

Version 1.2

July 2024

Use Cohesity DB Restore to Add an MS SQL Database to MS SQL AG

*Prepare a Standalone SQL Database for Adding
into a SQL AG Group*

ABSTRACT

Go beyond data management. Use your Cohesity Platform backups to seed a database for SQL AG. In this guide, we use Cohesity's SQL database restore feature to prepare a standalone SQL database for its addition into a SQL AG group.

Table of Contents

Introduction to Cohesity MS SQL Database Restore	4
Think About Other Use Cases.....	4
Know All the Details	4
What is Microsoft SQL Server Always On?	5
Terminology	5
Real-World Example.....	6
Visualize Database Restore	7
Add a Database into a SQL AG Relationship	8
Map the Workflow to Success	8
Jump In with an Existing AG Relationship	8
Check the Prerequisites	9
Restore the Database to a Replica Host.....	10
Shield the Restored Database	14
Pause the Protection Job	14
Make the Databases Identical	16
Take Log Backup from Primary Database	16
Restore Log Backup to Secondary Database	17
Tie the Knot with SQL AG	18
Add Database to Existing AG Group on Primary AG Host.....	18
Join Database to Existing AG Group on Secondary AG Host	22
Your Feedback	24
About the Authors.....	24
Document Version History.....	24

Figures

Figure 1: Example—Setting Up a SQL AG Database Between Data Centers	6
Figure 2: Database Restore Workflow	7
Figure 3: AG Nodes—Primary and Secondary Replicas	9

Tables

Table 1: Restore Concepts	5
Table 2: Data Synchronization Options	21

Introduction to Cohesity MS SQL Database Restore

SQL Always On availability groups (AG) is fast becoming the SQL replication technology of choice. Leverage your snapshots to introduce a database into an existing SQL AG group. To date, this has often been tricky, but now, Cohesity Platform's database restore feature makes it easy.

Partnering Cohesity with Microsoft SQL Server® has three major phases:

- Discovery
- Deployment Considerations
- Solution Procedures

You are at the Solution Procedures phase of pairing Cohesity Platform with Microsoft SQL Server.

This guide focuses on the process of preparing, seeding, and introducing a SQL database into an existing SQL AG group, and is aimed at IT and database administrators who manage data protection for Microsoft SQL Servers.

Think About Other Use Cases

The use case in this guide is to introduce a stand-alone database into an existing SQL AG relationship.

There are other implications for this process. Restoring a database to a second SQL server and then making them transactionally identical has several use cases. They are:

- Seeding a database for SQL database mirroring.
- Seeding a database for SQL log shipping.
- Moving a large database from one instance to another for load balancing or upgrading a hardware upgrade.
- Pre-staging databases at a different regional disaster-recovery (DR) site.
- Scaling out the AG relationship by adding a replica.

Know All the Details

There is a lot to know about Microsoft SQL AG technology and how it is applied in the field. Cohesity Platform is uniquely suited to work with SQL AG.

Find out more by reading these articles:

- [Cohesity with MS SQL Server AlwaysOn Availability Group Solution Guide](#)
- [Prerequisites, Restrictions, and Recommendations for Always On availability groups](#)
- [Protect SQL Server with Cohesity—Deployment Configurations Guide](#)
- [SQL Server Instance Prerequisites and Restrictions](#)
- [What is SQL Server AlwaysOn?](#)

What is Microsoft SQL Server Always On?

The SQL Always On availability groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. SQL AG maximizes the availability of a set of user databases for an enterprise. An availability group supports a failover environment for a discrete set of user databases, known as availability databases, that failover together. An availability group supports a set of read-write primary databases and one to four sets of corresponding secondary databases. For more, read [Overview of Always On Availability Groups \(SQL Server\)](#).

Terminology

There are several concepts and terms that are important to understand as you learn about the Microsoft SQL Server backup features in Cohesity Platform.

Table 1: Restore Concepts

TERM	DEFINITION
Protection Job	A Cohesity Protection Job is a backup job that runs repeatedly, based on an associated Policy, to back up data from a source and store it on the Cohesity cluster. A Protection Job can also store the data on External Targets that can be other clusters, cloud object storage, or tape.
AG Replica (Replica)	The term "replica" typically refers to availability replicas. For example, "primary replica" and "secondary replica" always refer to availability replicas. ¹ An <i>availability database</i> is sometimes called a <i>database replica</i> in Transact-SQL, PowerShell, and SQL Server Management Objects (SMO) names.
Primary Replica	The primary replica is assigned the primary role and hosts read-write databases, which are known as <i>primary databases</i> .
Secondary Replica	A secondary replica hosts read-only databases, known as secondary databases. Every availability replica is assigned an initial role—the primary or secondary role, which is inherited by the availability databases of that replica. The role of a given replica determines whether it hosts read-write databases or read-only databases. ²

¹ [Overview of Always On Availability Groups \(SQL Server\)](#), Microsoft

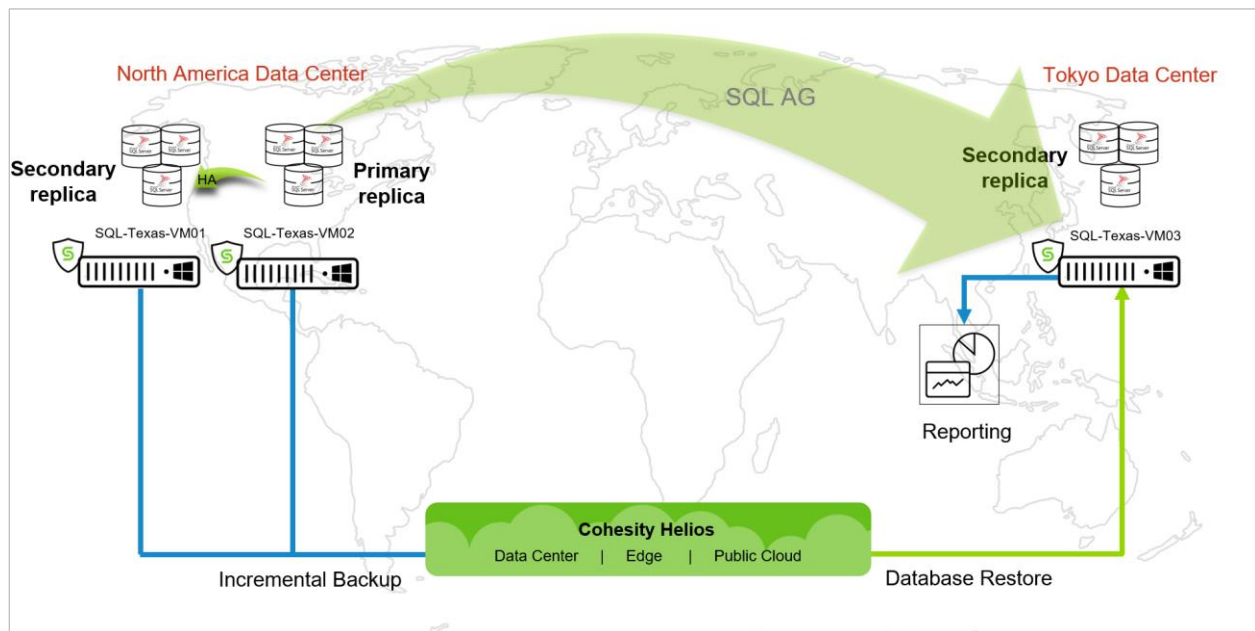
² Ibid

Real-World Example

In our example, the customer is headquartered in Dallas and has a branch in Tokyo.

Both locations have a data center (DC); one in Texas (North America) and the other in Tokyo (Japan). The customer wants to set up an AG database between Dallas and Tokyo.³ Linking the two locations is an existing SQL AG relationship.

Figure 1: Example—Setting Up a SQL AG Database Between Data Centers



In order to establish a SQL AG relationship between two SQL databases located in Dallas and Tokyo, you must have a transactionally identical copy of the database on both the primary (Dallas) and secondary (Tokyo) servers.

The database administrators (DBAs) in Dallas calculate that it will take three days (72 hours) to back up, copy, and restore the database to the Tokyo data center.

In the time that it takes the DBAs to perform this task, the primary database in Dallas continues to receive transactions. Once the copy of the database is restored to Tokyo, how do we make the secondary database in Tokyo identical to the primary database in Dallas, as required for an AG relationship?

In this case, we will establish a copy of the database in Tokyo using Cohesity's DB Restore feature. After that, establishing a SQL AG relationship is easy.

³ There are several reasons to set up an AG database between Dallas and Tokyo: 1) establish a disaster-recovery site, 2) keep business data in Tokyo up-to-date, and 3) offload business reporting to its Tokyo branch.

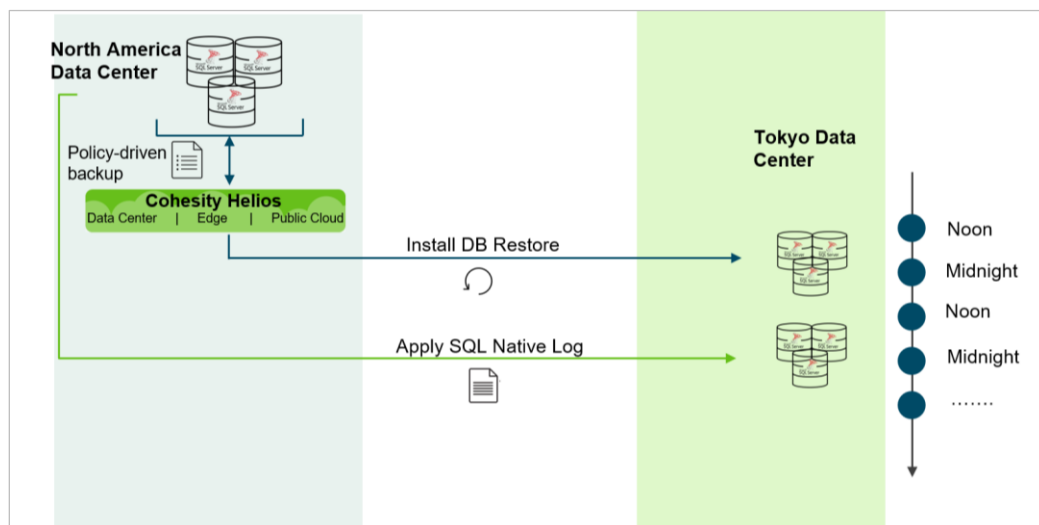
Visualize Database Restore

To continue with our example, the database restore feature is used to copy the database from Dallas over to Tokyo. This is accomplished in a few steps, starting with a restore.

The process is:

1. Initial seeding (restore) of the database from a snapshot on Cohesity Platform over to the SQL instance in Tokyo.
 - The database is attached to the target SQL instance.
 - The database is left in a “recovering” state.
2. Make the databases equivalent.
 - Take a SQL native log backup of the primary database.
 - Apply the SQL native backup to the secondary database (replica).

Figure 2: Database Restore Workflow



Once the restore completes, the Tokyo copy of the database is almost identical to the one in Dallas. The last step will be to make the database in Tokyo identical to the database in Dallas.

Let's start with adding a database into a SQL AG relationship, in the [next chapter](#).

Add a Database into a SQL AG Relationship

In the following sections, we will seed a database (which we'll call "NewBraunfels") over to a second SQL instance using Cohesity's database restore feature, and then introduce it into a SQL AG relationship.

Map the Workflow to Success

Before jumping into the process, it's important to know what all the stages involved.

To add a database into a SQL AG relationship using Cohesity's database restore feature:

1. [Jump in with an existing SQL AG.](#)
2. [Check the prerequisites.](#)
3. [Restore the database to a host replica.](#)
4. [Make the databases identical.](#)
5. [Tie the knot with SQL AG.](#)

IMPORTANT: Cohesity *strongly* recommends you thoroughly test this process in a test environment before applying it in a production environment.

Jump In with an Existing AG Relationship

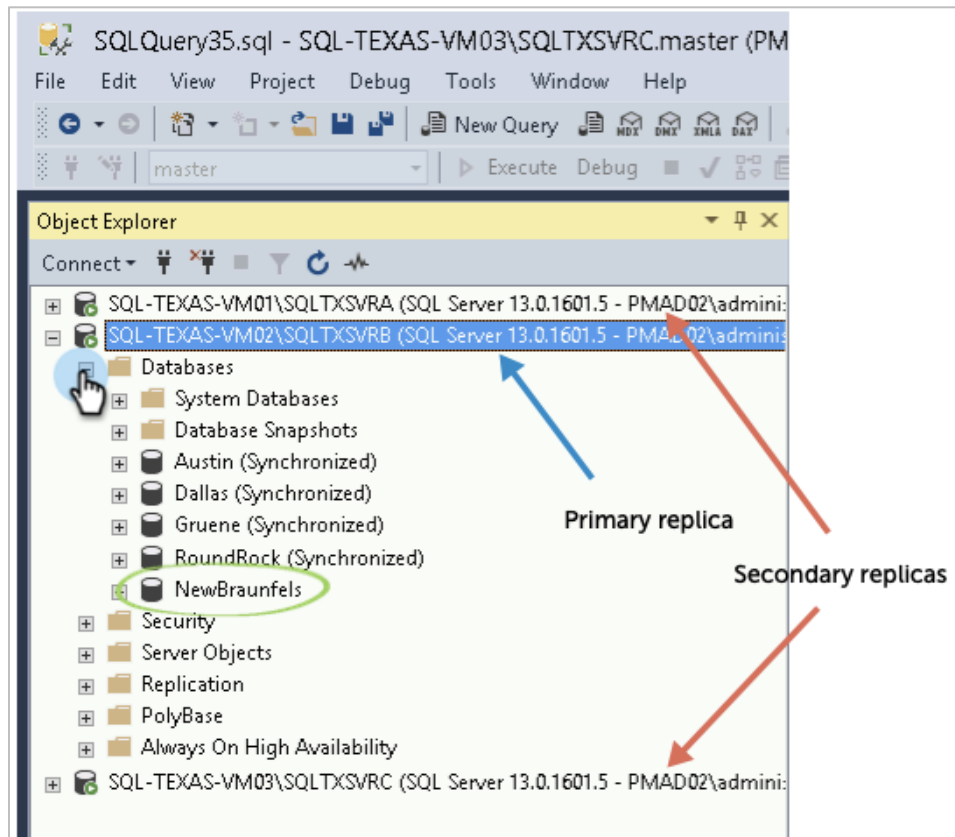
Chances are you are already working with a set of servers that has an existing AG relationship already built. Your first step is knowing the lay of the land so that you can properly place the new database and introduce it into the AG relationship.

We start with an existing SQL AG setup that has three nodes:

- SQL-TEXAS-VM01 (Called "Texas 1") — currently a secondary replica.
- SQL-TEXAS-VM02 (Called "Texas 2") — currently a primary replica in Dallas.
- SQL-TEXAS-VM03 (Called "Texas 3") — currently a secondary replica in Tokyo.

In Figure 3 below, we can see the existing AG nodes, including the primary replica (SQL-TEXAS-VM02) and the two secondary replicas (SQL-TEXAS-VM01 and SQL-TEXAS-VM03). In the primary replica, we find the NewBraunfels database that we wish to bring into the AG relationship.

Figure 3: AG Nodes—Primary and Secondary Replicas



Check the Prerequisites

Because we are starting with an existing SQL AG setup, most of these prerequisites are already established, but it's always good to check them anyway.

- All the SQL instances and the Cohesity cluster must be joined in an Active Directory Domain.
- All the versions of SQL instances are compatible.
- All the instances are participating in a SQL AG relationship.
- All the SQL instances participating in the SQL AG relationship are registered with your Cohesity cluster.
- The database you are working with is a qualified database.
- The database must be in a normal and open state on the Primary SQL AG instance.
- The database is not already participating in a SQL AG relationship.⁴
- The database has a recent file-based incremental snapshot under a Cohesity Protection Job.

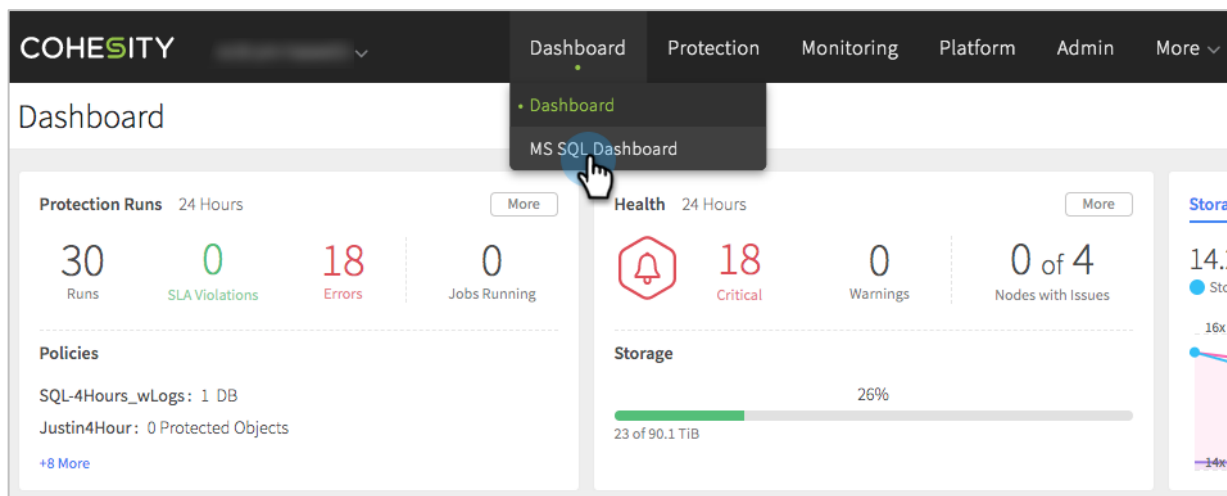
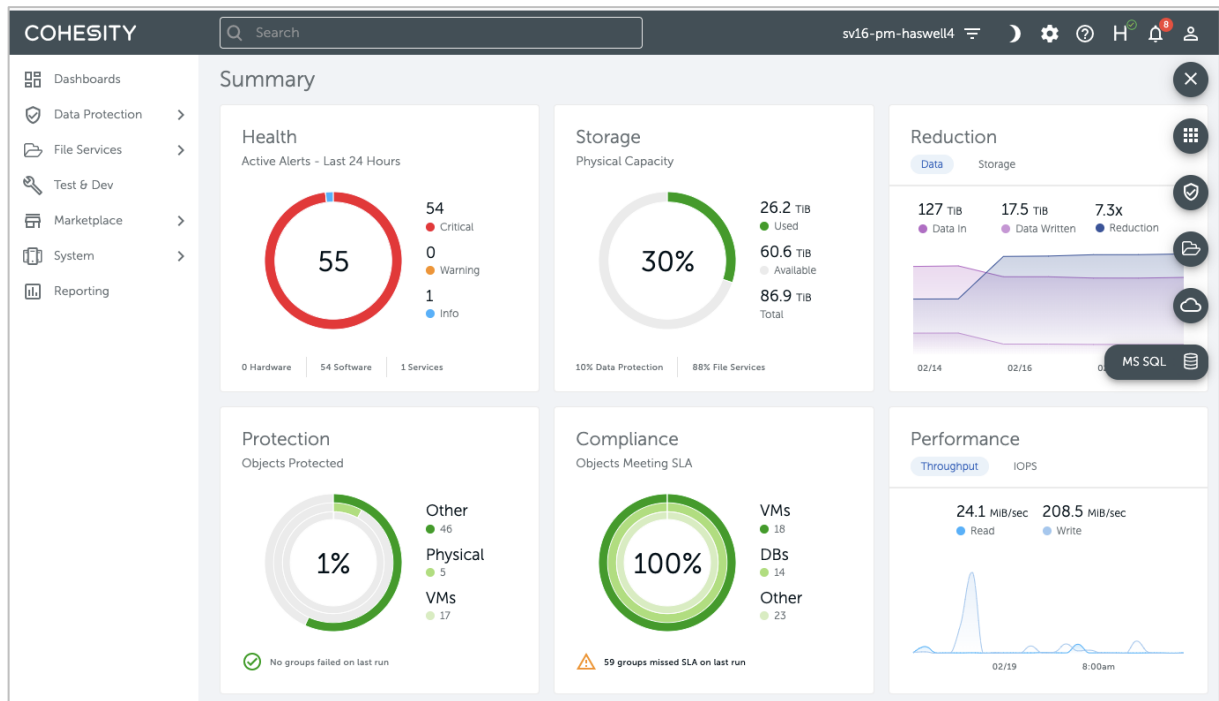
⁴ This is an artificial prerequisite used to maintain simplicity in this guide. There are valid scenarios where a user might want to add another database replica to an existing AG group.

Restore the Database to a Replica Host

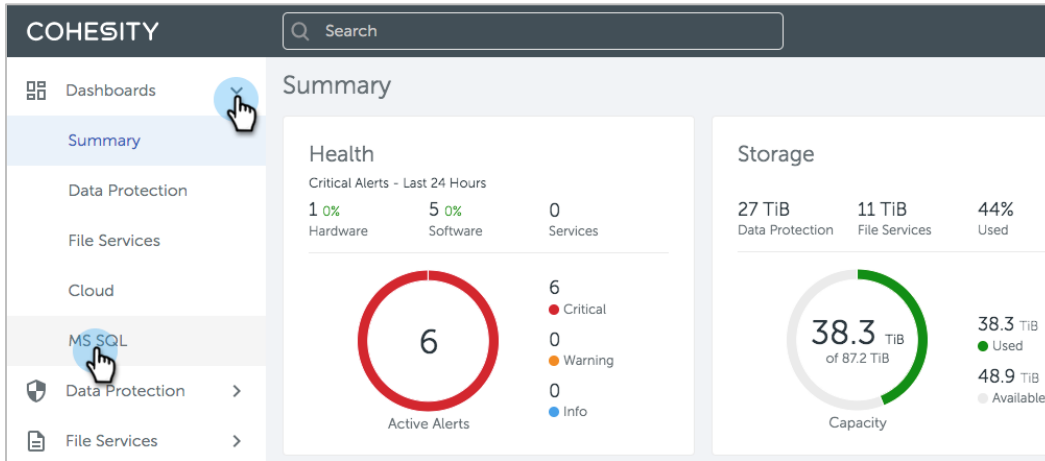
The next step is to restore the database to get an initial copy onto the secondary host.

To restore the database:

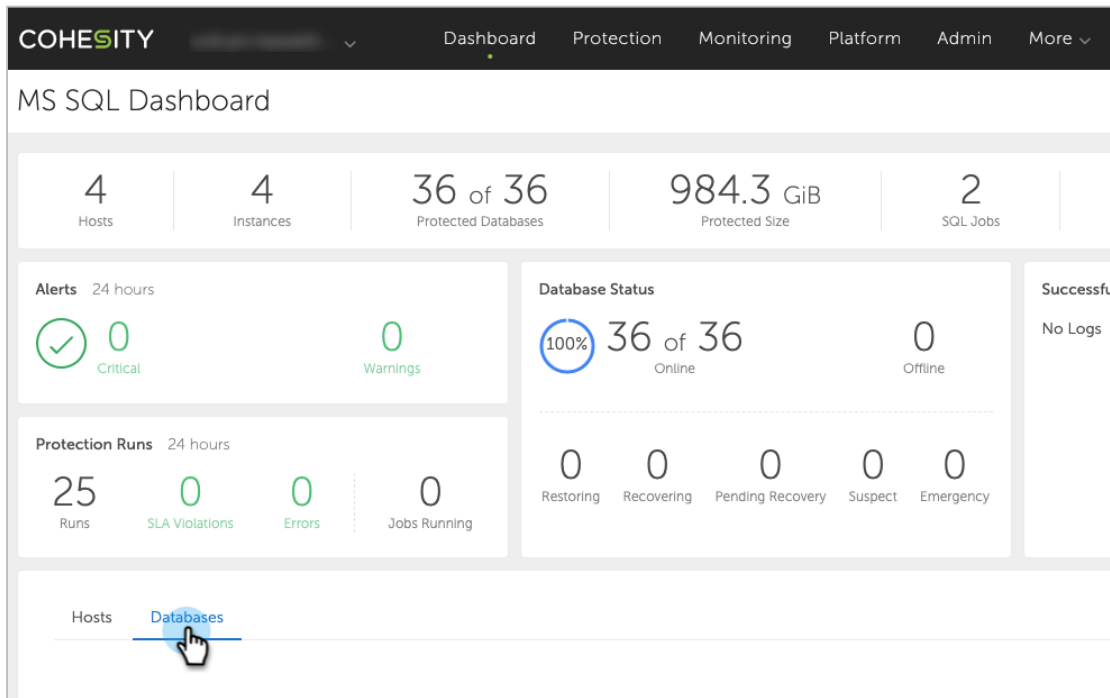
1. Log in to Cohesity Platform.
2. Select **Dashboard > MS SQL Dashboard**.



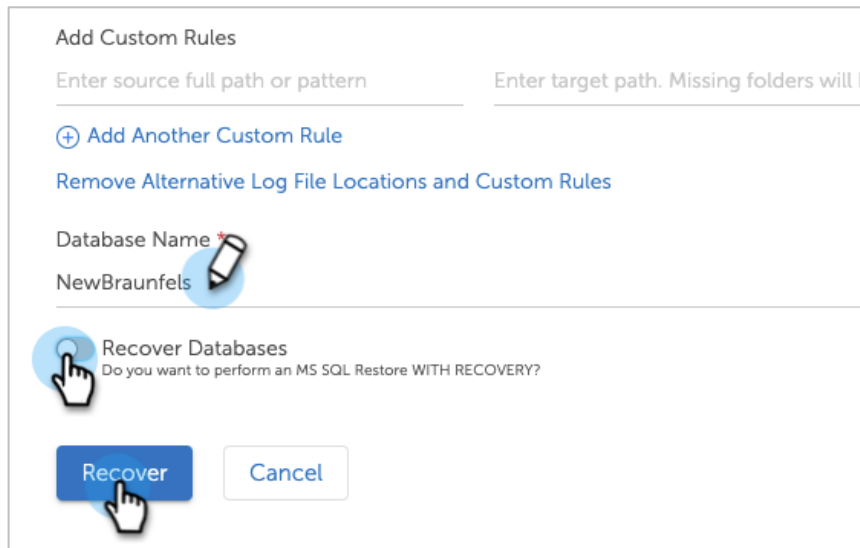
a) In Cohesity version 6.4 and higher, select **Dashboards > MS SQL**.



3. Click the **Databases** tab.



6. In the same form, enter a **Database Name**, disable **Recover Databases**, and click **Recover**.



The screenshot shows a web form for adding a custom rule. At the top, there is a section titled "Add Custom Rules" with two input fields: "Enter source full path or pattern" and "Enter target path. Missing folders will be created". Below these fields are two links: "+ Add Another Custom Rule" and "Remove Alternative Log File Locations and Custom Rules". The "Database Name" field is filled with "NewBraunfels" and has a pencil icon. Below the database name is a checkbox labeled "Recover Databases" with the text "Do you want to perform an MS SQL Restore WITH RECOVERY?". The checkbox is currently unchecked, and a hand cursor is pointing at it. At the bottom of the form are two buttons: "Recover" (highlighted with a hand cursor) and "Cancel".

CAUTION: Do NOT enable **Recover Databases** in this form. The SQL AG database must remain in the “restoring” state.

Shield the Restored Database

During this time, nothing should change the state of the primary database. Therefore, no Protection Job or any other action on the database that would change the state of the database should occur.

Examples of actions that change the state of the database are:

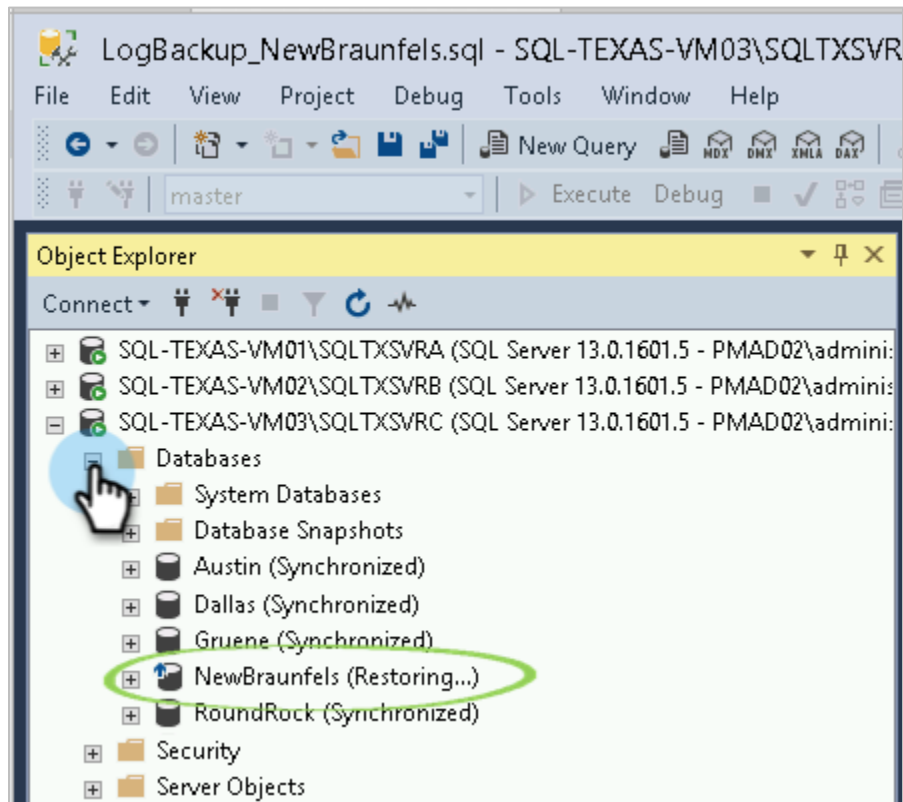
- A database log backup.
- A database full backup.
- Taking the database offline.
- Changing the *recovery* model.
- Changing the database to *read only*.

Pause the Protection Job

Because you do not want a Protection Job to run against the database during the time you are restoring and introducing the database into the AG relationship, you must pause the Cohesity Protection Job that protects the SQL AG servers.

NOTE: Application transactions might still be occurring on the primary database (NewBraunfels, in our example). These will be captured and applied to the secondary copy in the [next task](#). If the application generates a large volume of transactions during this short time, consider briefly pausing the application.

In your [SQL Server Management Studio](#) (SSMS), the restored database should now appear on the TEXAS 03 secondary replica, as a database in the “**Restoring...**” state.



Make the Databases Identical

Finally, because application transactions can still be occurring on the primary database during restore, we need to capture and apply these newest transactions to the secondary copy to ensure that the databases are fully identical.

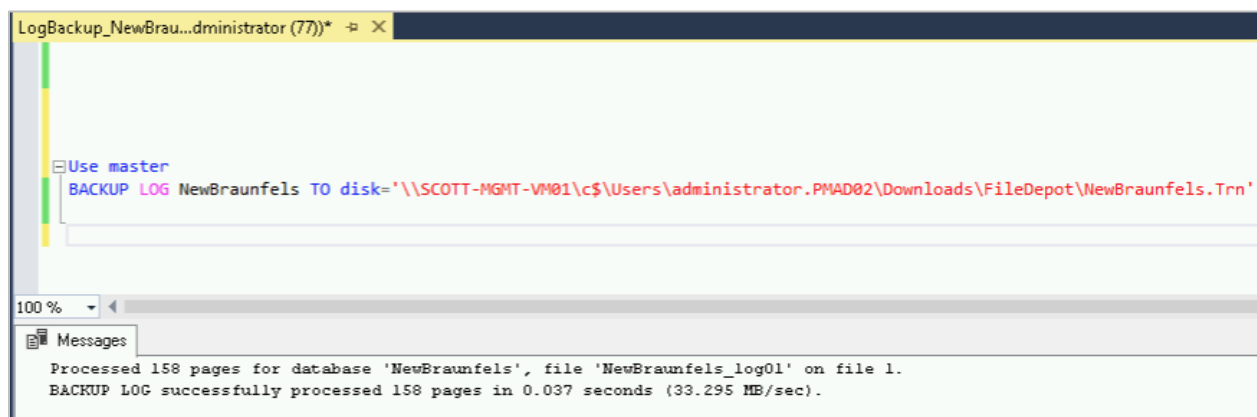
To capture the outstanding transactions:

1. [Take a SQL native log backup from the primary database.](#)
2. [Restore the log backup to the secondary database.](#)

Take Log Backup from Primary Database

Go back to the primary database and use SSMS to take a log backup of that database. This captures the changes and internal differences between the primary and the restored database.

For example:



The T-SQL command used in this example is:

```
BACKUP LOG NewBraunfels TO disk= '\\SCOTT-MGMT-VM01\FileDepot\NewBraunfels.Trn' WITH COPY_ONLY
```

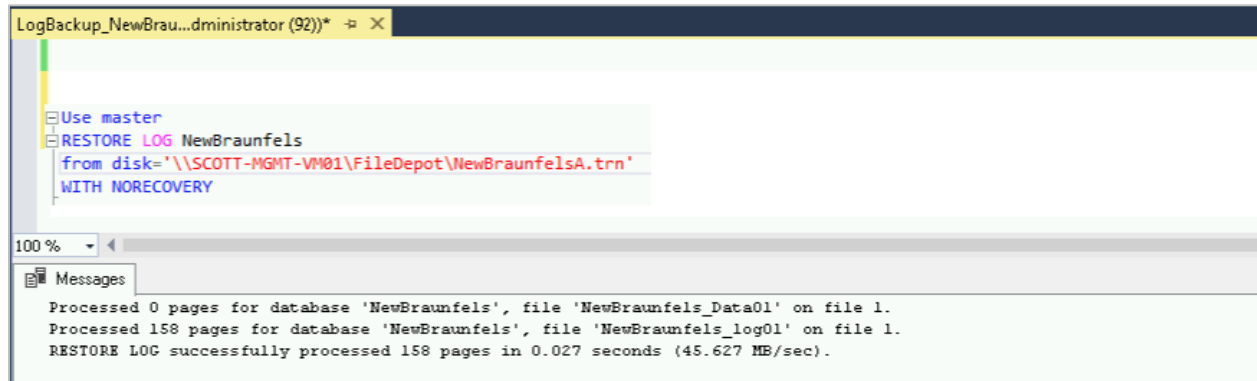
IMPORTANT: You must specify “COPY_ONLY” on the log backup. If you do not, you will break the log chain.⁵

⁵ For a deeper discussion on the COPY_ONLY option, see Microsoft’s [Copy-Only Backups, Microsoft](#).

Restore Log Backup to Secondary Database

Now use SSMS to restore the SQL native log to the secondary database. Because the database is in the “Restoring” state, it will accept the changes captured in the transaction log backup. Apply the log file and leave the database in a restoring state.

For example:



```
LogBackup_NewBraun...dministrator (92)* -> X
```

```
Use master
RESTORE LOG NewBraunfels
from disk='\\SCOTT-MGMT-VM01\FileDepot\NewBraunfelsA.trn'
WITH NORECOVERY
```

100 %

Messages

```
Processed 0 pages for database 'NewBraunfels', file 'NewBraunfels_Data01' on file 1.
Processed 158 pages for database 'NewBraunfels', file 'NewBraunfels_log01' on file 1.
RESTORE LOG successfully processed 158 pages in 0.027 seconds (45.627 MB/sec).
```

The T-SQL command used in this example is:

```
RESTORE LOG NewBraunfels from disk=
'\\SCOTT-MGMT-VM01\FileDepot\NewBraunfels.trn'
WITH NORECOVERY
```

Both the primary database and the restored database are now transactionally identical, and the database is now ready to be introduced into the SQL AG relationship.

Tie the Knot with SQL AG

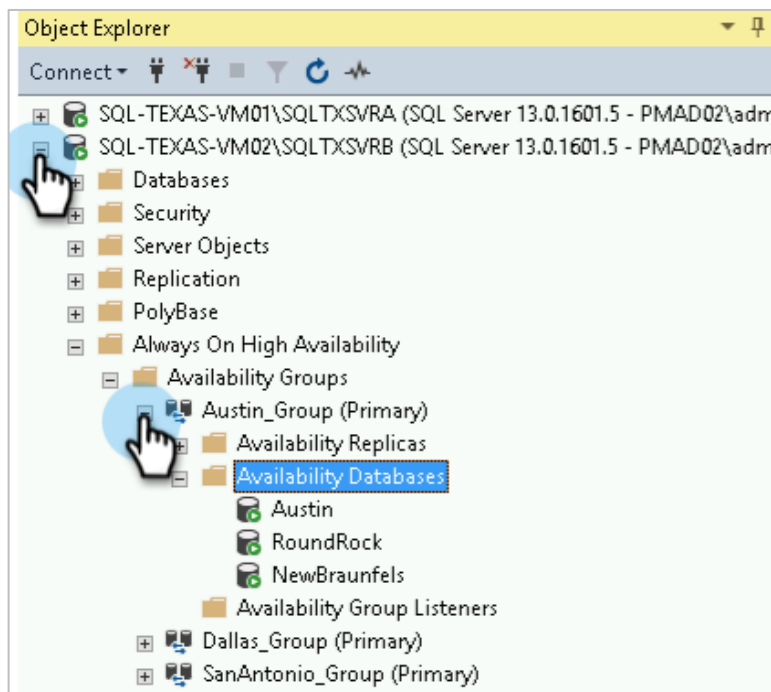
Now that our two databases are identical⁶, we can introduce them into the SQL AG group.

Add Database to Existing AG Group on Primary AG Host

Before we can join the database to AG on the secondary host, we must add it to an AG group on the primary host.

To add a database to an existing AG group on the primary AG host:

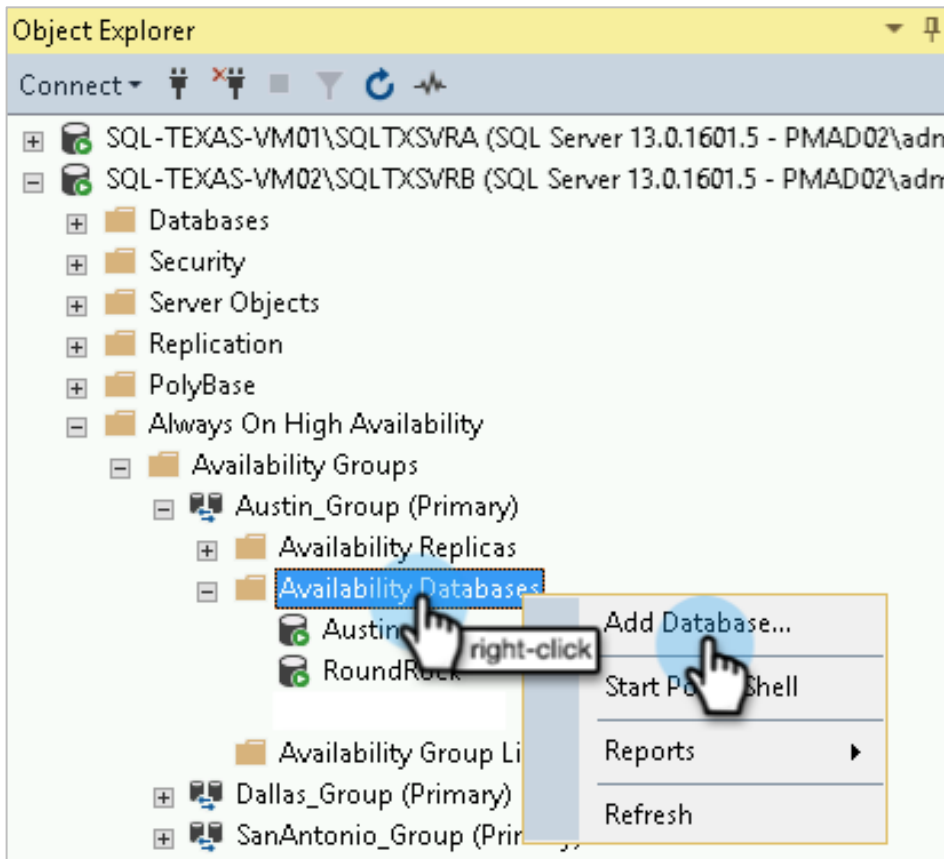
1. Open SSMS and, under the primary AG host (TEXAS 02 in our example), navigate to **Always On High Availability > Availability Groups > <AG_Name> (Primary) > Availability Databases**. In our example, we'll use the "Austin_Group" AG.



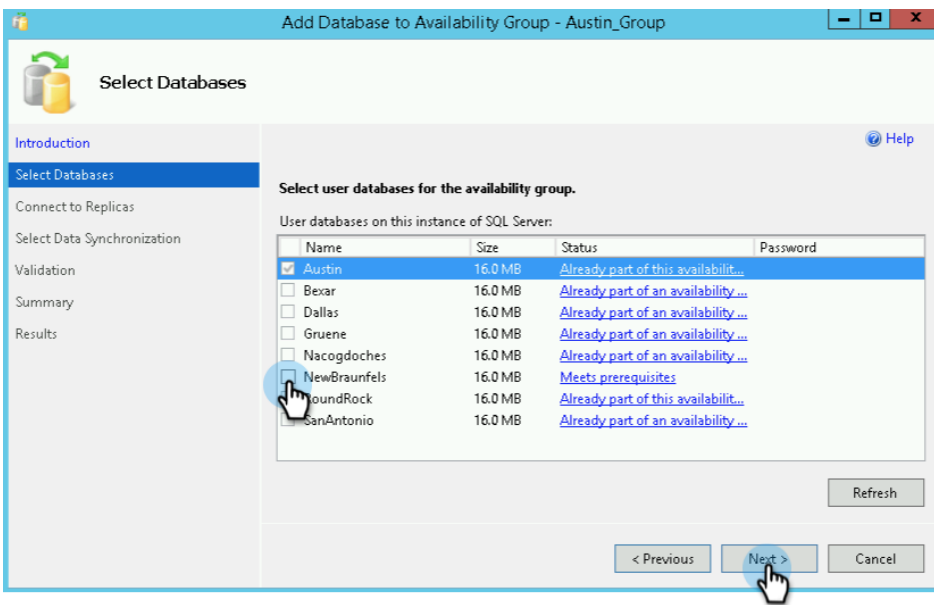
IMPORTANT: You must start the introduction with the AG group labeled **Primary**. In the example above, the primary group resides on SQL instance TEXAS 02. Later, we will introduce the restored copy into the AG group on the other end, where it sits on TEXAS 03.

⁶ The databases are not strictly identical in every sense of the word. However, they are close enough internally to be connected by an AG group. There is a brief moment during the introduction to AG where outstanding transactions on the primary replica are synchronized with the secondary replica, ultimately making them identical.

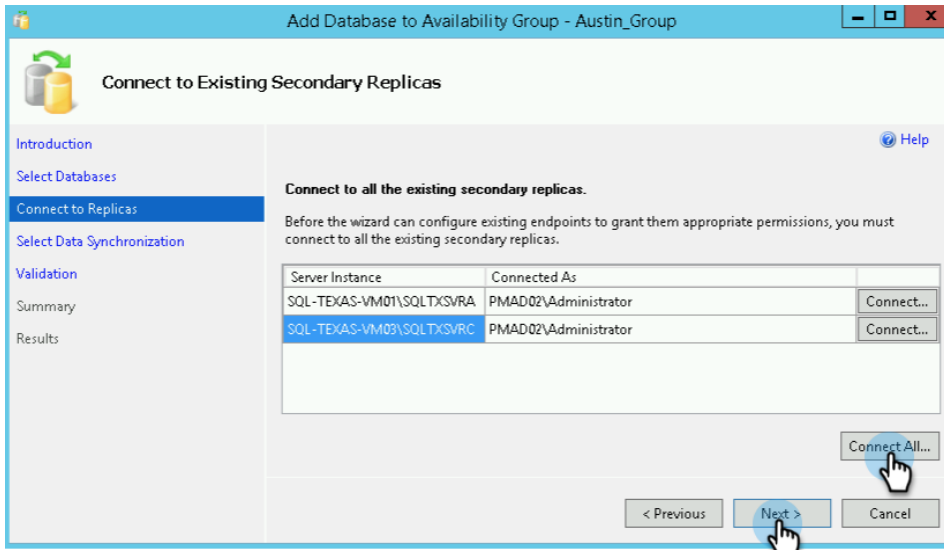
2. Right-click **Availability Databases** and select **Add Database....**



3. Select the database and click **Next**. (Note that in our example, the NewBraunfels database is listed as **Meets prerequisites**.)



- Connect to the replica hosts. In this existing SQL AG relationship, we have two replica hosts. Click **Connect All** to connect to all secondary SQL AG instances, then click **Next**.



- See Table 2 below to choose the right data synchronization option for you. Make your selection and click **Next**.

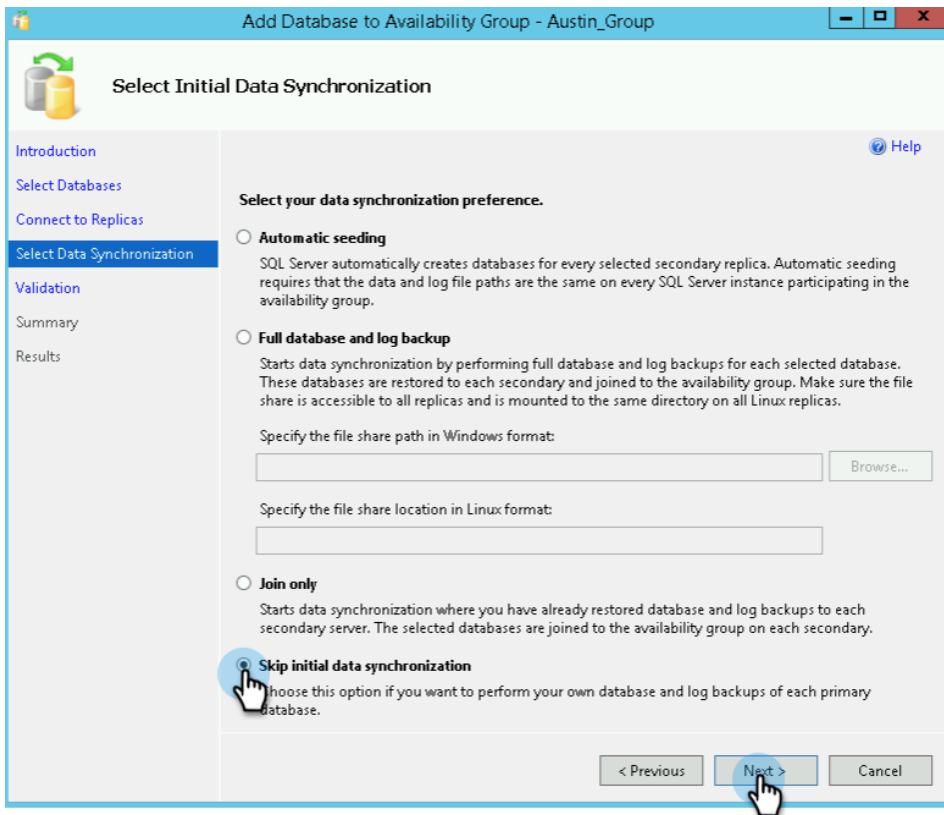
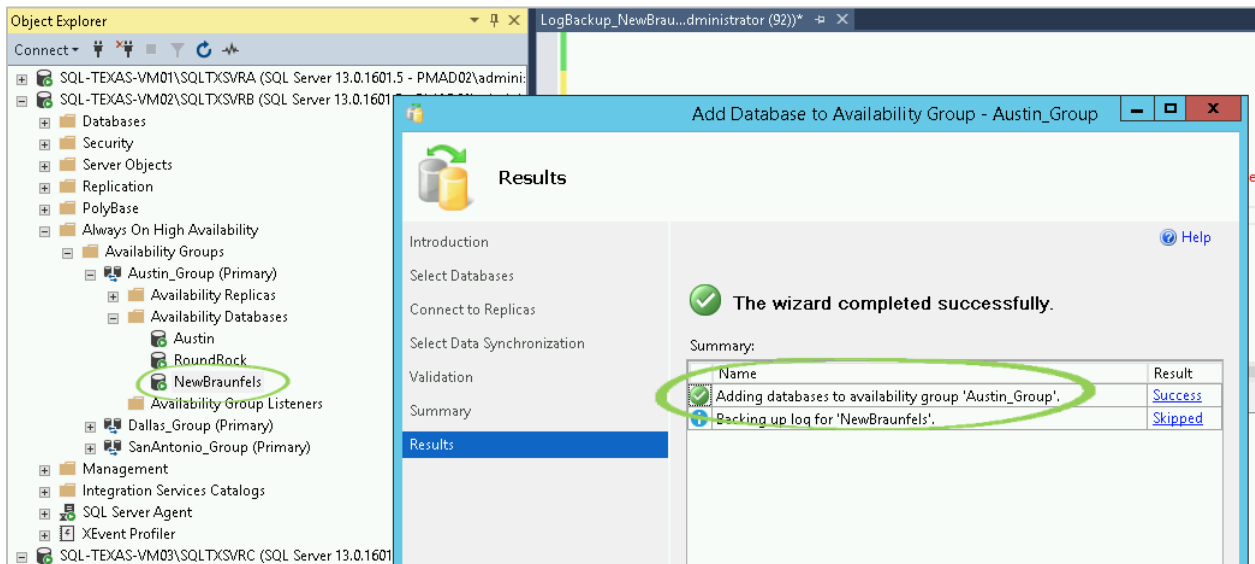


Table 2: Data Synchronization Options

OPTION	DEFINITION
Automatic seeding	SQL Server automatically creates the secondary replicas for every database in the group.
Full database and log backup	For each primary database, this option performs several operations in one workflow: create a full and log backup of the primary database, create the corresponding secondary databases by restoring these backups on every server instance that is hosting a secondary replica, and join each secondary database to availability group.
Join only	Select this option only if the new secondary databases already exist on each server instance that hosts a secondary replica for the availability group. The wizard will attempt to join each existing secondary database to the availability group.
Skip initial data synchronization	Select this option if you want to perform your own database and log backups of every primary database, restore them to every server instance that hosts a secondary replica. After you exit the wizard, you will then need to join every secondary database on every secondary replica.

For more, see Microsoft's [Select Initial Data Synchronization Page \(Always On Availability Group Wizards\)](#).

In our example, when the wizard completes successfully, it will show that the NewBraunfels database has been added to the primary side of the the “Austin_Group” AG.

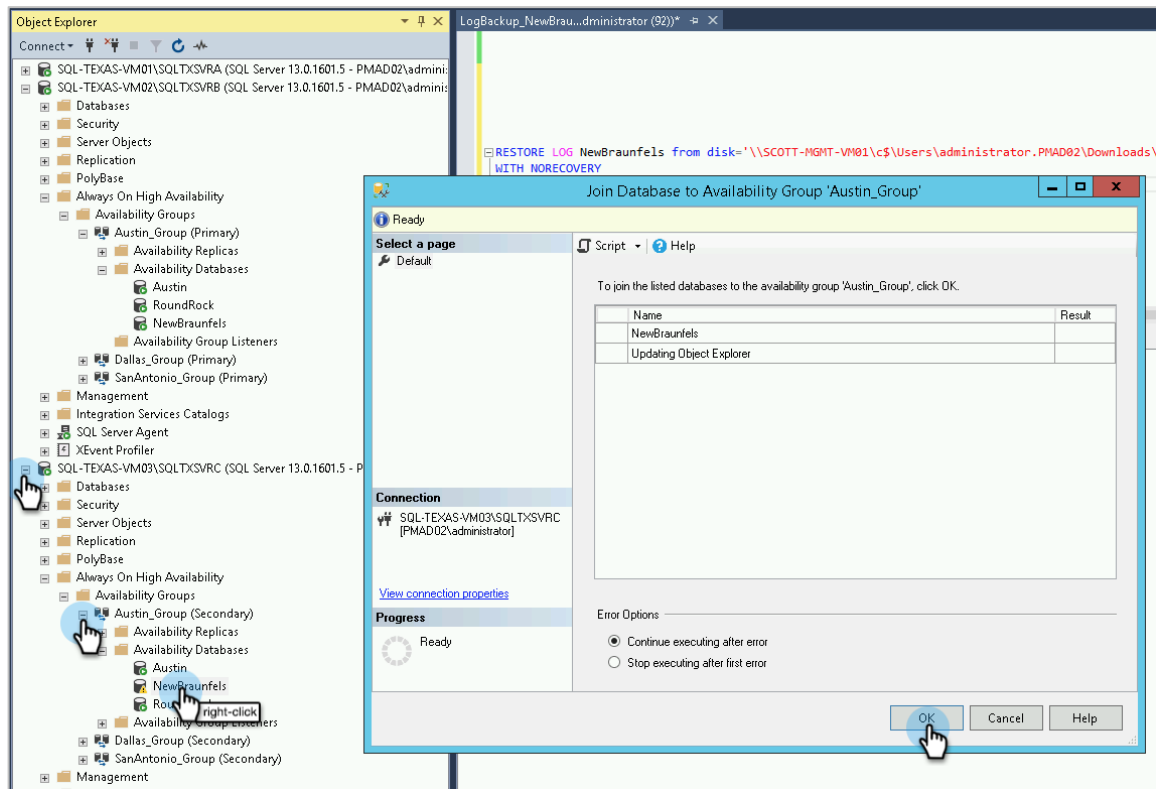


Join Database to Existing AG Group on Secondary AG Host

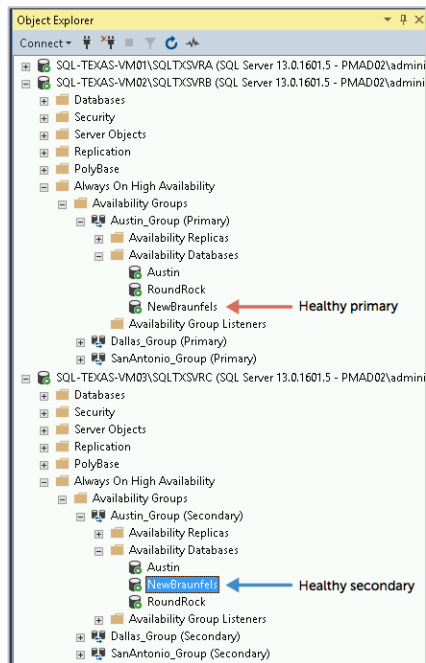
After adding the database to an AG group on the primary host, we are now ready to join it to an AG group on the secondary host.

To join a database to an existing AG group on the secondary AG host:

1. Open SSMS and, under the secondary AG host (TEXAS 03 in our example), navigate to **Always On High Availability > Availability Groups > <AG_Name> (Primary) > Availability Databases**. In our example, we're using the "Austin_Group" AG. Right-click the database name (**NewBraunfels**) and when prompted, click **OK** to join the database.



2. Confirm that you have successfully created the AG database in SSMS.



IMPORTANT: You should go back and enable the Protection Job that protects these servers. Check the Protection Job and confirm that the NewBraunfels database is protected.

Your Feedback

Was this document helpful? [Send us your feedback!](#)

About the Authors

Scott Lorenz is a SQL Solutions Engineer at Cohesity. In his role, Scott focuses on business-critical applications, MS SQL Server databases, and data protection with Enterprise and Cloud Storage.

Other essential contributors include:

- Bart Abicht, Senior Technology Writer and Editor at Cohesity
- Freddy Grahn, Senior Technical Field Enablement Engineer

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.2	July 2024	Republishing
1.1	Feb 2022	Rebranding updates
1.0	Dec 2019	First release

ABOUT COHESITY

[Cohesity](#) is a leader in AI-powered data security and management. Aided by an extensive ecosystem of partners, Cohesity makes it easier to protect, manage, and get value from data – across the data center, edge, and cloud. Cohesity helps organizations defend against cybersecurity threats with comprehensive data security and management capabilities, including immutable backup snapshots, AI-based threat detection, monitoring for malicious behavior, and rapid recovery at scale. Cohesity solutions are delivered as a service, self-managed, or provided by a Cohesity-powered partner. Cohesity is headquartered in San Jose, CA, and is trusted by the world's largest enterprises, including six of the Fortune 10 and 44 of the Fortune 100.

Visit our [website](#) and [blog](#), follow us on [Twitter](#) and [LinkedIn](#) and like us on [Facebook](#).

© 2024. Cohesity, Inc. All Rights Reserved. The information supplied herein is the confidential and proprietary information of Cohesity and may only be used (a) by the intended recipients and (b) in conjunction with validly licensed Cohesity software and services. Find the terms of Cohesity licenses at www.cohesity.com/agreements.

Cohesity, the Cohesity logo, SnapTree, SpanFS, DataPlatform, DataProtect, Helios, the Helios logo, DataGovern, SiteContinuity, DataHawk, and other Cohesity marks are trademarks or registered trademarks of Cohesity, Inc. in the US and/or internationally. Other company and product names may be trademarks of the respective companies with which they are associated. This material (a) is intended to provide you information about Cohesity and our business and products; (b) was believed to be true and accurate at the time it was written, but is subject to change without notice; and (c) is provided on an "AS IS" basis. Cohesity disclaims all express or implied conditions, representations, warranties of any kind.