



Technical Report

SQL Server on Google Cloud Platform Using Cloud Volumes Service

Deployment guide

August 2021 | TR-4898
Pat Sinthusan, NetApp

Akaash Rampersad
Customer Engineer, Infrastructure Modernization
Google Cloud

In partnership with

Google Cloud

Abstract

This document provides guidance on how to deploy Microsoft SQL Server leveraging NetApp® Cloud Volumes Service on virtual machines (VMs) in Google Cloud Platform.

TABLE OF CONTENTS

Intended audience	3
Microsoft SQL Server on Google Cloud Platform overview	3
Before you begin.....	3
Prerequisites.....	4
SMB and continuous availability	4
How to enable continuous availability on Cloud Volumes Service	4
SQL Server over SMB requirements	5
Create Cloud Volumes Service for SQL Server	7
Create a Cloud Volumes Service volume	7
Assign permissions and create a database over SMB on Cloud Volumes Service	10
Deploy Always On failover cluster over SMB on Cloud Volumes Service	13
Create Windows Server Failover Cluster.....	13
Install a new SQL Server failover cluster	17
Commonly known errors.....	24
Where to find additional information	25
Version history.....	25

LIST OF FIGURES

Figure 1) Enable CA for all volumes.....	5
Figure 2) Join Active Directory and provide domain account which elevated privilege.....	6
Figure 3) Specify CVS path for system database files to reside.....	7

Intended audience

This technical report provides an overview of how to configure a highly resilient Google Cloud Platform environment for [SQL Server installed on SMB fileshare storage](#) using [NetApp Cloud Volumes Service \(CVS\) for Google Cloud](#).

This document is intended for NetApp Cloud Volumes Service and/or SQL Server database administrators who are responsible for deploying Microsoft SQL Server in Google Cloud Platform. It is assumed that the reader is familiar with the various components of the solution.

Microsoft SQL Server on Google Cloud Platform overview

Today, many customers migrate to Google Cloud Platform to accelerate their SQL Server deployments, reduce costs, and provide increased agility for their business processes. These key benefits enable IT decision makers to adopt a cloud-first strategy for a critical infrastructure such as SQL Server.

Moreover, moving the SQL Server estate to Google Cloud and running it on Google Compute Engine (GCE) provides full compatibility with existing tools and workflows such as SQL Server Management Studio (SSMS), SQL Server Integration Services (SSIS), SQL Server Reporting Services (SSRS), and Visual Studio.

In preparation for migrating production SQL Servers into the cloud, customers often migrate their development, test, staging, and other nonproduction SQL Servers first to test and certify their systems. GCE offers instances of various sizes along with multiple deployment options to accommodate these various systems. Deployments can scale from small, single host configurations to complex multizone/region configurations.

Memory-optimized GCE instances such as the M1-family and M2-family offer excellent compute capabilities for demanding database workloads such as SQL Server volume location databases (VLDBs) or SAP HANA while other standard workload databases run equally well on general-purpose GCE instances such as the N1-family, N2-family, and E2-family. Block storage for these machine families is provided by Google Cloud Persistent Disks, which are networked together and provide IOPs based on the size of the disk along with the number of vCPUs on the instance. NetApp Cloud Volumes can help overcome these conditions by providing IOPs at different performance levels regardless of the size of the volume.

Traditional SQL Server Always On instances usually require each node to maintain a copy of the data, which can quickly become costly as the databases grow. For example, in a traditional SQL Server Always On deployment, if the databases are 10TB in size, each replica must have at least 10TB of storage provisioned (primary replica, 10TB; secondary replica, 10TB; and so on). One of the benefits of having a Windows Server Failover Cluster (WSFC) deployed for SQL Server is only needing to have one copy of the data, which is shared between the nodes of the cluster. Therefore, running SQL Server Always On in a failover cluster instance (FCI) on a WSFC can immediately realize cost savings by reducing the amount of storage required.

NetApp Cloud Volumes Service and Cloud Volumes ONTAP are two cloud storage solutions for running high performance SQL Server workloads in combination with GCE instances.

Before you begin

Before you complete the steps in this report, you should already have:

- A Google Cloud Platform subscription
- A Windows Active Directory domain

- A domain user account to run the SQL Server service and login capability into the GCE instance to mount file shares
- A domain user account for installing SQL Server
- DNS configured on the Google Cloud Platform network, pointing to the domain controller

Prerequisites

You should have a domain user account that has permissions to create objects on both GCE instances and in Active Directory.

You should also have an account to create Cloud Volumes Service and a basic understanding of how to install SQL Server single instance.

Before you create a volume in Cloud Volumes Service, you should have a solid understanding of [Cloud Volumes Service service levels](#) and [CVS-Performance service type](#). For information about service-level performance of CVS-Performance volumes, see the following resources:

- [Performance expectations](#)
- [Selecting the appropriate service level and allocated capacity for NetApp Cloud Volumes Service](#)

SMB and continuous availability

SMB3 with the continuously available share property enabled provides a very high level of resiliency between the GCE instances and the storage service. SMB Transparent Failover enables maintenance operations on the Cloud Volumes Service without interrupting connectivity to server applications storing and accessing data on SMB volumes. To support SMB Transparent Failover, Cloud Volumes Service supports the SMB continuous availability (CA) shares option for use with SQL Server applications over SMB running on GCE instances. SMB CA shares enable SQL Server workloads on Cloud Volumes Service, which provides performance improvements, scale, and cost benefits for single instance, Always On failover cluster (AOFC) instance, and Always On availability group deployments.

How to enable continuous availability on Cloud Volumes Service

NetApp recommends that you enable CA for all the volumes that host the SQL Server data files. This option can be set during volume creation, as shown in Figure 1.

Figure 1) Enable CA for all volumes.

The screenshot shows the Google Cloud Platform interface for creating a file system. The left sidebar contains a menu with 'Volumes', 'Backups', 'Snapshots', 'Active Directories', and 'Volume Replication'. The main area is titled 'Create File System' and includes a section for 'Volume Details'. In this section, 'Allocated Capacity' is set to 1024 GiB, and 'Protocol Type' is set to SMB. Below these, there are three checkboxes: 'Make snapshot directory (.snapshot) visible', 'Enable SMB Encryption', and 'Enable CA share support for SQL server'. The 'Enable CA share support for SQL server' checkbox is checked and highlighted with a red rectangular box. The text below this checkbox reads: 'Enable this option only for sql server workloads that require continuous availability.'

SQL Server over SMB requirements

Starting with SQL Server 2012 (11.x), system databases (master, model, MSDB, and TempDB) and database engine user databases [can be installed with Server Message Block \(SMB\) file server as a storage option](#). This applies to both SQL Server stand-alone and SQL Server FCIs.

These allow you to leverage Cloud Volumes Service with all its performance and data management capabilities, such as volume capacity, performance scalability, and data protection features, which the SQL Server can take advantage of.

To install SQL Server over SMB shares, the installer must meet the following requirements:

- The installer has been granted the SeSecurityPrivilege setting through Google Cloud Platform Portal.
- The installer has read/write access to the share.
- The installer is a member of the local Administrators Group on the Windows host.

Figure 2) Join Active Directory and provide domain account which elevated privilege.

Google Cloud Platform Cloud Heroes Search

Cloud Volumes

- Volumes
- Backups
- Snapshots
- Active Directories
- Volume Replication

← Edit Active Directory connection

NetBIOS name of the server that will be created.

Organizational Unit
OU=CVS,OU=GCP,OU=NetApp

Name of the Organizational Unit(OU) within Windows Active Directory the user belongs to in order from leaf OU to root OU.

Kerberos Realm Details

Provide the AD server name and KDC IP address needed to create the service principal name(SPM) machin account used by Cloud Volumes Service.

AD Server Name
Active Directory Server Name used for the Kerberos realm.

KDC IP
Key Distribution Center IP address used for the Kerberos realm.

Region

Cloud Volume Service supports only one AD connection per Google cloud region. You can only associate a cloud volume in a region with the AD connection in the same region.

Region
us-central1

The region to which the Active Directory credentials are associated. Only one connection can be configured for a single region.

Security Privilege Users

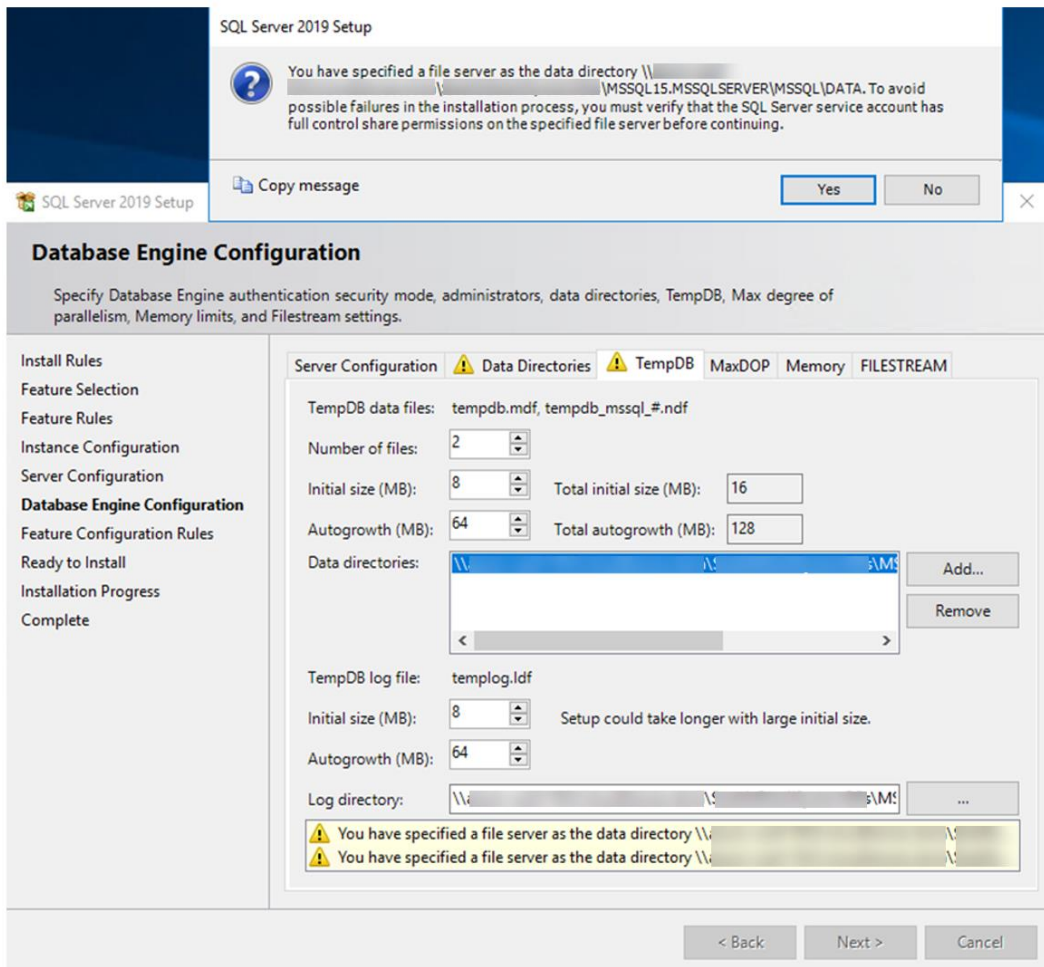
Provide a list of comma seperated domain user accounts that require elevated privileges to manage security log for the Active Directory associated with Cloud Volumes Service.

Accountname
SQLInstaller

SAVE CANCEL

The installation process is the same as for the typical block storage. The only exception is that the data root directory for system database files can be pointed to an SMB share during the database engine configuration step.

Figure 3) Specify CVS path for system database files to reside.




Create Cloud Volumes Service for SQL Server

Create a Cloud Volumes Service volume

To create a Cloud Volumes Service volume, complete the following steps:

1. From the Google Cloud Platform portal, select Cloud Volumes Service or enter Cloud Volumes Service in the search box.



2. To [create an SMB volume](#), select Volumes and Create.

Google Cloud Platform

Cloud Heroes

Cloud Volumes

Volumes

CREATE

DELETE

Volumes

Backups

Snapshots

Active Directories

Volume Replication

Quick reference for Cloud Volumes
[Private Service Access](#)

Filter
Search for volumes by name, ID, region, etc.

<input type="checkbox"/>		ID	Name
<input type="checkbox"/>		0e23bba9-4c8c-2bc7-9d5f-6009b540e77d	prabueurope
<input type="checkbox"/>		d32d904f-aeb3-059a-35f8-b128270c5a8b	prabueurope
<input type="checkbox"/>		ce227f08-3961-dfe7-d203-f775h935669h	prabutesteu

- Provide the volume name, service type, region, service level, allocated capacity, and protocol type. Enable CA share support for SQL Server, then click Save.

Google Cloud Platform

Cloud Heroes

Cloud Volumes

Create File System

Volumes

Backups

Snapshots

Active Directories

Volume Replication

Service Type

Cloud Volumes Service is offered as two service types: CVS and CVS-Performance. Select the service type that matches your workload needs. [Region availability](#) varies by service type. [Learn more](#)

CVS

Offers volumes created with zonal high availability.

CVS-Performance

Offers 3 performance levels and improved latency to address higher performance application requirements.

Volume Replication (BETA)

Secondary

Select to create volume as a destination target for volume replication. Applicable only to CVS-performance volumes.

Region

Region availability varies by service type.

Region *

us-central1

Volume will be provisioned in the region you select.

Volume Path *

ecstatic-brave-easley

Must be unique to the project.

Service Level

Select the performance level required for your workload.

Standard

Up to 16 MiB/s per TiB

Premium

Up to 64 MiB/s per TiB

Extreme

Up to 128 MiB/s per TiB

Snapshot

The snapshot to create the volume from.

Volume Details

Allocated Capacity *

2048

GiB

Allocated size must be between 1 TiB (1024 GiB) and 100 TiB (102400 GiB)

Protocol Type *

SMB

Make snapshot directory (.snapshot) visible

Enable this option to make .snapshot directory visible when you view your volume directories

Enable SMB Encryption

Enable this option only if you require encryption of your SMB data traffic.

Enable CA share support for SQL server

Enable this option only for sql server workloads that require continuous availability.

9

SQL Server on Google Cloud Platform Using
Cloud Volumes Service Deployment Guide

© 2021 NetApp, Inc. All Rights Reserved.

Assign permissions and create a database over SMB on Cloud Volumes Service

To assign permissions and create a database over SMB on Cloud Volumes Service, complete the following steps:

1. From the Google Cloud console, select Cloud Volumes and select the volume that was just created.

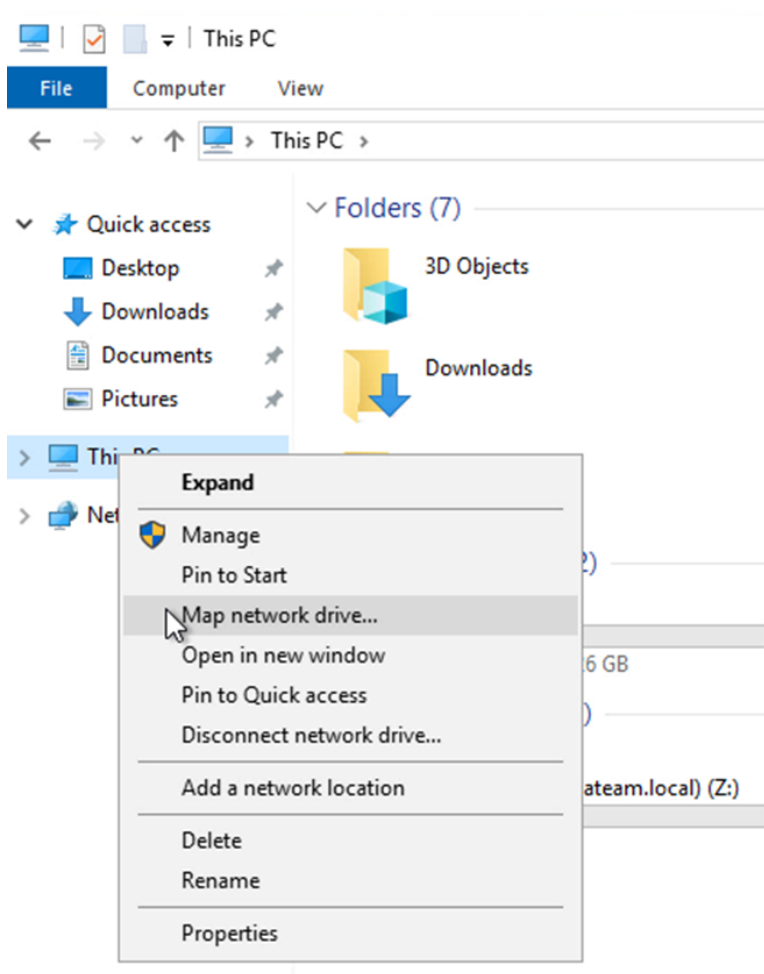
The screenshot shows the Google Cloud Platform console interface. The top navigation bar includes the Google Cloud Platform logo and the user 'Cloud Heroes'. The left sidebar shows the 'Cloud Volumes' section with sub-items: Volumes, Backups, Snapshots, Active Directories, and Volume Replication. The main content area displays the 'Volume Details' for a specific volume. The details include:

- Name: gcpsqlproddata01
- Service Type: CVS-Performance
- Storage Type: Primary
- ID: 603bc897-1d0c-f5cf-2615-91112fe30cb4
- Region: us-central1
- Life Cycle: available
- VPC: herocore-vpc
- Created: 2021-04-07T21:22:14.000Z
- Service Level: Premium
- Allocated Capacity: 1,024 GiB
- Protocol Type: SMB

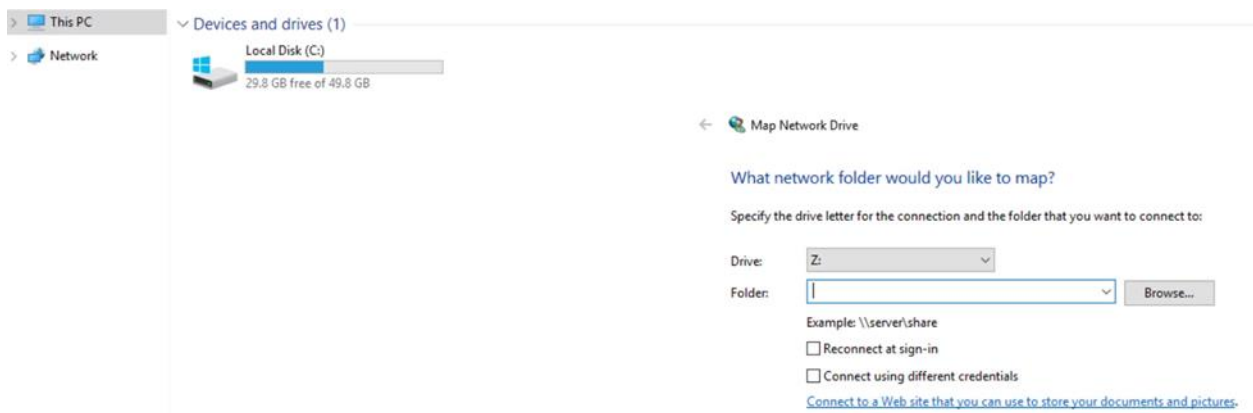
The 'Mount Targets' section is highlighted with a red box and contains the following table:

Protocol Type	Mount Target
SMB	\\NAS-9deb.cloudheroes.dom\\brave-wizardly-fermi

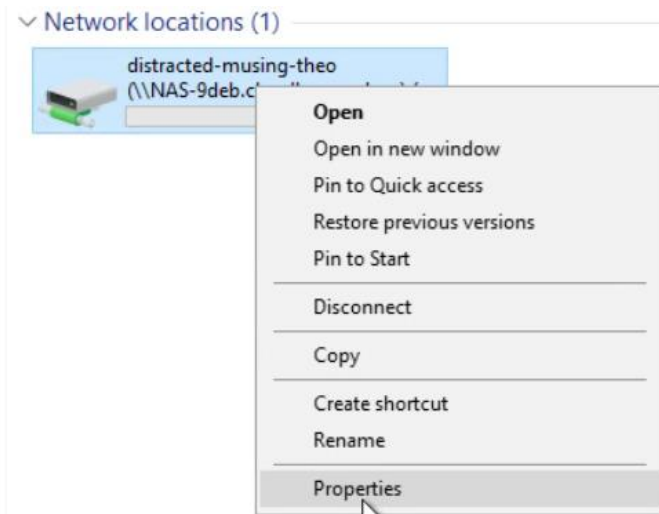
2. On the Google Cloud Platform VM, open Windows Explorer.
3. Right-click This PC.
4. Select Map Network Drive.



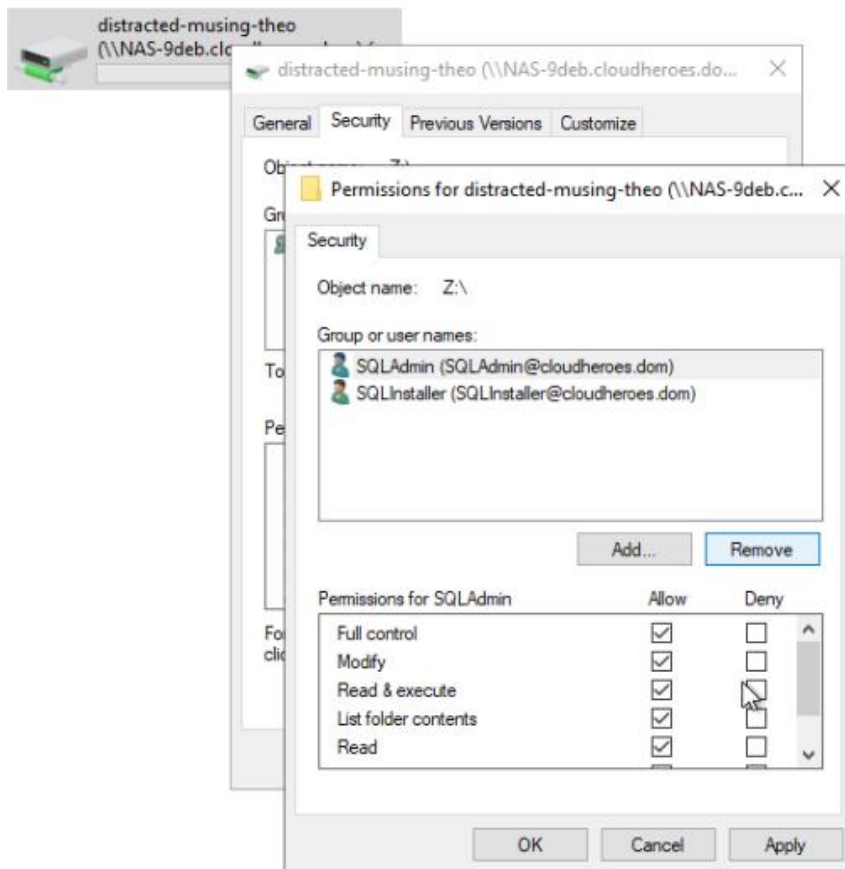
5. Paste the copied SMB path in the folder.




6. At sign-in, disable the Reconnect option.
7. Click Finish.
8. Right-click the mapped network drive and select Properties.



9. From the Security tab, click Edit.
10. Remove Everyone and add the SQL Server installer and SQL Server service accounts with full control permissions.



11. Create the database with data and log files residing on the Cloud Volumes Service SMB shares by using SQL Server Management Studio (SSMS).



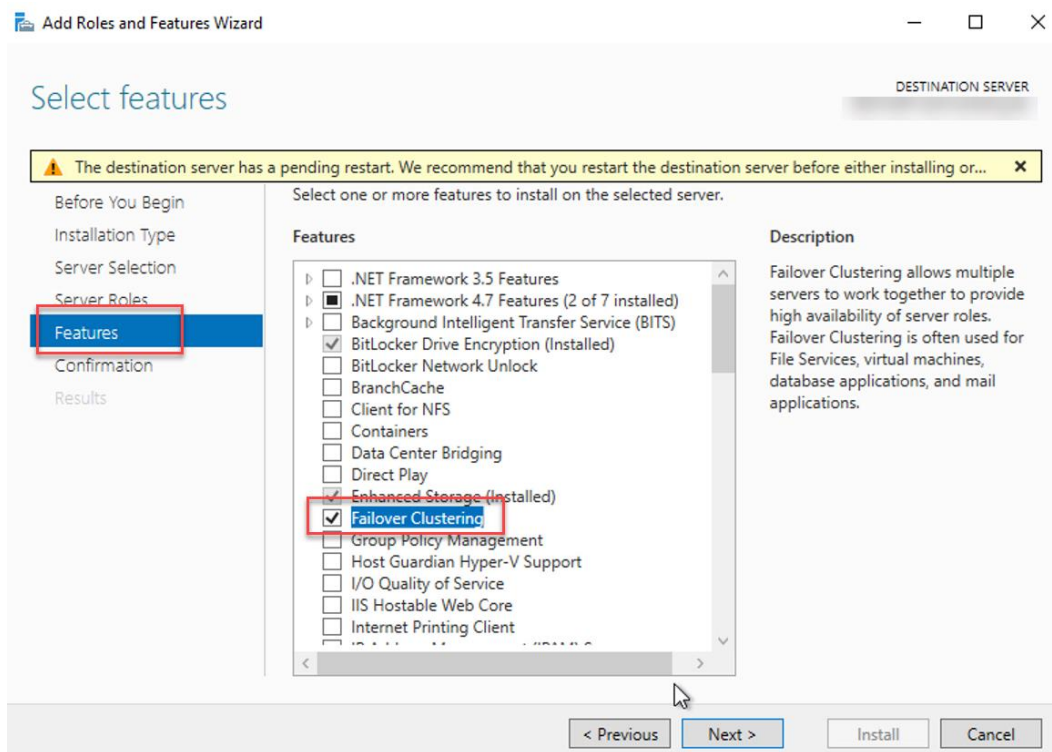
Deploy Always On failover cluster over SMB on Cloud Volumes Service

Google Cloud and NetApp recommend creating GCE instances in different zones. The following section describes how to deploy AOFC over SMB on Cloud Volumes Service.

Create Windows Server Failover Cluster

To create WSFC, complete the following steps:

1. Using Server Manager, add the Failover Clustering feature. Click Next.



PowerShell with administrator rights can also be used to add the failover clustering feature by running the following command:

```
Install-WindowsFeature -Name Failover-Clustering -IncludeManagementTools -Restart
```

2. Validate the configuration and create Windows failover cluster by adding participating servers.

If you are planning to use SQL Server 2016, do not use the user interface to create Windows failover cluster because this will cause dictionary key errors. Instead, use the following PowerShell command to create and validate Windows failover cluster.

```

$ClusterName = "SQLProdCluster"
$node1 = "SQLProd01"
$node2 = "SQLProd02"
$sqlvip = "10.1.1.68"
$quorum = '\\gcpsmb\quorum' #This is CVS path that create with GCP Portal

New-Cluster -Name $ClusterName -Node $node1,$node2 -NoStorage -StaticAddress $sqlvip -
managementpointnetworktype singleton
Set-ClusterQuorum -FileShareWitness $quorum -Credential $(Get-Credential)
Test-Cluster -Node $node1, $node2

```

Create Cluster Wizard



Select Servers

Before You Begin

Select Servers

Validation Warning

Access Point for Administering the Cluster

Confirmation

Creating New Cluster

Summary

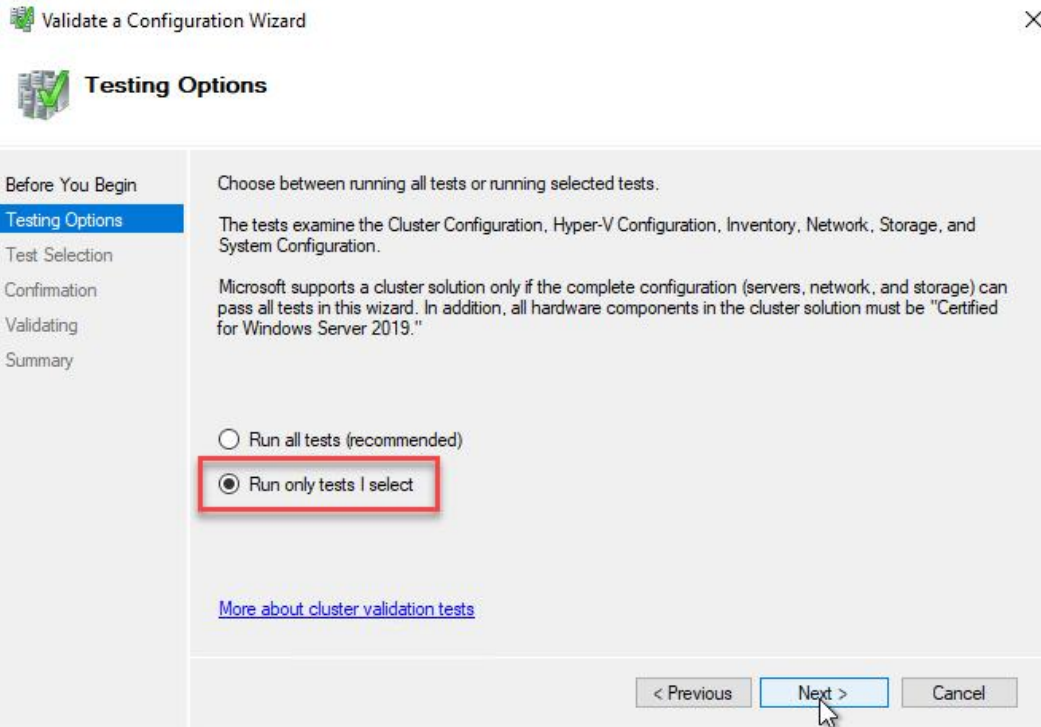
Add the names of all the servers that you want to have in the cluster. You must add at least one server.

Enter server name:

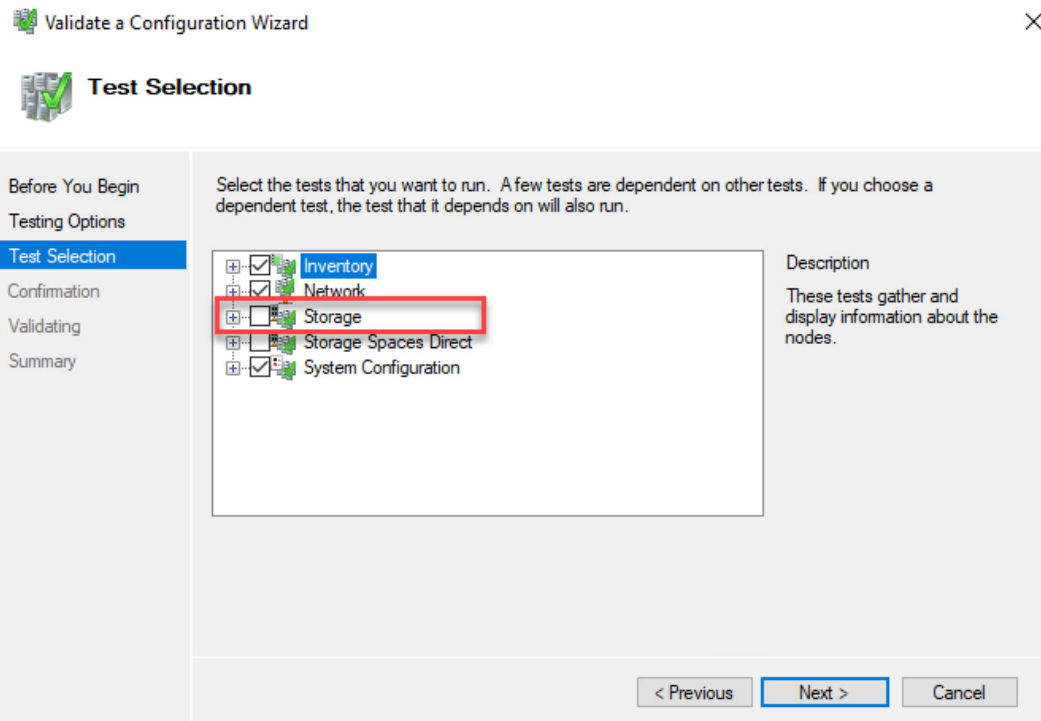
Selected servers:

SQLProd01: ☐
SQLProd02: ☐

3. In the Validate a Configuration Wizard, under Testing Options, select Run Only Tests I Select. Click Next.



4. Under Test Selection, clear the Storage option because local storage is not going to be used.




5. Under Summary, click Finish to complete validation configuration.



Summary

Before You Begin
Testing Options
Test Selection
Confirmation
Validating
Summary

 Testing has completed for the tests you selected. You should review the warnings in the Report. A cluster solution is supported by Microsoft only if you run all cluster validation tests, and all tests succeed (with or without warnings).

Node	
SQLProd01	Validated
SQLProd02	Validated

Result	
List BIOS Information	Success
List Environment Variables	Success
List Fibre Channel Host Bus Adapters	Success
List Host Guardian Service client configuration	Success
List iSCSI Host Bus Adapters	Success
List Memory Information	Success

To view the report created by the wizard, click View Report.
To close this wizard, click Finish.

[View Report...](#)

[Finish](#)

- In the Create Cluster Wizard, under Access Point for Administering a Cluster, select Create Cluster and provide a cluster name. Click Next.




Access Point for Administering the Cluster

Before You Begin
Select Servers
Access Point for Administering the Cluster
Confirmation
Creating New Cluster
Summary

Type the name you want to use when administering the cluster.

Cluster Name:

 The NetBIOS name is limited to 15 characters. One or more DHCP IPv4 addresses were configured automatically. All networks were configured automatically.

[< Previous](#) [Next >](#) [Cancel](#)

- Under Confirmation, clear the Add All Eligible Storage to the Cluster option and click Next.

**Confirmation**

Before You Begin

Select Servers

Access Point for Administering the Cluster

Confirmation

Creating New Cluster

Summary

You are ready to create a cluster.
The wizard will create your cluster with the following settings:

Cluster	ProdCluster
Node	SQLProd01, SQLProd02
Cluster registration	DNS and Active Directory Domain Services

☐ Add all eligible storage to the cluster.

To continue, click Next.

< Previous **Next >** Cancel

8. Under Summary, click Finish.

Summary

Before You Begin

Select Servers

Access Point for Administering the Cluster

Confirmation

Creating New Cluster

Summary

You have successfully completed the Create Cluster Wizard.

Node	SQLProd01, SQLProd02
Cluster	ProdCluster
IP Address	DHCP address on

To view the report created by the wizard, click View Report.
To close this wizard, click Finish.

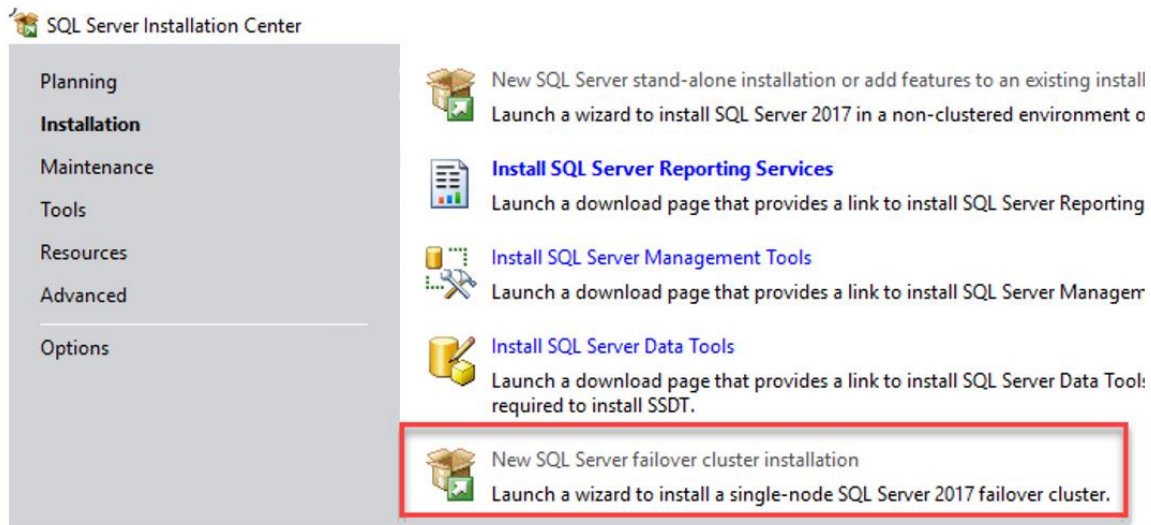
View Report...

Finish

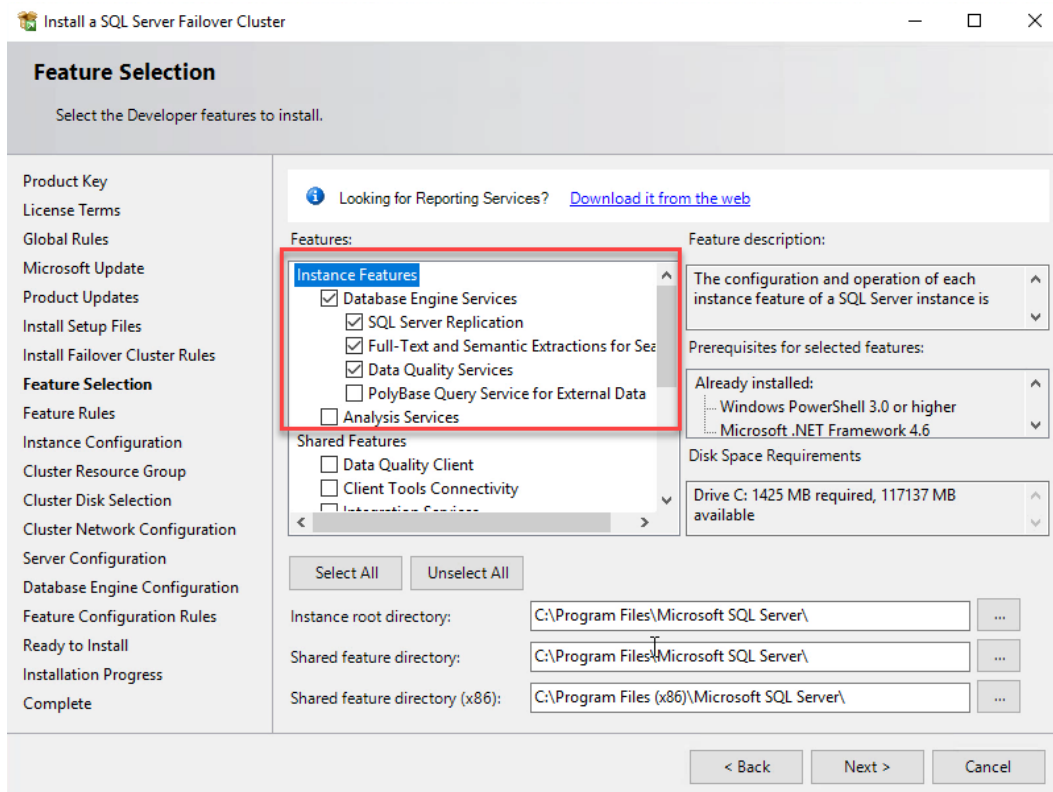
Install a new SQL Server failover cluster

To install a new SQL Server failover cluster, complete the following steps:

1. In SQL Server Installation Center, select Installation and New SQL Server Failover Cluster Installation.



2. Provide the product key and accept the license terms.
3. Under Feature Selection, select Database Engine Services.



4. Provide the SQL Server network name and click Next.

Instance Configuration

Specify the name and instance ID for the instance of SQL Server. Instance ID becomes part of the installation path.

Product Key
License Terms
Global Rules
Microsoft Update
Product Updates
Install Setup Files
Install Failover Cluster Rules
Feature Selection
Feature Rules
Instance Configuration
Cluster Resource Group
Cluster Disk Selection
Cluster Network Configuration
Server Configuration
Database Engine Configuration
Feature Configuration Rules
Ready to Install
Installation Progress
Complete

Specify a network name for the new SQL Server failover cluster. This will be the name used to identify your failover cluster on the network.

SQL Server Network Name:

☒ Default instance
☐ Named instance:

Instance ID:

SQL Server directory: C:\Program Files\Microsoft SQL Server\MSSQL14.MSSQLSERVER

Detected SQL Server instances and features on this computer:

Instance	Cluster Network Name	Features	Edition	Version	Inst
↔					

< Back
Next >
Cancel

5. In Database Engine Configuration, select the Data Directories tab and provide the data root directory with Cloud Volumes Service SMB volume paths.
6. In the information dialog box, click Next and then Yes.

Database Engine Configuration

Specify Database Engine authentication security mode, administrators, data directories and TempDB settings.

Product Key
License Terms
Global Rules
Microsoft Update
Product Updates
Install Setup Files
Install Failover Cluster Rules
Feature Selection
Feature Rules
Instance Configuration
Cluster Resource Group
Cluster Disk Selection
Cluster Network Configuration
Server Configuration
Database Engine Configuration
Feature Configuration Rules
Ready to Install
Installation Progress
Complete

Server Configuration **Data Directories** TempDB FILESTREAM

Data root directory:

System database directory:

Install a SQL Server Failover Cluster

You have specified a file server as the data directory \\NAS-9deb.cloudheroes.dom\distracted-musing-theo\MSSQL13.MSSQLSERVER\MSSQL\Data. To avoid possible failures in the installation process, you must verify that the SQL Server service account has full control share permissions on the specified file server before continuing.

Yes No

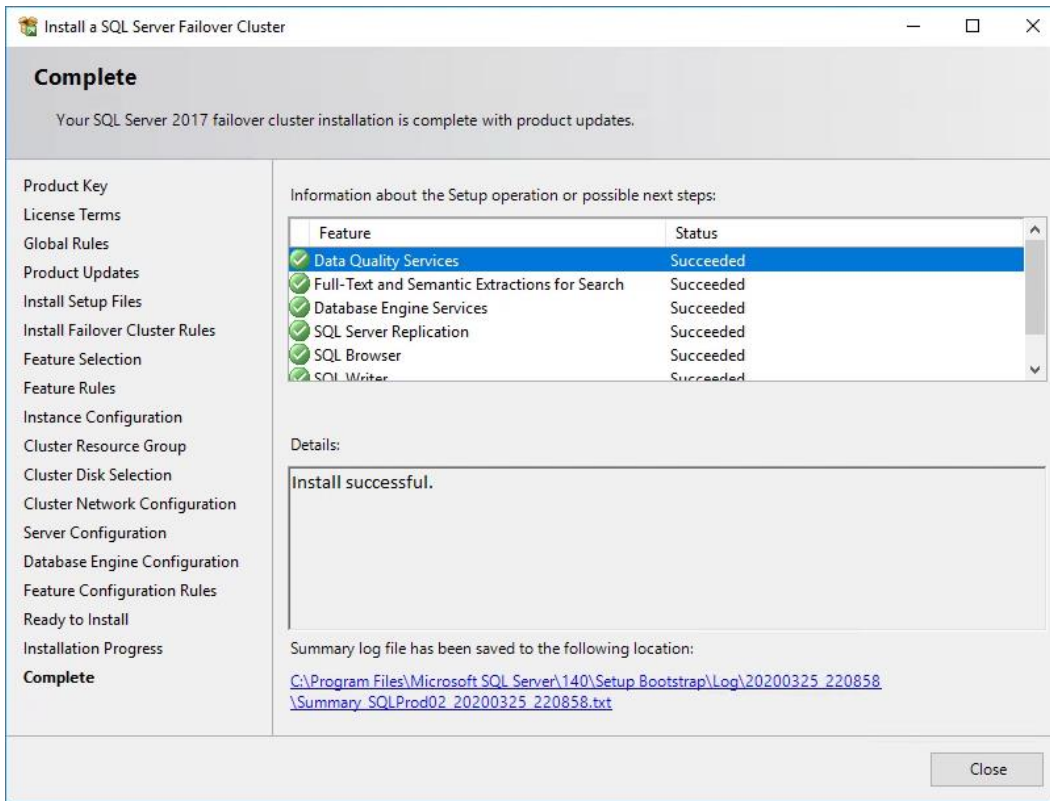
You have specified a file server as the data directory \\NAS-9deb.cloudheroes.dom\distracted-musing-theo\MSSQL13.MSSQLSERVER\MSSQL\Data.

You have specified a file server as the data directory \\NAS-9deb.cloudheroes.dom\distracted-musing-theo\MSSQL13.MSSQLSERVER\MSSQL\Data.

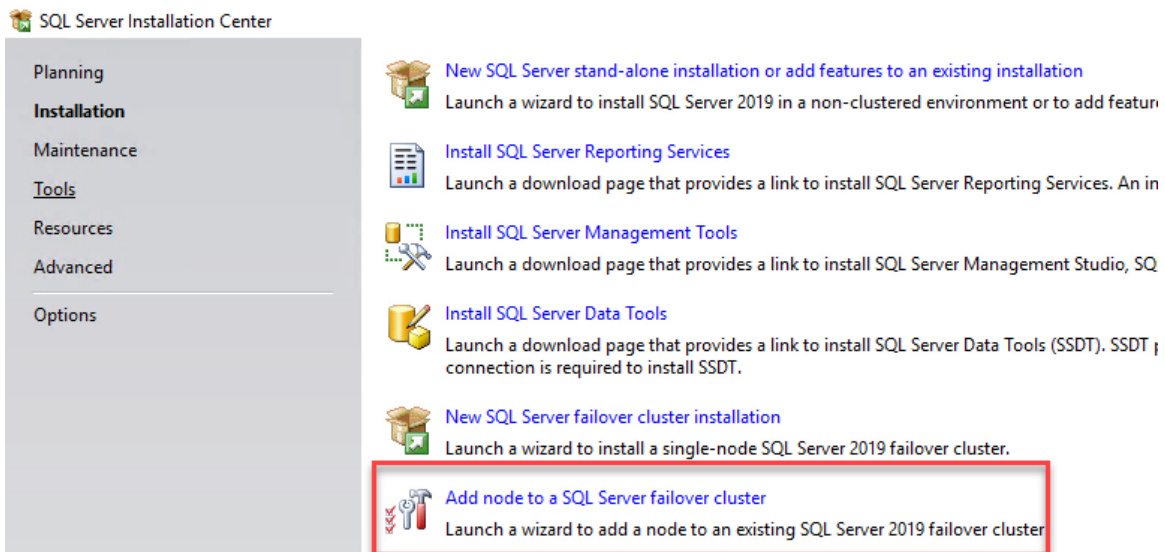
You have specified a file server as the data directory \\NAS-9deb.cloudheroes.dom\distracted-musing-theo\MSSQL13.MSSQLSERVER\MSSQL\Data.

< Back
Next >
Cancel

7. Finish the installation and click Close.



8. To install SQL Server in the second node, in the SQL Server Installation Center, select Installation then select Add Node to a SQL Server Failover Cluster .



9. Provide the product key and accept the license terms.
10. Provide the passwords for the SQL Server database engine and the SQL Server agent service accounts. Click Next.

Add a Failover Cluster Node

Service Accounts

Specify the service accounts and collation configuration.

Microsoft recommends that you use a separate account for each SQL Server service.

Service	Account Name	Password	Startup Type
SQL Full-text Filter Daemon Launcher	NT Service\MSSQLFDLaun...		Manual
SQL Server Database Engine		Manual
SQL Server Browser	NT AUTHORITY\LOCAL SE..		Automatic
SQL Server Agent		Manual

☒ Grant Perform Volume Maintenance Task privilege to SQL Server Database Engine Service

This privilege enables instant file initialization by avoiding zeroing of data pages. This may lead to information disclosure by allowing deleted content to be accessed.

[Click here for details](#)

< Back Next > Cancel

11. To finish the installation, click Install.

Add a Failover Cluster Node

Ready to Add Node

Verify the SQL Server 2019 features to be installed as part of the add node operation.

Ready to add this node to the SQL Server 2019 failover cluster:

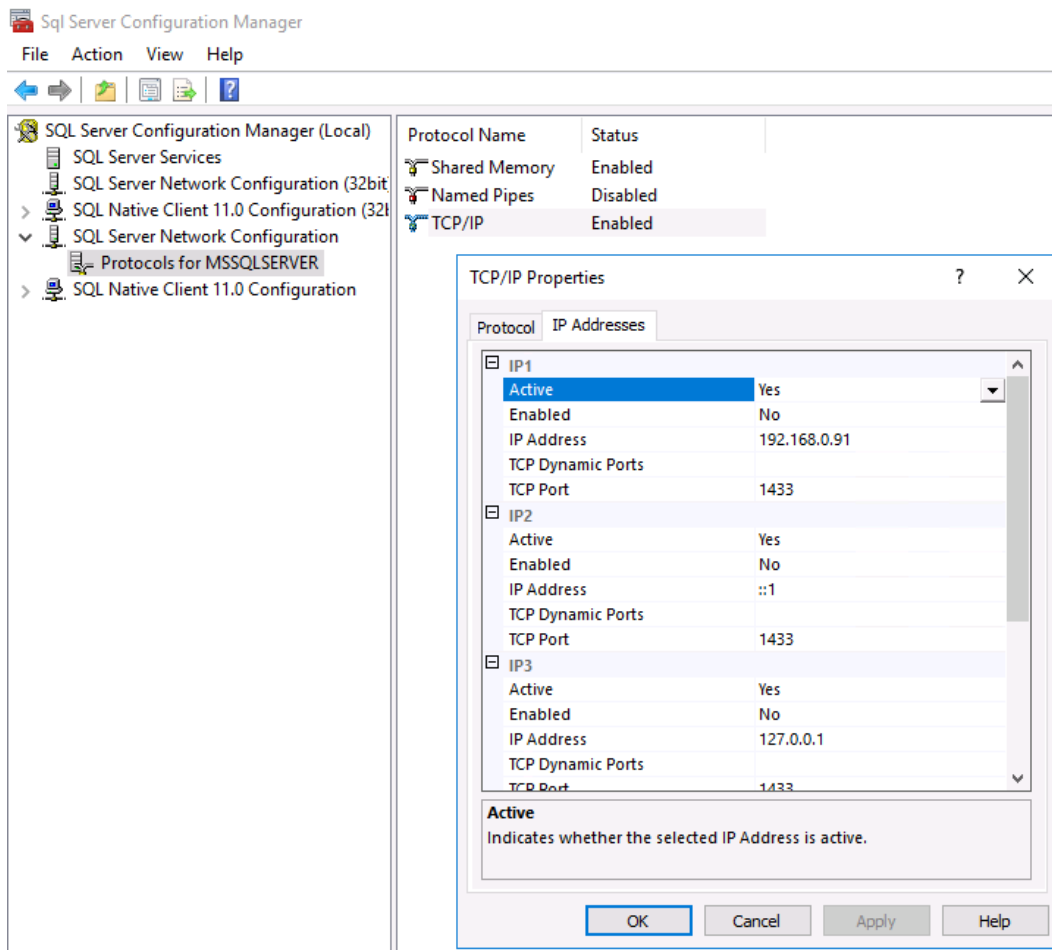
- Summary
 - Edition: Developer
 - Action: AddNode (Product Update)
 - Prerequisites
 - Already installed:
 - Windows PowerShell 3.0 or higher
 - To be installed from media:
 - Microsoft Visual C++ 2017 Redistributable
 - General Configuration
 - Features
 - Database Engine Services
 - SQL Server Replication
 - Full-Text and Semantic Extractions for Search
 - Data Quality Services
 - Instance configuration
 - Instance Name: MSSQLSERVER
 - Instance ID: MSSQLSERVER
 - Instance IDs
 - SQL D... F... MSSQL15-MSSQLSERVER

Configuration file path:

C:\Program Files\Microsoft SQL Server\150\Setup Bootstrap\Log\20200527_011007\ConfigurationFile.ini

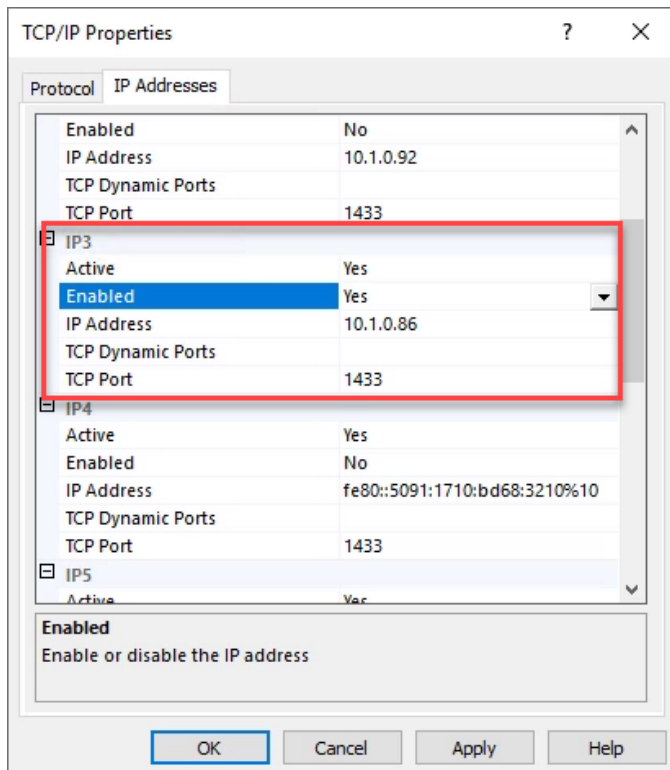
< Back Install Cancel

12. To enable virtual IP for SQL Server FCI, start SQL Server Configuration Manager.
13. Select SQL Server Network Configuration and then select Protocols for MSSQLSERVER.
14. Right-click TCP/IP and then select Properties.



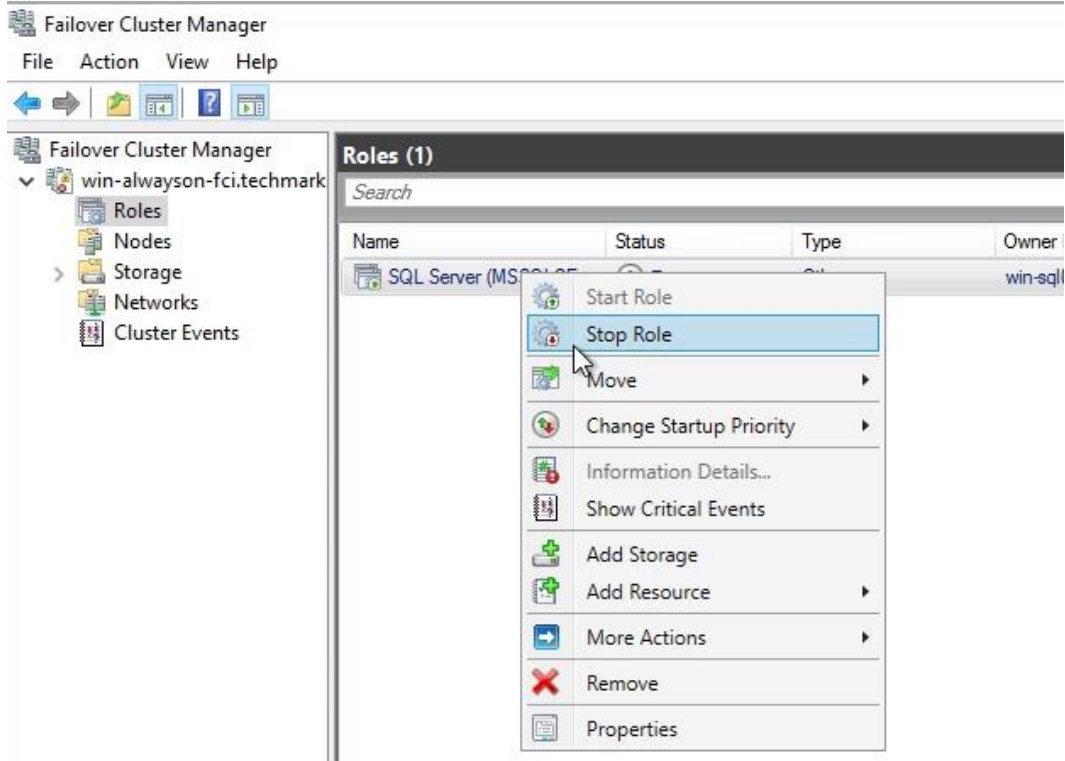
15. Select IP Addresses and enable IP3 for virtual IP Address and click OK.

Note: Make sure that the port is set to 1433 or the default port of SQL Server during installation.



16. In Failover Cluster Manager, restart SQL Server by selecting Roles.

17. Restart the SQL Server Cluster Role by selecting Stop Role and then Start Role.



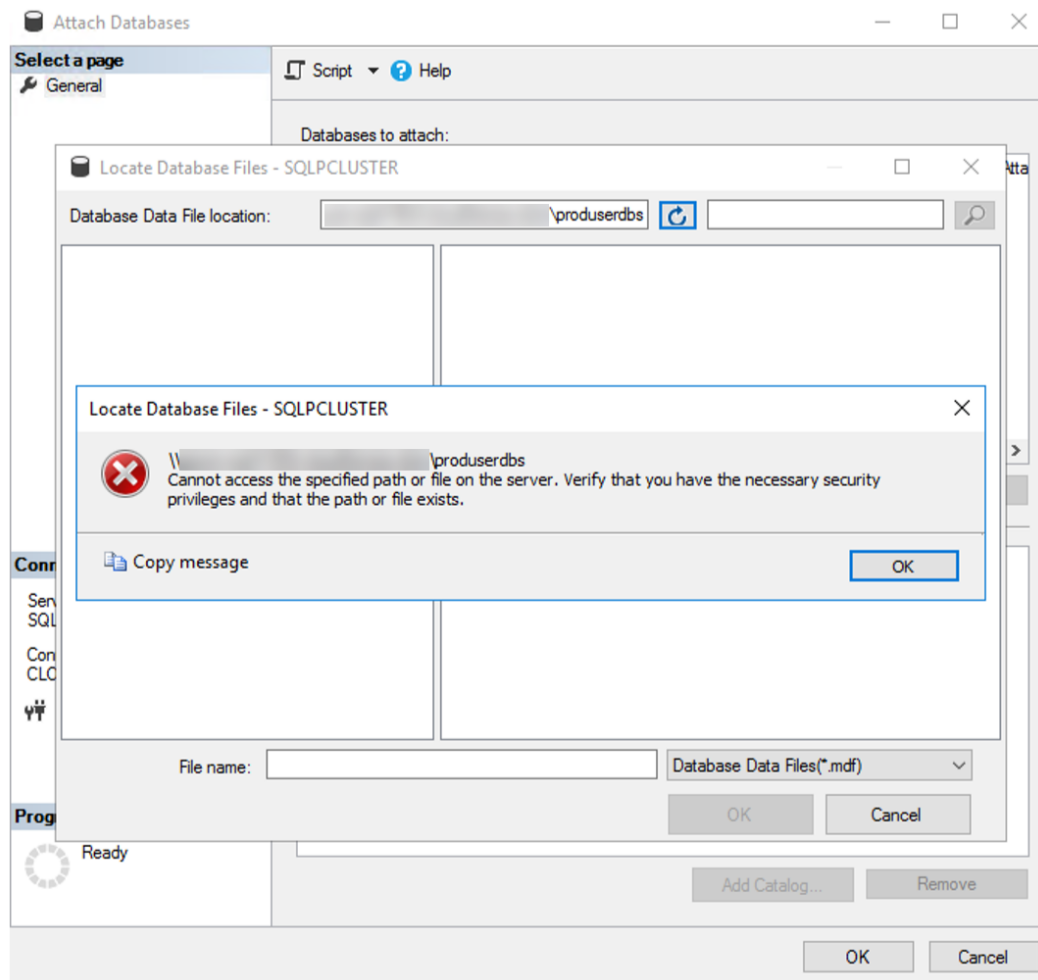
Note: You must deploy an internal load balancer that uses a health check to ensure that traffic is directed to the active node of the Windows Failover Cluster.

Commonly known errors

If you use T-SQL to attach or restore databases, a message like this might appear in query result and log in SQL Server log.

```
Msg 5120, Level 16, State 101, Line 1 Unable to open the physical file  
"\\servername\sharename\filename.mdf". Operating system error 5: (Access is denied.).
```

If you are using SQL Server Management Studio, an error message like this might occur.



To resolve, follow these steps:

1. Ensure that the SQL Server service account has full access to the Cloud Volumes Service volumes.
2. Ensure that SQL Server service has full access to the data and log files.
3. In [SQL Server Configuration Manager](#) (for information about 1802, see [Trace Flags](#)), use this startup option to turn on trace flag 1802. For more information about how to change the startup parameters, see [Database Engine Service Startup Options](#).
4. To reattach the database instead of SQL Server Management Studio, run the following T-SQL command:


```
exec sp_attach_db DatabaseName, '\\Network-attached storage_Path\DatabaseMDFFile.mdf',  
'\\Network-attached storage_Path\DatabaseLDFFile.ldf'  
go
```

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp Product Documentation
<https://www.netapp.com/support-and-training/documentation/>

Version history

Version	Date	Document version history
Version 1.0	August 2021	Initial release.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2021 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4898-0821