

This document describes the best ways to protect Oracle databases in a virtualized environment using Zerto Virtual Replication version 4.0 or higher.

See the following sections:

- [“The Challenges of Protecting an Oracle Database”, on page 1](#)
- [“Zerto Virtual Replication or Oracle Data Guard”, on page 1](#)
- [“Oracle Licensing Issues When Using Zerto Virtual Replication”, on page 2](#)
- [“Zerto Virtual Replication Support for Oracle Database Features”, on page 4](#)
- [“Optimizing Replication Using Zerto Virtual Replication”, on page 5](#)
- [“Failover Testing”, on page 10](#)
- [“Failover, Move, and Failback”, on page 10](#)

The Challenges of Protecting an Oracle Database

Protecting Oracle Database Servers is challenging given the large amounts of data being written on a constant basis in large databases, sometimes terabytes in size. Many different BC/DR technologies can be used to protect Oracle Database Servers, common examples being storage based replication, log shipping, or database mirroring software. Each of these technologies introduces risks to the recovery of Oracle Database Servers.

When protecting Oracle Database Servers, the applications using the databases are often overlooked or not included in the same technology or recovery plan. This leads to complications in recovery and lack of visibility into the protection status of the entire application stack, such as SAP. This can make disaster recovery testing incomplete as the entire stack is not tested together, leading to increased risk in not being able to recover in a timely and efficient manner.

Zerto Virtual Replication or Oracle Data Guard

Oracle Data Guard is an application-based database protection solution for BC/DR that requires a target virtual machine to be permanently running to provide protection of Oracle databases. It requires a separate Oracle license.

When protecting Oracle databases with Zerto Virtual Replication instead of Oracle Data Guard, the same level of continuous protection is provided without the complexity of managing an additional solution that often requires manual intervention to maintain replication nor with the need for an additional Oracle license.

Note: An Oracle database can be protected by Oracle Data Guard and Zerto Virtual Replication simultaneously, as Oracle Data Guard runs inside the OS and Zerto Virtual Replication runs in the hypervisor, so there is no conflict between the two solutions.

Oracle Licensing Issues When Using Zerto Virtual Replication

This section serves solely as recommendations and does not represent any legally binding certification of compliance for Oracle licensing. For further clarification, Zerto recommends that you contact your Oracle licensing manager or utilize a third party Oracle licensing specialist.

The basic tenet of an Oracle Licensing Service Agreement (OLSA) is that Oracle databases need to be licensed where installed and/or running, not where they can run. This phrasing can cause ambiguity when configuring virtualization and disaster recovery for virtualized Oracle databases.

Note: The storage on which an Oracle database VM runs, or the hypervisor hosts that have access, have no bearing on the OLSA and so shared storage configurations are not factors that need to be configured in the OLSA.

See the following sections:

- [“Production Site Licensing”, on page 2](#)
- [“Recovery Site Licensing”, on page 2](#)
- [“VMware Support”, on page 3](#)
- [“Hyper-V Support”, on page 4](#)
- [“Amazon Web Services \(AWS\) Support”, on page 4](#)

Production Site Licensing

Oracle is typically licensed per physical processor and it is required to license all processors on the hypervisor host on which the Oracle database VM was installed. The Oracle VM should be kept running on this host for this reason. It is recommended to license all of the processors in a second hypervisor host to allow High Availability failover, vMotion (in VMware vSphere environments), or Live Migration (in Microsoft Hyper-V environments) between hypervisor hosts without breaching Oracle licensing terms.

If HA and DRS rules are utilized to keep the Oracle database VMs on both the hypervisor hosts that have Oracle licensing, then these hypervisor hosts can form part of a larger cluster. It is not a requirement to maintain a separate Oracle hypervisor cluster, or license all of the hosts in a cluster, only ensure that the hosts on which the Oracle database VMs run are licensed.

Note: It is technically possible to license only the processors in one hypervisor host, but the Oracle VMs would need to be shut down when performing host maintenance and would not have any high availability capabilities to avoid breaching licensing terms; therefore, this is not recommended.

Recovery Site Licensing

Zerto Virtual Replication maintains a replica of the virtual machine data and a journal of the data being changed within the time frame specified in the recovery site. No virtual machine is registered or powered-on in the inventory until a failover, move, or test operation is initiated. From a licensing perspective, a Zerto Virtual Replication replica virtual machine is a cold offline backup, but this does not mean the recovery site hypervisor host should not be licensed.

If a failover test is performed, the Oracle database VM is running in the recovery site host and by the basic tenet of the OLSA, this requires the processors to be licensed. As Zerto recommends regular failover testing to ensure a successful recovery from disaster, Zerto also recommends licensing one recovery site hypervisor host. Zerto Virtual Replication should be configured to replicate only the Oracle VMs to the licensed hypervisor host. It is possible to change the target host in the Zerto Virtual Manager to continue replication during hypervisor host maintenance without invalidating Oracle licensing, but a failover, move, or test operation should not be performed until the replication target has been changed back to the Oracle licensed hypervisor host.

Virtual machines being tested for failover by Zerto Virtual Replication cannot be vMotioned between hosts in a recovery site and therefore it is not required to license all the hosts in the recovery site hypervisor cluster. Neither is it required to maintain a separate cluster for Oracle VMs in a recovery scenario.

If a Zerto Virtual Replication failover or move operation is performed, HA and DRS rules should be applied to the failover virtual machine to ensure that the virtual machine only runs on licensed hosts in the recovery site. This can be done manually or as part of a post-failover script on the VPG.

If an Oracle database VM is never failed over, migrated, or run in a test failover operation, the recovery site hypervisor does not need to be licensed. However, Zerto recommends licensing one hypervisor host in the recovery site to enable testing, moving, and failing over the Oracle database virtual machines.

VMware Support

Zerto Virtual Replication can protect Oracle Database virtual machines running on VMware vSphere.

VMware states the following with regards to Oracle Database support:

“Oracle has a support statement for VMware products that is honored around the world. While there has been much public discussion about Oracle’s perceived position on support for VMware virtualization, our experience is that Oracle Support upholds its commitment to customers, including those using VMware virtualization in conjunction with Oracle products.

The following are some of the key facts about Oracle Support:

- *Oracle RAC support is now included for Database 11.2.0.2 and later.*
- *Known issues - Oracle Support will accept customer support requests for Oracle products running on VMware virtual infrastructure if the reported problem is already known to Oracle. This is crucial—if you are running Database 9i, 10g, or another product with a long history, the odds are in your favor that Oracle has seen your problem before. If they have already seen it, they will accept it.*
- *New issues - Oracle Support reserves the right to ask customers to prove that “new issues” attributed to Oracle are not a result of an application being virtualized. This is reasonable, as this is essentially the same policy that other ISVs use to some degree. It is key to look at the history of Oracle Support with regard to new issues.*
- *Certification - VMware vSphere is a technology that resides under the certified Oracle stack (unlike other virtualization technologies that alter the OS and other elements of the stack). As a result, Oracle cannot certify VMware virtual infrastructure. However, VMware is no different in this regard from an x86 server—Oracle doesn’t certify Dell, HP, IBM, or Sun x86 servers.*

VMware recommends that customers take a logical approach and test Oracle’s support statement. Begin with pre-production systems, and as issues are encountered and SRs are filed, track Oracle’s response. Our experience is that customers see no difference in the quality and timeliness of Oracle Support’s response.”

Best Practices with VMware

The recommendations from VMware for working with Oracle are:

- Create a computing environment optimized for vSphere.
- Create golden images of optimized operating systems using vSphere cloning technologies.
- Allow vSphere to choose the best virtual machine monitor based on the CPU and guest operating system combination.
- Set memory reservations equal to the size of the Oracle SGA.
- Use large memory pages.
- Use as few virtual CPUs (vCPUs) as possible.
- Enable hyperthreading for Intel Core i7 processors.
- Enable jumbo frames for IP-based storage using iSCSI and NFS.
- Create dedicated datastores to service database workloads.
- Use VMware vSphere VMFS for single instance Oracle database deployments.
- Align VMFS properly.
- Use Oracle automatic storage management.
- Use your storage vendor’s best practices documentation when laying out the Oracle database.
- Avoid silos when designing the storage architecture.
- Use Paravirtualized SCSI adapters for Oracle data files with demanding workloads.
- Use the VMXNET family of Paravirtualized network adapters.
- Separate infrastructure traffic from virtual machine traffic for security and isolation.
- Use NIC teaming for availability and load balancing.
- Take advantage of Network I/O Control to converge network and storage traffic onto 10GbE.
- Use vCenter and/or the esxtop/resxtop utility for performance monitoring in the virtual environment.
- To minimize time drift in virtual machines follow guidelines in the relevant VMware Knowledge Base articles.

For more details about running Oracle databases on VMware, refer to

https://www.vmware.com/files/pdf/partners/oracle/Oracle_Databases_on_VMware_-_Best_Practices_Guide.pdf.

Hyper-V Support

Zerto Virtual Replication can protect Oracle Database VMs running on Microsoft Hyper-V Server 2012 R2.

For more details about running Oracle databases and Hyper-V support, refer to https://blogs.oracle.com/cloud/entry/oracle_and_microsoft_join_forces.

Amazon Web Services (AWS) Support

Zerto Virtual Replication can protect Oracle Database VMs to AWS for recovery as EC2 instances.

For more details about running Oracle databases and AWS support, refer to <http://www.oracle.com/technetwork/topics/cloud/faq-098970.html#support>.

Zerto Virtual Replication Support for Oracle Database Features

This section describes the interaction between Zerto Virtual Replication and commonly used Oracle Database features.

See the following sections:

- [“Automatic Storage Management \(ASM\)”, on page 4](#)
- [“Oracle Clustered File System \(OCFS\)”, on page 4](#)
- [“Oracle Real Application Clusters \(RAC\)”, on page 5](#)

Automatic Storage Management (ASM)

ASM provides a clustered file system with volume management for Oracle databases. Disk groups are limited by the performance of the slowest member and so it is recommended to place all protected virtual disks used in a disk group on storage with a similar I/O profile.

Oracle best practice dictates that two ASM disk groups should be created, one with a sequential I/O profile for log files, and one with a random I/O profile for data files. By selecting an Oracle database virtual machine for protection, Zerto Virtual Replication automatically protects all virtual disks assigned to the virtual machine, and therefore all disk groups are included in the protection. As Zerto Virtual Replication maintains write-order fidelity of all the disks within a virtual machine, and between virtual machines in the same VPG, protecting a virtual machine with multiple disk groups recovers all the disk groups to the exact same point-in-time. The journal never presents a checkpoint available for selection where the disk groups are at different points-in-time.

Zerto best practice: In a VMware vSphere environment, it is not recommend to use ASM failure groups. The Oracle Databases on VMware Best practices guide states:

“Do not use Oracle ASM failure groups. Oracle failure groups consume additional CPU cycles and can operate unpredictably after suffering a disk failure. When using external redundancy, disk failures are transparent to the database and consume no additional database CPU cycles, because this is offloaded to the storage processors.”

Oracle Clustered File System (OCFS)

OCFS is a shared disk cluster technology for Linux used with Oracle Real Application Clusters before ASM was released. Zerto Virtual Replication does not support protecting virtual machines utilizing OCFS.

Oracle Real Application Clusters (RAC)

Oracle RAC utilizes shared virtual disks for CRS, voting, data, and redo logs between 4 virtual machine nodes. Zerto Virtual Replication does not support protecting virtual machines utilizing shared disks, therefore Zerto Virtual Replication does not support protecting Oracle RAC nodes.

Optimizing Replication Using Zerto Virtual Replication

Zerto Virtual Replication replicates only the block level changes on a per-virtual machine basis and always works asynchronously. Only copies of the writes are sent asynchronously to the local VRA in the memory of the hypervisor host.

Writes to and from the local storage are not delayed, so there is *no* measurable performance impact on Oracle VMs when protecting them with Zerto Virtual Replication.

See the following sections:

- [“Oracle Temporary Database Optimization”, on page 5](#)
- [“WAN Sizing”, on page 5](#)
- [“Database Consistency”, on page 5](#)
- [“VRA Configuration Best Practices”, on page 8](#)
- [“VPG Configuration Best Practices”, on page 8](#)

Oracle Temporary Database Optimization

Zerto Virtual Replication utilizes a temp data disk feature to optimize the replication traffic of Oracle temporary databases and it is enabled on a per-virtual disk basis. When enabled, an initial copy of the virtual disk is replicated to the recovery site, but no subsequent changes to the disk are replicated. This ensures that any application that uses the virtual disk for an Oracle temporary database sees the disk and original files in a recovery scenario without requiring any manual reconfiguration. The temp data disk feature can save over 50% of replication traffic, journal space usage, and VRA load.

Zerto best practice: Keep all Oracle temporary databases on separate virtual disks and utilize the Zerto Virtual Replication temp data disk feature. You can test the impact of the temp data feature on bandwidth usage by editing the VPG settings to turn off the temp data feature.

WAN Sizing

All Oracle database virtual machines should follow standard Zerto sizing recommendations by using the Zerto WAN sizing tool, available on the self-service portal. It is necessary to capture the write rate on all disks in the Oracle database VM to calculate the required bandwidth between the sites.

Zerto best practice: If backups are made on the Oracle database virtual machine to a local disk, Zerto recommends having the backups stored on a separate virtual disk set as a temp data disk during VPG configuration, similar to the recommendation to use a temp data disk for temporary databases. Otherwise, Oracle backups can cause nightly bitmap syncs as the virtual machine is writing a copy of the entire database to a backup file that is being replicated. In this scenario, if the dataset is too large for the WAN link or for the VRA to keep up, then the virtual machine enters a bitmap sync mode.

Database Consistency

Zerto Virtual Replication replicates continuously, maintaining the write-order fidelity of all the block level changes between all virtual machines in a VPG and all disks within each virtual machine. This means that replication is always crash-consistent.

See the following sections:

- [“Database Consistency in a Windows Environment”, on page 6](#)
- [“Database Consistency in a Linux Environment”, on page 7](#)

Database Consistency in a Windows Environment

To maintain database consistency in Oracle virtual machines running Windows, Zerto recommends installing the Zerto PowerShell cmdlets. A PowerShell script can then be scheduled, to create database consistent points-in-time, checkpoints. Use the following process:

1. Windows Task Scheduler starts the PowerShell script.
2. The PowerShell script loads user configured variables.
3. The PowerShell script does the following:
 - a) Connects to the database using SQLPLUS.
 - b) Places the database into backup mode.
 - c) Inserts a checkpoint into the journal for the Oracle VPG using the Set-Checkpoint cmdlet to mark the database consistent point-in-time.
 - d) Ends the backup mode to resume normal database operations.

Note: For a blog entry describing how to use the Windows Task Scheduler to Run a Windows PowerShell Script, refer to <http://blogs.technet.com/b/heyscriptingguy/archive/2012/08/11/weekend-scripter-use-the-windows-task-scheduler-to-run-a-windows-powershell-script.aspx>.

The following is a sample PowerShell script for creating database consistent points-in-time in Windows:

```
# Start of script - Configure the variables below
$OracleUser = "system"
$OraclePassword = "password"
$OracleServer = "OracleServer"
$ZVMServerIP = "123.123.123.123"
$ZVMPort = "9080"
$ZVMUser = "administrator"
$ZVMPassword = "password"
$VPGName = "Oracle"
$CheckpointTag = "HotBackup"
# Nothing to configure below here, load Zerto PowerShell commands
add-pssnapin "Zerto.PS.Commands"
# Connect to Oracle and place the database in hot backup mode
$OracleConnectionString = $OracleUser + "/" + $OraclePassword + "@" + $OracleServer
sqlplus $OracleConnectionString
ALTER DATABASE BEGIN BACKUP;
EXIT
# Insert a checkpoint in the Zerto journal for the defined VPG
Set-Checkpoint -VirtualProtectionGroup $VPGName -Tag $CheckpointTag -ZVMIP $ZVMServerIP -ZVMPort
$ZVMPort -Username $ZVMUser -Password
$ZVMPassword
# Connect to Oracle and resume normal database operations
sqlplus $OracleConnectionString
ALTER DATABASE END BACKUP;
EXIT
# End of script
```

The frequency with which database consistent points-in-time can be inserted is dependent on the performance of the Oracle databases and services, not Zerto Virtual Replication. In theory, Zerto Virtual Replication could support inserting database consistent checkpoints every few seconds, since Zerto Virtual Replication just inserts a point-in-time marker while database consistency is handled by Oracle hot backup mode.

Zerto best practice: Start with daily database consistent points-in-time and build up to reach the maximum level that the databases and services can handle.

Upon recovering a Windows Oracle database virtual machine from a database consistent checkpoint it is required to release the database from backup mode to continue normal operation as it will not start automatically on boot. The following sample command shows how this can be done:

```
sqlplus / as sysdba
startup mount;
alter database end backup;
alter database open;
```

Database Consistency in a Linux Environment

To insert database consistent points-in-time in Linux virtual machines, Zerto recommends installing the Zerto PowerShell cmdlets on a Zerto Virtual Manager Windows server virtual machine, with a PowerShell SSH server to allow remote execution from Linux virtual machines. To write a database consistent point-in-time, use the following process:

1. An Oracle Linux virtual machine runs a bash script that places the database into backup mode.
2. Via SSH, the bash script connects to the Zerto Virtual Manager PowerShell server and executes a PowerShell script.
3. The PowerShell script on the Zerto Virtual Manager inserts a checkpoint into the journal for the Oracle VPG using the Set-Checkpoint cmdlet to mark the database consistent point-in-time.
4. The bash script disconnects the SSH session.
5. The bash script ends the backup mode to resume normal database operations.

The following is a sample Linux script for creating database consistent points-in-time in Linux:

```
#!/bin/bash
# Place database in backup mode
sqlplus "/as sysdba" <<EOF
alter database begin backup;
exit;
EOF
# Run the remote PowerShell script on the PowerShell server via SSH
ssh administrator@123.123.123.123 './oraclevpg.ps1' exit
ENDSSH
# Resume normal database operations
sqlplus "/as sysdba" <<EOF
alter database end backup;
exit;
EOF
# End of script
```

Zerto best practice: To enable remote SSH access to PowerShell cmdlets on the Zerto Virtual Manager, Zerto recommends using the PowerShell Server that can be downloaded from <http://www.powershellserver.com/download/>. For more information about PowerShell Server, see <http://www.powershellserver.com/getting-started>.

The following is a sample PowerShell script that can be run on the Zerto Virtual Manager, initiated from the Linux virtual machine via SSH:

```
#
# Start of script - Configure the variables below
#
$ZVMServerIP = "123.123.123.123"
$ZVMPort = "9080"
$ZVMUser = "administrator"
$ZVMPassword = "password"
$VPGName = "Oracle"
$CheckpointTag = "HotBackup"
#
# Insert the checkpoint into the journal
#
Set-Checkpoint -VirtualProtectionGroup $VPGName -Tag $CheckpointTag -ZVMIP $ZVMServerIP -ZVMPort
$ZVMPort -Username $ZVMUser -Password $ZVMPassword
#
# End of script
#
```

Upon recovering a Linux Oracle database virtual machine from a database consistent checkpoint it is required to release the database from backup mode to continue normal operation as it will not start automatically on boot. The following sample command shows how this can be done:

```
sqlplus / as sysdba
startup mount;
alter database end backup;
alter database open;
```

VRA Configuration Best Practices

Zerto best practice: Install a VRA on each source and target hypervisor node on which the protected Oracle VM can run. There will be a break in replication if the Oracle VM is moved to a hypervisor without a VRA.

Zerto best practice: If multiple Oracle virtual machines that are all highly transactional in their I/O profile require protection, distribute the load evenly between hypervisor hosts, and therefore VRAs, to minimize the I/O load placed on a single appliance.

VPG Configuration Best Practices

When defining a VPG to protect Oracle Database virtual machines, you need to consider the following:

- The virtual machines to include in the VPG and the boot order between these machines. See [“Virtual Machines in the VPG”, on page 8](#).
- [“SLA Considerations”, on page 9](#).
- [“Journaling Considerations”, on page 9](#)
- [“Storage Considerations”, on page 9](#)
- [“Network Considerations”, on page 9](#)

Virtual Machines in the VPG

Place all of the virtual machines that form a single application in the same VPG. If multiple applications use the same database and all require consistency then one VPG should be created with all the application virtual machines, even if the number of virtual machines is large.

If the application virtual machines don't require consistency to the database virtual machine, then create separate VPGs for the applications and the database virtual machine. If multiple database virtual machines update each other continuously and are interdependent then they should be placed in the same VPG for cross-database consistency.

Zerto best practice: Boot ordering within a VPG should be applied to ensure that the Oracle Database virtual machines boot first and that the database services are online before any application servers that use the database services begin booting.

SLA Considerations

The VPG SLAs should be configured to allow a sufficient amount of data retention for recovery to previous points-in-time, give sufficient leeway on the RPO to allow spikes in data change rates without raising false alerts, and to send reminders to test on a reasonable schedule.

Zerto best practice: Start with a five minute RPO and leave the VPG protected for at least one working day in this configuration. Configure email alerts to be sent immediately when there is a breach of the RPO SLA during this time period. After 24 hours, the RPO should be altered up or down depending on the SLA requirements of the business. That is, if the RPO regularly hits 5 minutes but typically lasts for 10 minutes before returning to normal, then the RPO alert and SLA should be changed to 15 minutes or the bandwidth between the sites should be increased.

Zerto does not recommend changing the RPO alert to less than 20 seconds as this leaves no leeway for any spikes in data change rates or network connectivity, and can result in unnecessary RPO alerts being generated. Neither should the RPO alert be configured to the exact SLA defined by the business as this gives no leeway on the replication alert vs. the SLA defined. For example, if the business requires an RPO of 30 minutes then best practice would be to configure the RPO alert at 20 minutes to allow a comfortable leeway. If the RPO alert was constantly breached in this scenario then the situation can be remedied by increasing available bandwidth without breaching the business SLAs.

Zerto best practice: The default test frequency of 6 months is recommended when first configuring the protection of a VPG. It should be increased or decreased based on the frequency of testing required by business defined SLAs.

Journaling Considerations

Performance of the journal storage is important as the protected virtual machine data is inserted with the same I/O profile and the journal can be heavily used during the pre-commit and commit phases of a failover, before committing the failover.

Maintaining space availability on the journal datastore is also key to maintaining replication. Filling the datastore will break replication, whereas multiple journals hitting their size limit will only reduce the time frame of the journal, but it will not break the replication.

Zerto best practice: When configuring journaling for Oracle Database virtual machines, start with a one hour journal. After one hour of production data has been replicated, use the size of the journal to calculate the space required over the time period the business dictates as an SLA requirement for data retention. The journal should then be increased to the desired time frame for retention.

Zerto best practice: The storage used for the journal should have the same I/O profile as the storage containing the replica virtual machine data.

Zerto best practice: The journal storage should have enough space to accommodate the data change rate over the time period specified. To cope with spikes in data change there should be at least 30% of disk space free on the storage and the maximum journal size should be configured on each journal to ensure no journal can fill the storage specified.

Storage Considerations

The datastore configuration of the protected Oracle database virtual machine should be matched in the recovery site to ensure the same performance in a recovery scenario. Since Zerto Virtual Replication is storage agnostic, it is possible to replicate from any storage to any storage, and therefore any LUN to any LUN configuration. However, this does not take into account the performance requirements of the protected virtual machines.

Zerto best practice: Use different storage for the journal and for the replica virtual machine data.

The number of virtual machines running per storage, as recommended by the storage vendor, should be taken into account if you are replicating multiple virtual machines to a consolidated number of storages in a target site.

Network Considerations

The networking configuration of the Oracle Database virtual machines running in production should be mirrored in the recovery site. If a different IP address is required in the recovery site, the VPG can be configured on a per virtual NIC basis to

change the IP address automatically as part of a failover, move, or test operation. If an IP change is required, ensure that the `listener.ora` and `tnsnames.ora` configuration files do not have hard coded IP addresses, but use hostnames.

For failover testing, a mirror of the production networking should be configured in an isolated network segment to allow failover testing with no impact on production.

Failover Testing

Zerto recommends performing failover testing to validate the recovery process, the recovery time objective, data consistency, and operation of the Oracle Database VM in a recovery scenario. Failover testing of Oracle Database virtual machines can be done in working hours since production is not shut down and there is no break in replication. Ensure that isolated networks are specified for failover testing to ensure that the Oracle Database virtual machines cannot communicate with production applications and clients. Otherwise the failover test can impact production as the failover test virtual machine could receive database updates from production that would be deleted when the failover test operation is stopped.

The size of a failover test operation for Oracle Database virtual machines is simply the size of the changes written to the test virtual machines during the test. All changes from the failover test virtual machines are deleted when the test is stopped.

Zerto best practice: Do not run a failover test for longer than the journal history is configured. If a longer test is required, then use the offsite-clone feature to create an entire VM-level clone of the virtual machines in the VPG onto any storage from any specified point-in-time available in the journal.

If recovering to a database consistent point-in-time, the failover test operation enables validation of the Oracle Database starting in backup mode to verify the consistency. It then enables validating removing the database from backup mode and mounting the database to ensure a successful recovery.

Zerto best practice: Perform test failovers to both crash and database consistent points-in-time to verify the recovery of both on a per Oracle database basis. The ability to recover from either should be recorded and included in a disaster recovery plan for future decision making in a failover scenario when deciding what point-in-time to recover to and the amount of data loss to accept.

If IP address changes are required as part of a failover or move operation, they should also be performed in the failover test operation to ensure a successful recovery. This allows the IP change process to be verified and the full disaster recovery scenario to be simulated.

Failover, Move, and Failback

To ensure a successful failover, move, or failback operation, the operation should always use a commit policy set to `None` to allow an indefinite time period to verify the successful recovery of Oracle Database virtual machines.

After a successful recovery has been confirmed, the VPG should be committed to the point-in-time selected. If the verification is not successful, then a different point-in-time should be selected from the journal and verification to this new point-in-time performed.

Zerto best practice: Before performing a move operation for Oracle Database virtual machines, take a database consistent checkpoint using the sample scripts in this document. This ensures that there is a guaranteed point-in-time in the journal if the database services do not shut down in a clean state.

In a failover operation you can failover to the most recent point-in-time to minimize data loss or to failover to the last database consistent point-in-time and accept the data loss for the time interval between the two points.

Deciding which point-in-time to use when performing a failover operation is both a user and business decision. Using the `None` commit policy allows you to verify failing over to the most recent crash-consistent point-in-time and, if this does point is not valid, you can perform a failover to the last database consistent point-in-time.

Zerto recommends ensuring that sufficient leeway is built into the RTO SLA defined for the business to allow for the virtual machines to be powered on, database services started, and the recovery verified as successful before re-commencing business operations. For example, if the virtual machine and VM service RTO take 15 minutes to start, a recommended RTO SLA for the

Best Practices When Protecting an Oracle Database - Version 6.5

business would be 45 minutes to allow verification and subsequent recovery to a previous point-in-time if the first point-in-time does not recover to a working state.

Sufficient time should be allowed when configuring failback to allow for a delta sync to read both source and target disks and then replicate the delta changes. The speed of a delta sync operation for Oracle Database VMs is dependent on the following factors:

- Amount of data changed since the failover operation
- Speed of the IP link and available bandwidth for replicating the change
- Speed of both the source and target storage
- Size of the virtual disks and amount of free space
- The latency to the source and target storage

The delta sync process does not run as fast as the VRAs can run, it runs as fast as possible without impacting the performance of the Oracle Database virtual machine. This is done by dynamically scaling up and down the read rate to the protected virtual machine disks depending on the latency to the storage to ensure no performance impact of the delta sync process. This ensures that the performance of Oracle Database virtual machines in a recovery scenario is the highest priority with subsequent failback to production a background task that can be left to run until the VPGs are in a protected state ready for failback.

A failback operation is typically a planned migration back to the production site.

Zerto best practice: Perform a failback out of working hours. The failback involves shutting down the virtual machines now running in the disaster recovery site. A database consistent point-in-time should be added before initiating the failback procedure. The failback should also be tested by a test failover process to validate the process and recovery time objective.

Failover, Move, and Failback

Zerto helps customers accelerate IT transformation by eliminating the risk and complexity of modernization and cloud adoption. Replacing multiple legacy solutions with a single IT Resilience Platform, Zerto is changing the way disaster recovery, data protection and cloud are managed. With unmatched scale, Zerto's software platform delivers continuous availability for an always-on customer experience while simplifying workload mobility to protect, recover and move applications freely across hybrid and multi-clouds. Zerto is trusted by over 6,000 enterprise customers globally, and is powering resiliency offerings for Microsoft Azure, IBM Cloud, AWS, Sungard and more than 350 cloud services providers.

For assistance using Zerto Virtual Replication, contact: @Zerto_Support.

Learn more at Zerto.com

Copyright © 2018, Zerto Ltd. All rights reserved.