

This document is intended to cover best practices of protecting a mission critical Microsoft SQL Server using Zerto Virtual Replication. All information is correct for Zerto Virtual Replication version 3.1 and higher.

Zerto Virtual Replication can replicate Microsoft SQL Server virtual machines with both crash- and application-consistent points in time. These are indicated as checkpoints in the journal in the recovery site.

**Note:** Zerto Virtual Replication does not support MSCS Clusters with an active/active cluster configuration. For full details about support for MSCS, refer to *Zerto Virtual Replication Best Practices When Protecting Virtual Machines Running MSCS*.

## Zerto Virtual Replication Handling of Swap Disks

In the *STORAGE* step of the *Create VPG* wizard, if you specify that storage is a swap disk, Zerto Virtual Replication performs an initial synchronization of the virtual disk (VMDK or VHD/X) and the files it contains, but does not replicate any subsequent changes to the selected virtual disk. This ensures that a replicated copy of the disk and the files on the disk are available when booting from a recovery operation so that configuration of the disk in the recovered virtual machine is not needed.

### Windows Page File

The Windows Page File should be placed on a separate virtual disk with a Paravirtual SCSI controller. The virtual disk should be indicated as a swap disk in Zerto Virtual Replication after the page file has been moved to the new virtual disk.

### TempDB

**Zerto best practice:** The SQL TempDB should be placed on a separate virtual disk with a Paravirtual SCSI adapter. As the TempDB is recreated upon a reboot, its virtual disk should be defined as a swap disk in Zerto Virtual Replication. For instructions on how to correctly move TempDB before indicating its virtual disk as swap, see:

[www.dbadiaries.com/how-to-move-tempdb](http://www.dbadiaries.com/how-to-move-tempdb)

**Note:** Not performing this configuration can double the replication of data. It can result in large RPO spikes during heavy transactional updates and will waste journal space.

## Disk Configuration

Zerto recommends using a different virtual disk size for each disk to aid identification from Windows.

**Note:** The VMware recommended disk configuration is C: drive on 0:0 with the default adapter, SQL database VMDKs on 1:0 Paravirtual Adapter, SQL Log VMDKs on 2:0 Paravirtual Adapter and the TempDB and Page File on 3:0/3:1 Paravirtual Adapter.

## Backups

Backups of the protected SQL databases are not required within the time frame configured for the Zerto Virtual Replication journal as you can recover SQL from a failover or failover test within that time window.

**Zerto best practice:** Backups should not be performed to a local virtual disk. If backups are required to a local virtual disk, create a new thin virtual disk and define this virtual disk as a swap disk in Zerto Virtual Replication before any backups are written to disk.

**Note:** Backing up SQL to a local virtual disk without indicating this as a swap disk in Zerto Virtual Replication will result in RPO spikes whenever a backup job is run and will fill the journal with redundant checkpoints.

## Log Shipping

Log shipping is not required on a protected virtual machine as Zerto Virtual Replication allows failover to any point in time up to 14 days in the past with increments every time something has changed.

**Zerto best practice:** If log shipping is required, it should not be performed to a local virtual disk on the protected VM. The virtual disk should be indicated as a swap disk in Zerto Virtual Replication before creating any logs on the virtual disk.

**Note:** Not performing this configuration will result in doubling the replication of data.

## Crash Consistency

Zerto Virtual Replication maintains write-order fidelity ensuring the disks in the target site are always crash-consistent. Only confirmed writes are replicated. This means Zerto Virtual Replication only replicates data when the OS has confirmed the change has been written to the local disk.

## Transaction Consistency

Zerto Virtual Replication maintains consistency between the virtual machines in a VPG and the virtual disks within a virtual machine. Transaction consistency requires the Zerto VSS agent to be installed on the SQL VM. Transaction consistent points-in-time are indicated as checkpoints in the journal in the recovery site.

**Note:** Only one transaction consistent point-in-time per database virtual machine in a VPG can be selected as the point-in-time to recover to. VPGs should therefore be designed accordingly.

## Scripting Transactional Consistency

Once the Zerto VSS agent has been installed, the following batch script should be run inside the SQL VM:

```
"C:\Program Files\Zerto\ZertoVssAgent\zertovssagent.exe" 1.x.x.x 9080 2.x.x.x 9080 Scripted-TransConsistent-Checkpoint
```

Where 1.x.x.x is the source ZVM IP address and 2.x.x.x is the recovery ZVM IP address. The text at the end is just a description of the checkpoint created.

**Note:** The default installation folder for ZertoVssAgent is C:\Program Files\Zerto for the 64-bit version and C:\Program Files (x86)\Zerto for the 32-bit version.

The script should be scheduled using the Windows Task Manager with *Run as administrator* privileges. Zerto recommends running this script as often as you can without impacting the production application.

### ABOUT ZERTO

Zerto is committed to keeping enterprise and cloud IT running 24/7 by providing scalable business continuity software solutions. Through the Zerto Cloud Continuity Platform, organizations seamlessly move and protect virtualized workloads between public, private and hybrid clouds. The company's flagship product, Zerto Virtual Replication, is the standard for protection of applications in cloud and virtualized datacenters.