
Microsoft SQL Recovery

<https://campus.barracuda.com/doc/78155621/>

The requirements for each type of SQL recovery objective varies. Read each section below to find the recovery objective that best meets your needs and follow the requirements and considerations closely before performing SQL database recovery.

All SQL Database Recovery Options

The following requirements and considerations apply to all SQL database recovery operations:

- The entire database is recovered with each recovery operation. Barracuda Networks does not support granular recovery of SQL database records.
- A database must be recovered to an SQL instance that is the same version or later than that of the original SQL instance. Databases cannot be recovered to an older SQL version.
- Any existing database of the same name that resides on the recovery target is overwritten during the recovery operation (even if you specify a different path). To retain the original database, you must modify the recovery database name in the **Restore** dialog.

The following sections detail the requirements and considerations for different types of recovery operations.

Full Backup of a Database

The following requirements and considerations apply to recovering full backups of SQL databases:

- A full backup can be recovered to the original location (the instance where the backup was taken) or to an alternate location.
- If applicable, see *System Databases*, *Recovery to a Clustered Instance*, and *Availability Groups* additional requirements.

Differential or Transaction Log Backup of a Database

The following requirements and considerations apply to recovering SQL differential or transaction log backups:

- The backup can only be recovered to the original location (the instance where the backup was

taken).

- During recovery, all previous backups in the log chain are also recovered. This means that when recovering a transaction backup, all previous transaction backups, the latest differential (if any), and the parent full are also recovered.
- If applicable, see *System Databases*, *Recovery to a Clustered Instance*, and *Availability Groups* for additional requirements.

System Databases

The following requirements and considerations apply to recovering SQL system databases:

- The master, model, and msdb system databases can only be recovered to their original SQL instances and names. The recovery job overwrites the existing database (even if you specify a different path).
- The master, model, and msdb databases must be recovered individually.
- To recover the master database, you must first stop the SQL instance. See the Microsoft TechNet article [How to Stop an Instance of SQL Server for details](#).

Recovery to a Clustered Instance

In an SQL cluster, an SQL Server failover cluster instance is installed into a Windows Server Failover Cluster (WSFC). The cluster's databases reside on shared storage that is accessible to each server node.

When recovering a SQL backup to a clustered instance, you must recover to a clustered storage resource that is accessible to all nodes in the cluster and is a dependency of the SQL clustered instance you are restoring to. As long as you leave the Specify Path field empty (on the Recover dialog), the backup is recovered to the original clustered storage resource. If you enter a path, be sure to specify a clustered storage resource that is a dependency of the SQL clustered instance you are restoring to. (If you specify a local volume, the recovery fails with a SQL VSS Writer error.)

Because it is the instance that is clustered, and not the hosted database, a backup taken on a clustered instance is no different than one taken on a non-clustered instance. As with all user database backups, a full backup can be recovered to any eligible instance (which can be clustered or non-clustered), and a differential or log backup must be recovered to the original location.

Availability Groups

The following additional requirements and considerations apply to recovering SQL availability group databases:

- The database must not be present in an availability group on the target server. If the database already exists, you must remove it from the availability group, then delete the mirrored copies from all secondary replicas.
- A database of the same name that is in the restoring state must not be present on the target instance. If found, you must delete the database before running the recovery.
- A differential or transaction log backup must be recovered to the original instance, and this instance must have an availability group that is configured with the same listener IP address as the backup's Availability Group asset.
- After recovery, you must manually reconfigure the mirrors and add the recovered database to the availability group. See Microsoft's documentation for details.

© Barracuda Networks Inc., 2024 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.