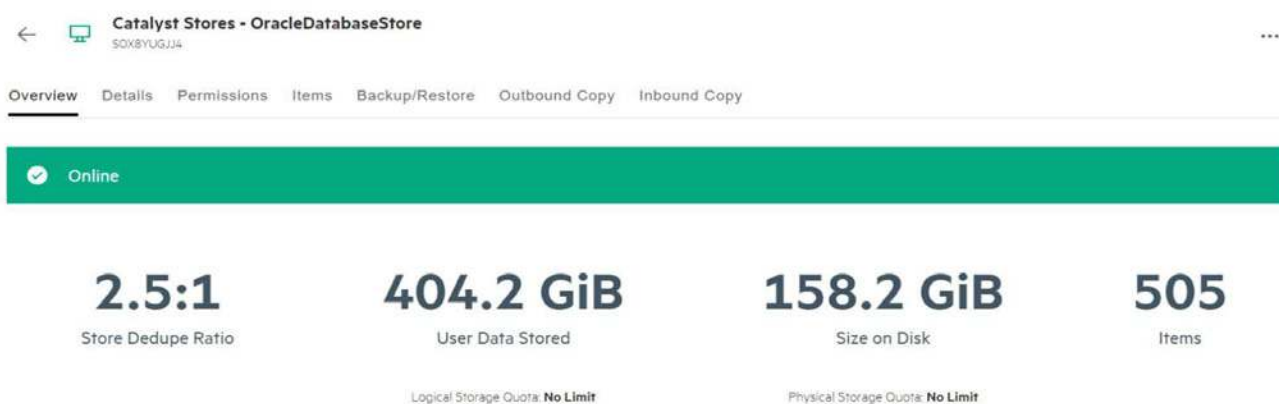


FIGURE 12. Catalyst store for Zerto long-term retention of data

## Data Services Cloud Console

HPE GreenLake's DSCC is a SaaS-based console that delivers a suite of cloud services designed not only to enable cloud operational agility for data infrastructure everywhere, but also to unify data operations across the data lifecycle. It is a powerful combination of capabilities made possible by a unique cloud-native architecture that automates and orchestrates integrated infrastructure and data workflows. This capability transforms complex data operations into a streamlined data management experience across clouds.

DSCC offers multiple cloud apps including Data Ops Manager (DOM). DOM is used to manage onsite storage from the cloud. The storage itself is in a hybrid cloud where it can provide the kind of agility and performance needed by apps and databases such as Oracle Database. The current HPE storage platforms managed with DOM are the HPE Alletra 6000, HPE Alletra 9000, and HPE Primera 600 arrays. These arrays continue to support command line interfaces (CLIs) and other APIs needed to communicate such as REST, Python, and PowerShell APIs. They are good candidates for cloud orchestrations through Kubernetes and other such systems.

Typical functions that can be performed with DOM are management of individual arrays, system access to arrays, and volume provisioning. You can create datastores for Oracle databases and Zerto recovery for Zerto and VMware.

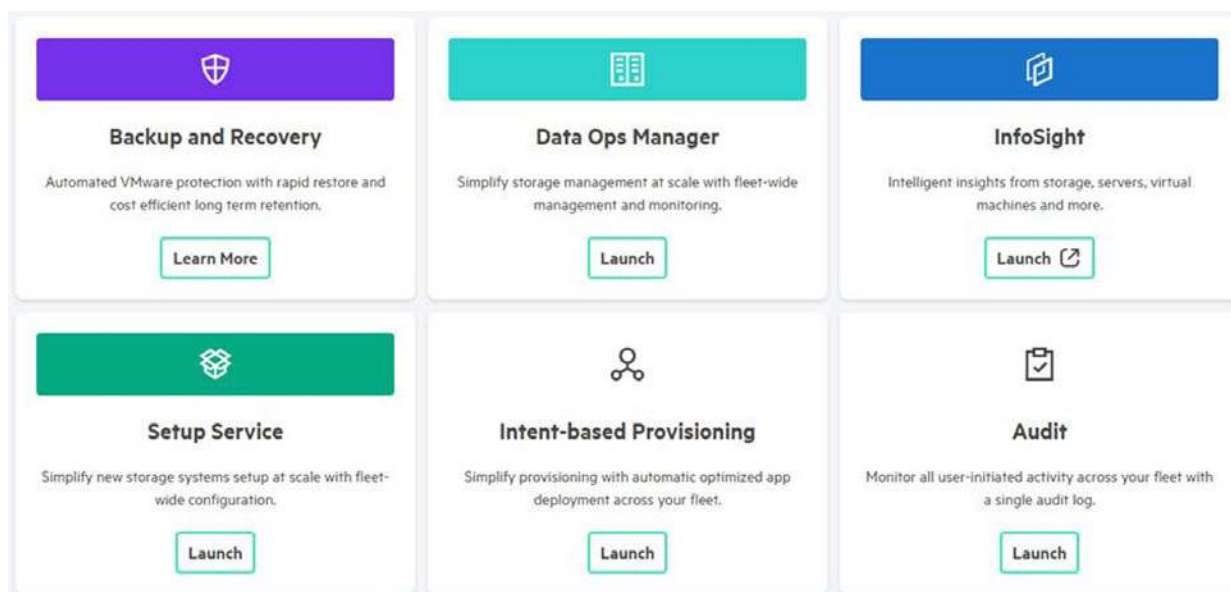


FIGURE 13. Data Ops Manager in DSCC

## PUTTING ZERTO MODES INTO SPECIFIC WORKFLOWS

With Zerto, it is possible to replicate the same Oracle database workload to up to three recovery sites simultaneously. One of these can also be a local copy. For example, a local copy could be available for quick restores, a second copy at an on-premises site, a third copy in the public cloud, and long-term retention. Replication is applicable for test environments, managed service environments, and the cloud.

### Failover

There are two different modes for failover:

- A test failover enables you to test the failover of the VM and datastores. Test VPGs can be preconfigured to failover on a test network, eliminating any disturbance to the primary site production systems.
- A live failover allows an actual failover to be done to the disaster recovery site on the production network.

## Restore

The restore workflow restores a VM and datastore that has been damaged by hardware, software, or human error.

### File restore

Figure 14 shows how file restore looks at the journal entry. Journal entries occur every five seconds.

	Point in Time	Source	Type	Na...
<input type="radio"/>	January 31, 2022 11:22:24 PM	Journal	Auto	
<input type="radio"/>	January 31, 2022 11:22:19 PM	Journal	Auto	
<input type="radio"/>	January 31, 2022 11:22:14 PM	Journal	Auto	
<input type="radio"/>	January 31, 2022 11:22:09 PM	Journal	Auto	
<input type="radio"/>	January 31, 2022 11:22:04 PM	Journal	Auto	
<input type="radio"/>	January 31, 2022 11:21:58 PM	Journal	Auto	

FIGURE 14. File restore journal entries

### VM/VPG restore

This type of restore occurs at the VM or VPG level. After the VM or VPG has been specified, Zerto looks in the repository and presents options for datasets to restore.

Restore VM Select ▼ ⓘ

Restore VPG Select ▼ ⓘ

FIGURE 15. Shows the options to restore by selecting either the VM or the VPG

## LONG-TERM RETENTION ON ZERTO

Long-term retention is changing; there are new solutions for customer data that provide much better tiering of data based on usage. Zerto long-term retention provides efficient data movement and extremely fast recovery times. If you want to back up Oracle data to the cloud, Zerto has the capabilities to make this happen. Data can be compressed for more efficient network payloads and data can be stored immutably on Azure blobs or in Amazon S3. Amazon S3 is used for cold unstructured data but is also suited for structured data, usually cold unused data. You can create Amazon S3 buckets from the AWS management console. For more information, see [Deploy & Configure Long-Term Retention for Amazon S3](#).

### Retention repository options

Zerto provides the additional details needed for a more comprehensive understanding of Zerto offerings, including various workloads for Oracle database, containerized workloads, and CDP. There are several options for providing long-term protection and even immutable storage options.

## HPE StoreOnce Catalyst and HPE StoreOnce VSA

Catalyst is a powerful option for long-term retention in a Zerto data protection environment based on VMware. Catalyst is a backup protocol for enterprise data backups requiring high capacity and performance data protection. It can provide deduplication and copy replication.

Integrating Catalyst with Zerto brings all the rich features of HPE StoreOnce to a Zerto solution. Many Oracle environments are already using Catalyst and existing HPE StoreOnce environments could be integrated into Zerto.

HPE StoreOnce VSA brings the functionality of HPE Store Once onto a VM and integrates well into the VMware environment. By spinning up the VSA on a VMware environment and providing a minimum of a 1 TB datastore, you can integrate the VSA into Zerto through an HPE StoreOnce Client connection.

Figure 16 shows the comprehensive HPE StoreOnce portfolio. For more information on HPE StoreOnce, see the HPE StoreOnce Next- Generation Platform Brochure.




FIGURE 16. HPE StoreOnce product portfolio

## ZERTO ANALYTICS

Zerto Analytics helps optimize the database environment and improve IT resilience. Zerto Analytics enables you to view Zerto deployments from a single console, even across multiple sites and when protecting both VMs and Kubernetes cluster. The Zerto converged architecture enables Zerto to provide complete visibility into a single data analytics pane. CDP can be realized across the entire VMware environment and proper sizing can be planned to use Zerto data analytics.

### Comprehensive analytics

You can view and monitor Zerto environments from a birds-eye view with workload planning. You can monitor the entire data center from anywhere at any time through SaaS and any mobile device. Zerto Analytics provides real-time and historical analyses with reporting from an easy-to-use dashboard. You can access up to 90-day network, RPO, and journal analyses of all Zerto environments, a definite advantage for critical environments such as Oracle databases.



FIGURE 17. Zerto Analytics dashboard

## SUMMARY

The Zerto solution for disaster recovery based on CDP can be the tool to provide a VMware Oracle database and application environment with all the comprehensive resources needed to effectively manage and protect mission-critical databases. Protect your data from anywhere, around the clock and at all levels, from the hybrid cloud to the public cloud.

## Resources, contacts, or additional links

Zerto [zerto.com](https://www.zerto.com)

HPE data protection solutions <https://www.hpe.com/us/en/storage/data-protection-solutions.html>

LEARN MORE AT

<https://www.hpe.com/us/en/storage/data-protection-solutions/disaster-recovery-zerto.html>

## About Zerto

Zerto, a Hewlett Packard Enterprise company, empowers customers to run an always-on business by simplifying the protection, recovery, and mobility of on-premises and cloud applications. Zerto's cloud data management and protection solution eliminates the risks and complexity of modernization and cloud adoption across private, public, and hybrid deployments. The simple, software-only solution uses continuous data protection at scale to converge disaster recovery, backup, and data mobility. Zerto is trusted by over 9,500 customers globally and is powering offerings for Microsoft Azure, IBM Cloud, AWS, Google Cloud, Oracle Cloud, and more than 350 managed service providers. [www.zerto.com](https://www.zerto.com)

Copyright 2022 Zerto. All information may be subject to change.