

How to Build a SANless SQL Server Failover Cluster Instance in Google Cloud Platform with SIOS DataKeeper



If you are going to host SQL Server on the Google Cloud Platform (GCP) you will want to make sure it is highly available. One of the best and most economical ways to do that is to build a SQL Server Failover Cluster Instance (FCI). Since SQL Server Standard Edition supports Failover Clustering, we can avoid the cost associated with SQL Server Enterprise Edition which is required for Always On Availability Groups. In addition, SQL Server Failover Clustering is a much more robust solution as it protects the entire instance of SQL Server, has no limitations in terms of DTC (Distributed Transaction Coordinator) support and is easier to manage. Plus, it supports earlier versions of SQL Server that you may still have, such as SQL 2012 through the latest SQL 2017. Unfortunately, SQL 2008 R2 is not supported due to the lack of support for cross-subnet failover.

Traditionally, SQL Server FCI requires that you have a SAN or some type of shared storage device. In the cloud, there is no cluster-aware shared storage. In place of a SAN, we will build a SANless cluster using SIOS DataKeeper Cluster Edition (DKCE). DKCE uses block-level replication to ensure that the locally attached storage on each instance remains in sync with one other. It also integrates with Windows Server Failover Clustering through its own storage class resource called a DataKeeper Volume which takes the place of the physical disk resource. As far as the cluster is concerned the SIOS DataKeeper volume looks like a physical disk, but instead of controlling SCSI reservations, it controls the mirror direction, ensuring that only the active server writes to the disk and that the passive server(s) receive all the changes either synchronously or asynchronously.

In this guide, we will walk through the steps to build a two-node failover cluster between two instances in the same region, but in different Zones, within the GCP as shown in Figure 1.

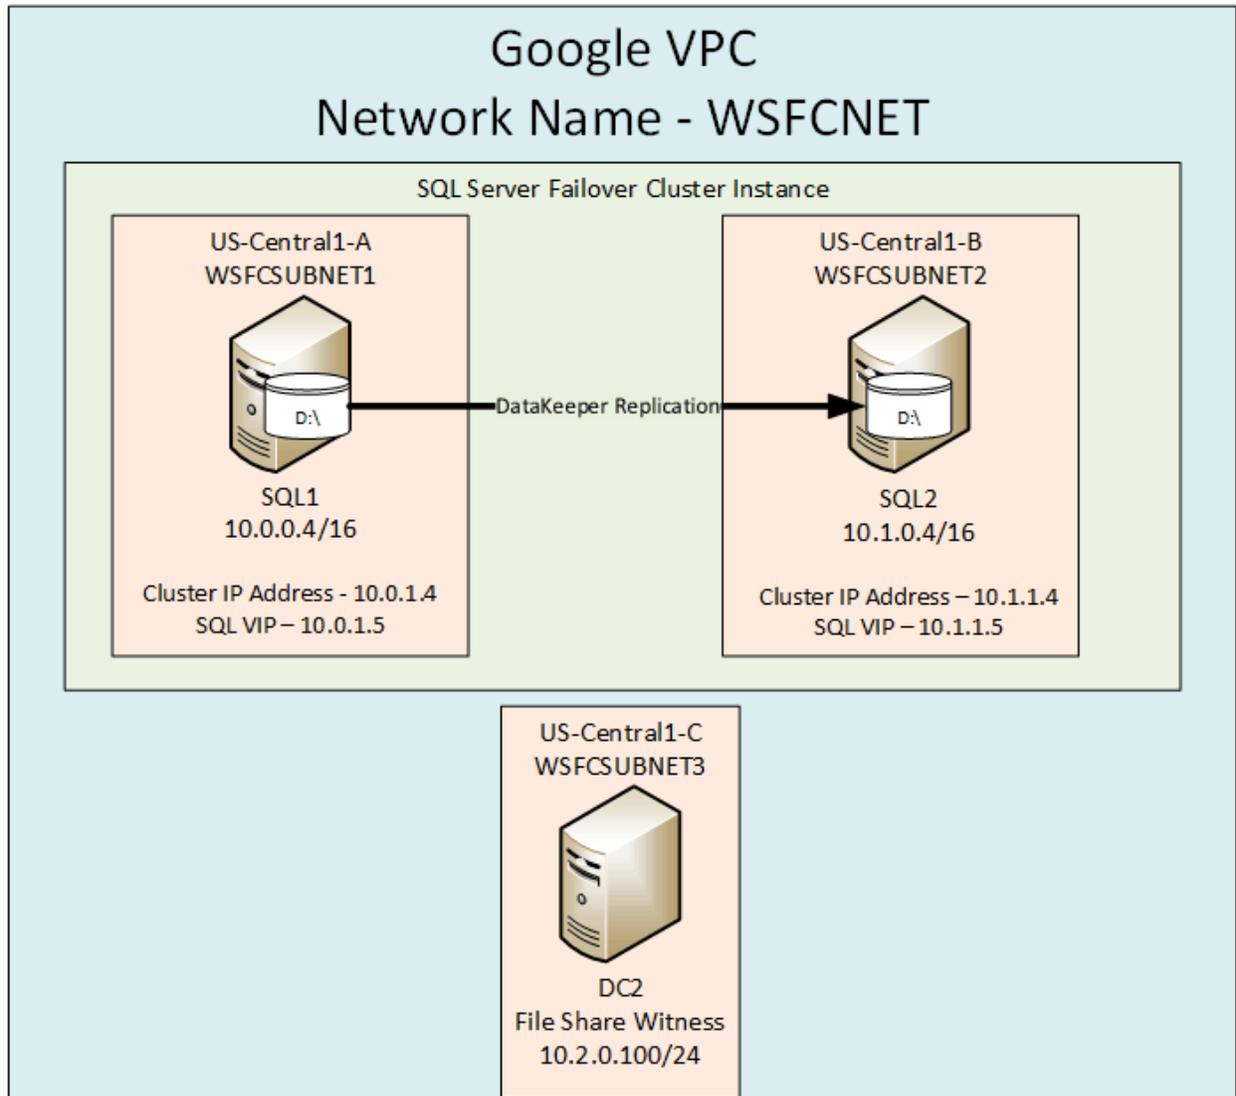


Figure 1 – SANless WSFC cluster created using SIOS DataKeeper with two nodes in the same region but different zones within the GCP.

Create a Custom Mode VPC Network

If you are familiar with failover clustering, you will be pleased to know that *most* of the process is the same in GCP. However, one of the biggest differences occurs at the network layer. In theory, you would think that you can just build one big subnet and put all of your nodes in it and use a virtual IP address, just like you would in your own data center. However, like some other cloud providers, such as Azure, the virtual network does not support connecting directly to the virtual IP address, due to the lack of support for gratuitous ARPs (address resolution protocols). In GCP they work around this for failover clustering by putting each node in a different subnet, and then creating host-specific routes for the cluster IP addresses. If networking is unfamiliar, try to follow the steps below as best you can; I'll explain each step as I go along.

First things first. We want to create our own virtual network. Using the GCP console, create your own custom network as shown below.

```
gcloud compute networks create wsfcnet --subnet-mode custom
```

1. Once the network is created you will add three subnets to the network. Notice that I am using a /24 subnet. I will explain more about that later.

```
gcloud compute networks subnets create wsfcsubnet1 --network wsfcnet \  
--region us-central1 --range 10.0.0.0/24
```

```
gcloud compute networks subnets create wsfcsubnet2 --network wsfcnet \  
--region us-central1 --range 10.1.0.0/24
```

```
gcloud compute networks subnets create wsfcsubnet3 --network wsfcnet \  
--region us-central1 --range 10.2.0.0/24
```

2. Create a firewall rule to allow traffic between the instances on internal IP addresses on the new VPC network.

```
gcloud compute firewall-rules create allow-internal-ports \  
--network wsfcnet --allow tcp:1-65535,udp:1-65535,icmp \  
--source-ranges 10.0.0.0/24,10.1.0.0/24,10.2.0.0/24
```

3. Create a firewall rule to allow RDP on port 3389 on the VPC network.

```
gcloud compute firewall-rules create allow-rdp --network wsfcnet \  
--allow tcp:3389 --source-ranges 0.0.0.0/0
```

You can verify that the rules are in place if you like.

 You have \$233.97 in credit and 358 days left in your free trial.

Google Cloud Platform SANLESS-FCI-2

VPC network ← VPC network details EDIT DELETE VPC NETWORK

VPC networks

- External IP addresses
- Firewall rules
- Routes
- VPC network peering
- Shared VPC

wsfcnet

Subnet creation mode
Custom subnets

Dynamic routing mode
Regional

Subnets Static internal IP addresses Firewall rules Routes VPC Network Peering

<input type="checkbox"/> Name ^	Region	IP address ranges	Gateway	Private Google access
<input type="checkbox"/> wsfcsubnet1	us-central1	10.0.0.0/24	10.0.0.1	Disabled
<input type="checkbox"/> wsfcsubnet2	us-central1	10.1.0.0/24	10.1.0.1	Disabled
<input type="checkbox"/> wsfcsubnet3	us-central1	10.2.0.0/24	10.2.0.1	Disabled

Firewall rules CREATE FIREWALL RULE REFRESH DELETE

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. [Learn more](#)

Note: App Engine firewalls are managed [here](#).

Ingress Egress

<input type="checkbox"/> Name	Targets	Source filters	Protocols / ports	Action	Priority	Network ^
<input type="checkbox"/> default-allow-icmp	Apply to all	IP ranges: 0.0.0.0/0	icmp	Allow	65534	default
<input type="checkbox"/> default-allow-internal	Apply to all	IP ranges: 10.128.0.0/9	tcp:0-65535, udp:0-65535, 1 more ▾	Allow	65534	default
<input type="checkbox"/> default-allow-rdp	Apply to all	IP ranges: 0.0.0.0/0	tcp:3389	Allow	65534	default
<input type="checkbox"/> default-allow-ssh	Apply to all	IP ranges: 0.0.0.0/0	tcp:22	Allow	65534	default
<input type="checkbox"/> allow-rdp	Apply to all	IP ranges: 0.0.0.0/0	tcp:3389	Allow	1	wsfcnet
<input type="checkbox"/> allow-internal-ports	Apply to all	Subnetworks: wsfcsubnet1, 2 more ▾	tcp:1-65535, udp:1-65535, 1 more ▾	Allow	1000	wsfcnet

Now we are ready to create our instances.

Create Instances

Here we will create three instances. DC1, SQL1, and SQL2. Each instance will reside in a different subnet and in a different zone. We will also assign a static IP address and enable IP forwarding on each instance. Later, when we configure each instance, we will use *netsh* to change the subnet mask permanently to a /16 address which will be required to support the host-specific routing needed for clustering (as mentioned earlier). That step will be explained later in this document. DC1 will be our domain controller and also act as a file share witness for the cluster.

For DC we will use a standard Windows Server 2016 image. For SQL1 and SQL2 we will use a Windows Server 2016 image that has SQL Server pre-installed on it. We will have to uninstall that standalone instance of SQL Server and reinstall it on each server as a clustered instance once we build the basic cluster. In this example, we will use SQL Server 2016 Standard Edition. However, you can build a SQL Server Failover Cluster Instance (FCI) with any version of Windows (2012 R, 2016) or any version of SQL Server Standard or Enterprise Edition (2012, 2014, 2016, 2017). SQL 2008 R2 and earlier versions are not supported in Google as cross-subnet failover support was not added until SQL 2012. You also have the option of bringing your own license of SQL Server to the cloud as well. Just be sure to check with your Microsoft licensing specialists.

We will also add an additional data disk to both SQL1 and SQL2, which we will mirror with SIOS DataKeeper to be used as the cluster storage. You may add more than one disk if you like, but for simplicity, we will use one disk.

Below I walk you through the process of provisioning the instances via the GUI. You should consider scripting this to make the process easier moving forward.

← Create an instance

Name [?](#)

Zone [?](#)

Machine type

Customize to select cores, memory and GPUs.

7.5 GB memory

[Customize](#)

[Upgrade your account](#) to create instances with up to 96 cores

Container [?](#)

Deploy a container image to this VM instance. [Learn more](#)

Boot disk [?](#)



New 50 GB standard persistent disk

Image

Windows Server 2016

[Change](#)

If you are using Windows and intend to run additional Microsoft software, please fill out the [License Verification Form](#)

[Learn more](#) about Microsoft license mobility requirements

The important thing to note here is to make sure we are adding it to the right subnet and also enabling IP forwarding.

Network interfaces ?

Network interface

Network ?
wsfcnet

Subnetwork ?
wsfcsubnet3 (10.2.0.0/24)

Primary internal IP ?
Ephemeral (Automatic)

⌵ Show alias IP ranges

External IP ?
Ephemeral

IP forwarding ?
On

Public DNS PTR Record ?
 Enable
PTR domain name

Done Cancel

Notice that I am choosing zone “a” for this instance. Each of the instances and the file share witness should reside in a different zone for maximum redundancy.

← Create an instance

Name

Zone

Machine type

Customize to select cores, memory and GPUs.

7.5 GB memory

[Customize](#)

[Upgrade your account](#) to create instances with up to 96 cores

Container

Deploy a container image to this VM instance. [Learn more](#)

Boot disk



New 50 GB standard persistent disk

Image

SQL Server 2016 Standard on Window...

[Change](#)

If you are using Windows and intend to run additional Microsoft software, please fill out the [License Verification Form](#)

[Learn more](#) about Microsoft license mobility requirements

Identity and API access

Service account

Access scopes

- Allow default access
- Allow full access to all Cloud APIs
- Set access for each API

Firewall

Add tags and firewall rules to allow specific network traffic from the Internet

- Allow HTTP traffic
- Allow HTTPS traffic

Create a disk

Name [?](#)

sql1data

Description (Optional)

cluster disk 1

Disk Type [?](#)

SSD persistent disk

Source type [?](#)

Image

Snapshot

None (blank disk)

Size (GB) [?](#)

100

Estimated performance [?](#)

Operation Type	Read	Write
Sustained random IOPS limit	3,000.00	3,000.00
Sustained throughput limit (MB/s)	48.00	48.00

Encryption [?](#)

Automatic (recommended)

Network interface 🗑️ ✕

Network ?

wsfcnet ▾

Subnetwork ?

wsfcsubnet1 (10.0.0.0/24) ▾

Primary internal IP ?

sql1ip (10.0.0.4) ▾

[⌵ Show alias IP ranges](#)

External IP ?

Ephemeral ▾

IP forwarding ?

On ▾

Public DNS PTR Record ?

Enable

PTR domain name

Network interfaces ?

Network interface ? 🗑 ✕

Network ?

wsfcnet ▼

Subnetwork ?

wsfcsubnet2 (10.1.0.0/24) ▼

Primary internal IP ?

sql2ip (10.1.0.4) ▼

[⌵ Show alias IP ranges](#)

External IP ?

Ephemeral ▼

IP forwarding ?

On ▼

Public DNS PTR Record ?

Enable

PTR domain name

Done Cancel

Deletion rule

Delete boot disk when instance is deleted

Encryption 

Automatic (recommended) 

Additional disks  (Optional)

Name	Mode	When deleting instance	
sql2data	<input type="text" value="Read/write"/> 	<input type="text" value="Keep disk"/> 	

 [Less](#)

VM instances

 [CREATE INSTANCE](#)

 [IMPORT VM](#)

 [REFRESH](#)



 Filter VM instances

<input type="checkbox"/>	Name 	Zone	Recommendation	Internal IP	External IP	Connect	
<input type="checkbox"/>	 dc1	us-central1-c		10.2.0.2	35.226.233.3	RDP 	
<input type="checkbox"/>	 sq1	us-central1-a		10.0.0.4	35.225.21.84	RDP 	
<input type="checkbox"/>	 sq2	us-central1-b		10.1.0.4	35.192.91.241	RDP 	

Configure the Servers

Connect to DC1 and promote it to a domain controller. Of course, you can skip this step if you already have an existing Active Directory Domain.

I recommend following the Google instructions and using the following PowerShell script to create your Active Directory Domain.

```
$DomainName = "datakeeper.local";
$DomainMode = "Win2012R2";
$ForestMode = "Win2012R2";
$DatabasePath = "C:\windows\NTDS";
$LogPath = "C:\windows\NTDS";
$SysvolPath = "C:\windows\SYSVOL";
Install-WindowsFeature -Name AD-Domain-Services -IncludeManagementTools
Install-ADDSForest -CreateDnsDelegation:$false -DatabasePath $DatabasePath -LogPath
$LogPath -SysvolPath $SysvolPath -DomainName $DomainName -DomainMode $DomainMode -
ForestMode $ForestMode -InstallDNS:$true -NoRebootOnCompletion:$false -
SafeModeAdministratorPassword ((Get-Credential).Password) -Force:$true
```

Update IP Addresses

Now that the Domain is created, we need to add SQL1 and SQL2 to the domain. Before we do that, we need to update the IP settings on each server. Earlier I said that even though the subnet has a /24 subnet, we are going to force our SQL1 and SQL2 VMs to use a /16 subnet so that we can create a host-specific route to the Cluster IP addresses running on those servers. Here is the first part of that step.

On SQL1 run the following command.

```
netsh interface ip set address name=Ethernet static 10.0.0.4 255.255.0.0 10.0.0.1 1
netsh interface ip set dns Ethernet static 10.2.0.100
```

On SQL2 run the following command.

```
netsh interface ip set address name=Ethernet static 10.1.0.4 255.255.0.0 10.0.0.1 1
netsh interface ip set dns Ethernet static 10.2.0.100
```

You may get some warnings about the DNS server not being a DNS server. This is most likely due to the Windows-based firewall. You will have to open the right ports or just turn off the firewall on DC1, SQL1, and SQL2 to complete the next step. At this point, join SQL1 and SQL2 to the domain as you normally would.

Update Routes

I mentioned earlier that we need to add some host-specific routes so that failover clustering will route traffic to the cluster IP address correctly. We are going to add those routes now. The following lines of code will add four routes. Notice that we are referencing some addresses that we have not yet seen. These are the addresses we will use when we configure our cluster, which will be a multisite cluster because each node is in a different subnet. While I'd rather have them all in the same subnet to simplify things, due to the network restrictions that I described earlier this is the best we can do. These commands should be run from the GCP console.

```
gcloud compute routes create cluster-sql1-route --network wsfcnct \  
--destination-range 10.0.1.4/32 --next-hop-instance sql1 \  
--next-hop-instance-zone us-central1-a --priority 1
```

```
gcloud compute routes create cluster-sql2-route --network wsfcnct \  
--destination-range 10.1.1.4/32 --next-hop-instance sql2 \  
--next-hop-instance-zone us-central1-b --priority 1
```

```
gcloud compute routes create cluster-sql1-route-listener --network wsfcnct \  
--destination-range 10.0.1.5/32 --next-hop-instance sql1 \  
--next-hop-instance-zone us-central1-a --priority 1
```

```
gcloud compute routes create cluster-sql2-route-listener --network wsfcnct \  
--destination-range 10.1.1.5/32 --next-hop-instance sql2 \  
--next-hop-instance-zone us-central1-b --priority 1
```

Create the Cluster

In the following steps we will create a basic cluster, and then install SQL Server into the cluster. Note the following IP addresses that we will be using for this process. These are the same addresses we used when we created the custom routes in the previous step.

SQL1

Cluster Core IP Address - 10.0.1.4

SQL VIP – 10.0.1.5

SQL2

Cluster Core IP Address – 10.1.1.4

SQL VIP – 10.1.1.5

Run the following PowerShell on both SQL1 and SQL2 to enable the failover clustering feature on both nodes.

```
Install-WindowsFeature Failover-Clustering -IncludeManagementTools
```

Run the following Powershell on SQL1 to validate the cluster.

```
Test-Cluster -Node sql1, sql2
```

```

PS C:\Windows\system32> Test-Cluster -Node sql, sql2
WARNING: System Configuration - Validate All Drivers Signed: The test reported some warnings..
WARNING: System Configuration - Validate Software Update Levels: The test reported some warnings..
WARNING: Network - Validate Network Communication: The test reported some warnings..
WARNING:
Test Result:
HadUnselectedTests, ClusterConditionallyApproved
Testing has completed for the tests you selected. You should review the warnings in the Report. A cluster solution is supported by Microsoft o
nly if you run all cluster validation tests, and all tests succeed (with or without warnings).
Test report file path: C:\Users\dave.DATAKEEPER\AppData\Local\Temp\Validation Report 2017.12.21 At 18.50.59.htm

Mode                LastWriteTime         Length Name
----                -
-a----            12/21/2017   6:51 PM           645360 Validation Report 2017.12.21 At 18.50.59.htm

```

Create the cluster by running the following PowerShell command from SQL1 or SQL2

```
New-Cluster -Name cluster1 -Node sql1,sql2 -NoStorage -StaticAddress 10.0.1.4,10.1.1.4
```

```

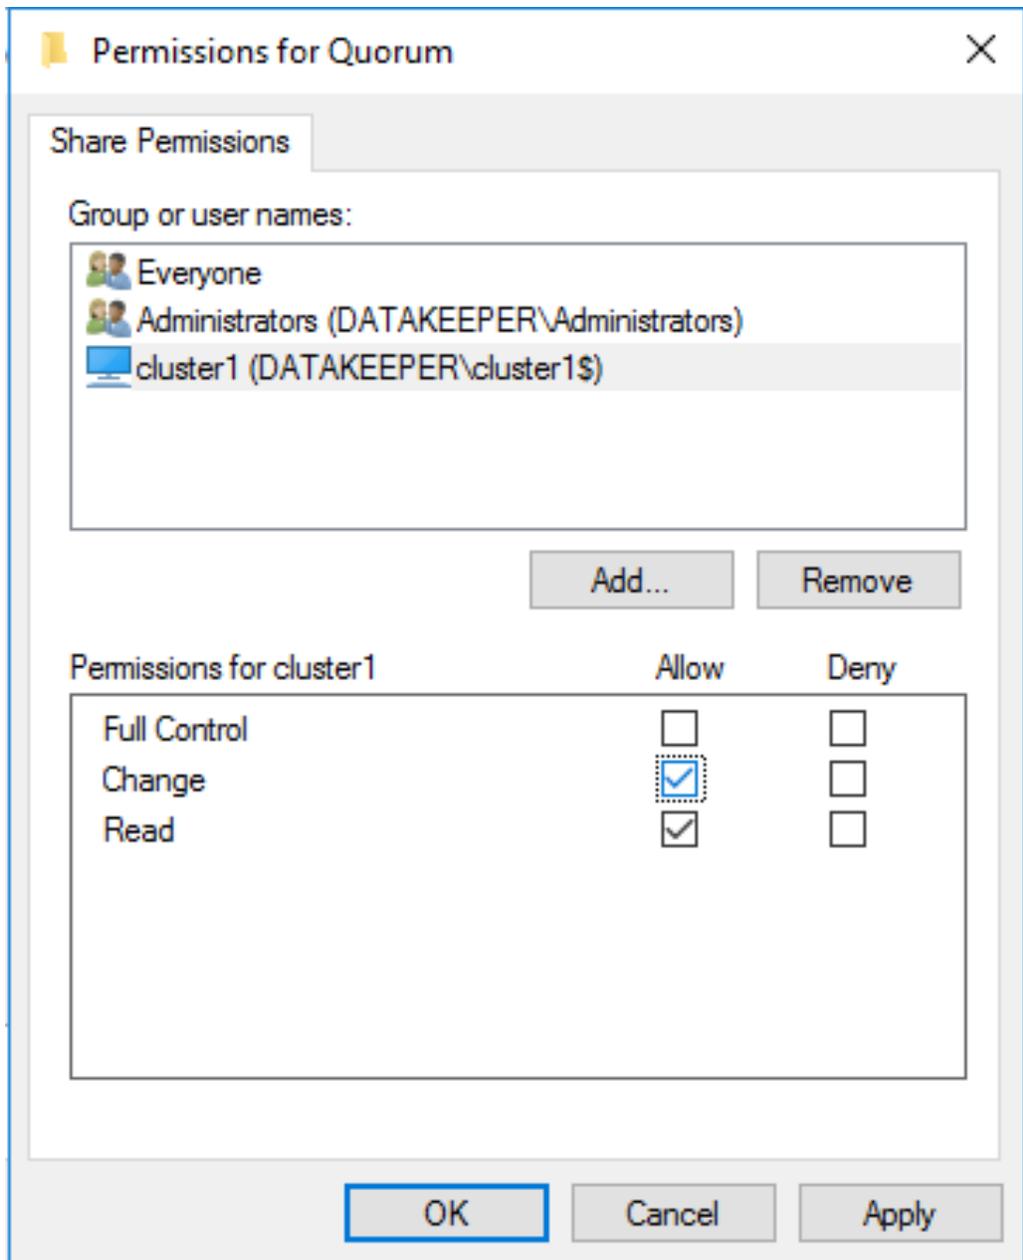
PS C:\Windows\system32> New-Cluster -Name cluster1 -node sql,sql2 -NoStorage -StaticAddress 10.0.1.4,10.1.1.4
WARNING: There were issues while creating the clustered role that may prevent it from starting. For more information view the report file below
WARNING: Report file location: C:\Windows\cluster\Reports\Create Cluster Wizard cluster1 on 2017.12.21 At 19.36.53.htm

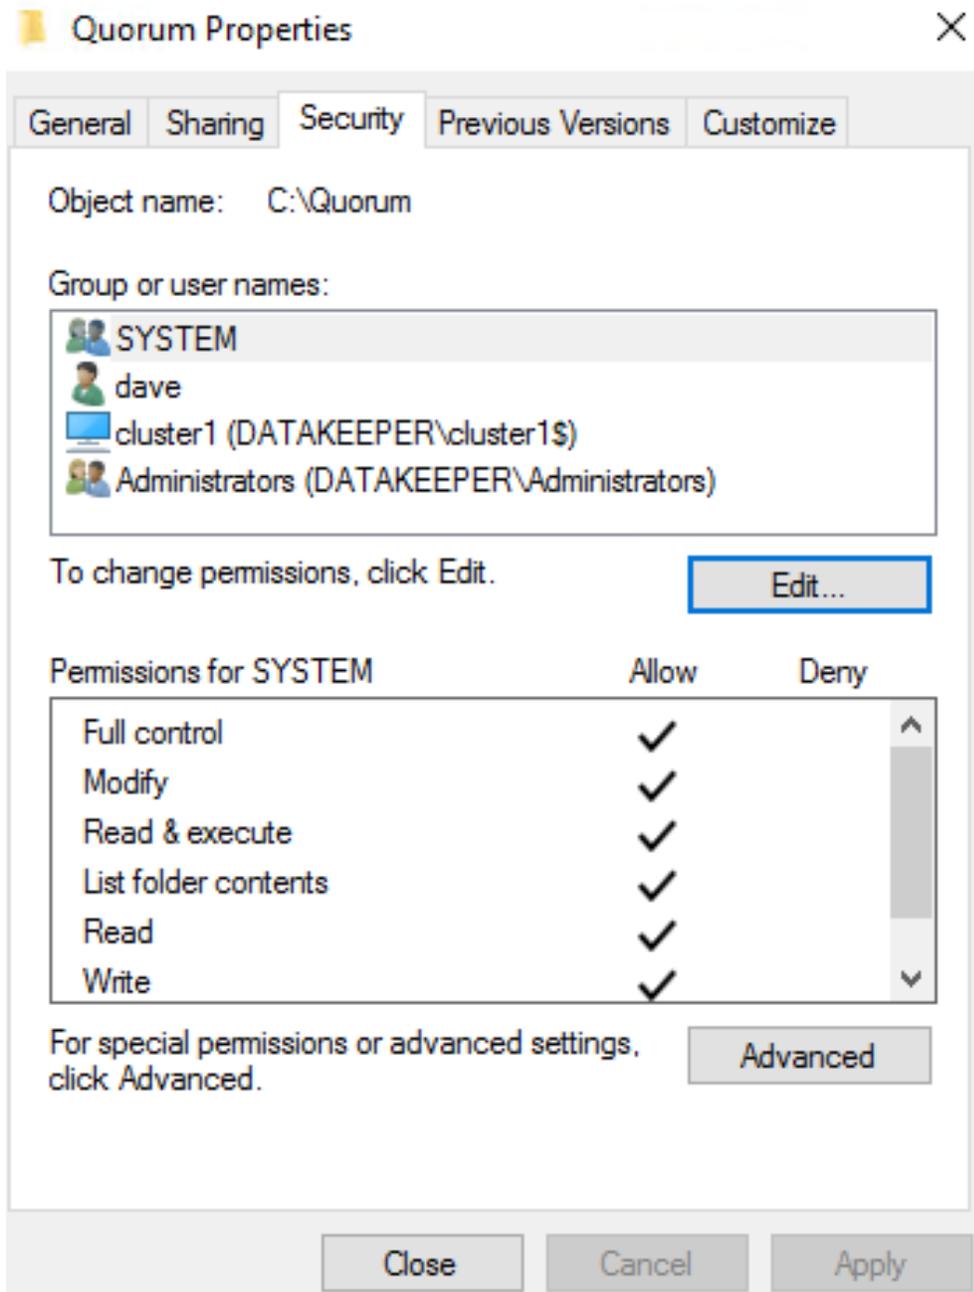
Name
----
cluster1

```

Once the cluster is created we will need to create a file share on DC1 and give the cluster computer object we just created (cluster1) read-write permissions at both the Share and NTFS level. Create and share this folder on DC1.







With the share created on DC1, we will now use PowerShell to add a File Share Witness to the cluster. Run this command from one of the cluster nodes.

```
set-clusterquorum -NodeAndFileshareMajority \\dc1\quorum
```

```
PS C:\Windows\system32> Set-ClusterQuorum -NodeAndFileShareMajority \\dc1\quorum

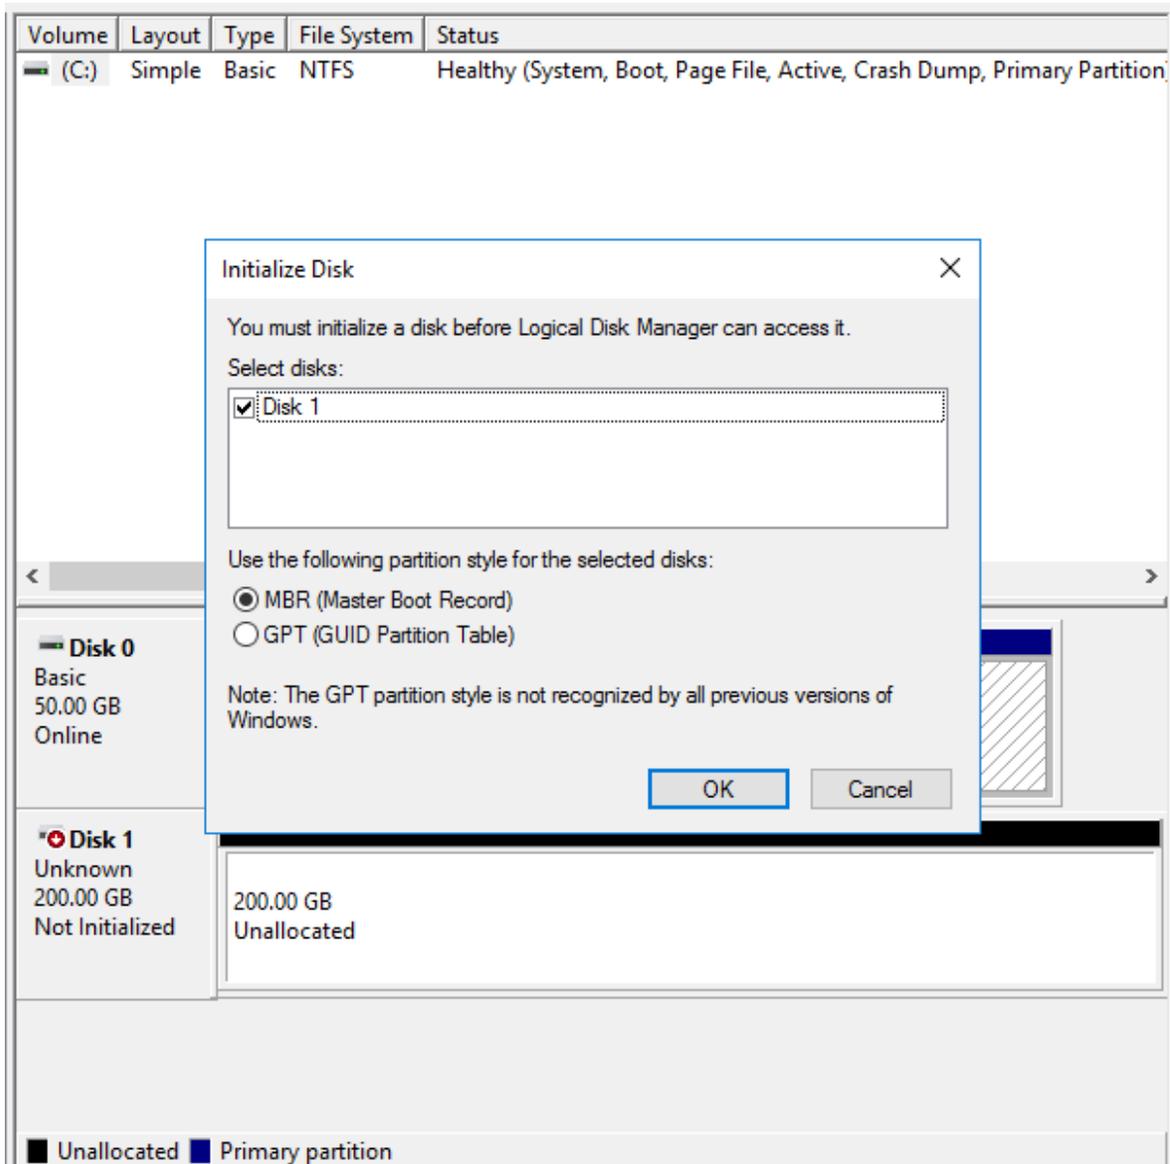
Cluster          QuorumResource
-----          -
cluster1         File Share Witness

PS C:\Windows\system32>
```

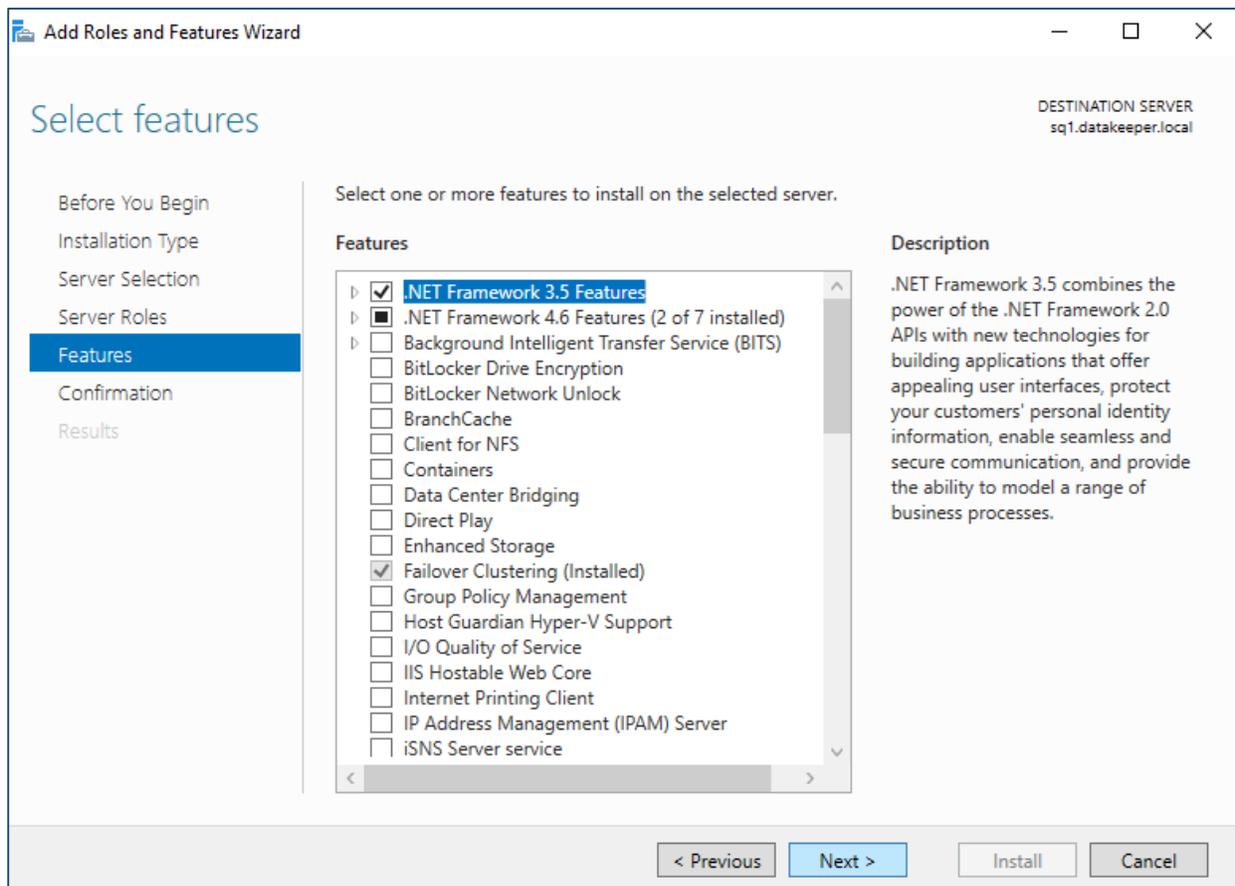
Now the basic cluster is configured, we are ready to install SQL into the cluster. However, before we do that, we need to configure the storage to be used by the cluster. In this case, we will use SIOS DataKeeper to replicate the locally attached disk from one server to another and register a DataKeeper Volume Resource in Failover Clustering. To create the DataKeeper Volume, follow the steps below.

If you have not already done so, attach an additional volume or volumes to each instance. Mount the volumes, partition and format them NTFS. You can use just one disk, multiple disks, or even put multiple disks in a Storage Pool and create a Storage Space that combines the IOPS and capacity of multiple disks together. If you decide to create a Storage Pool it is advisable to do that BEFORE you create the cluster. For our purposes, we are going to simply use one basic disk.

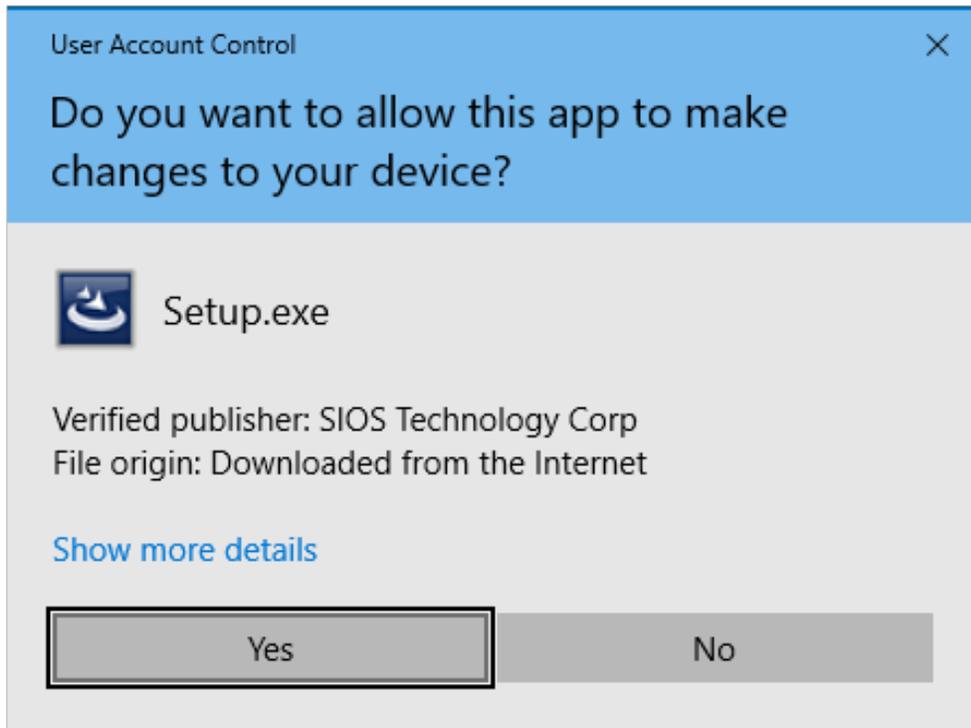




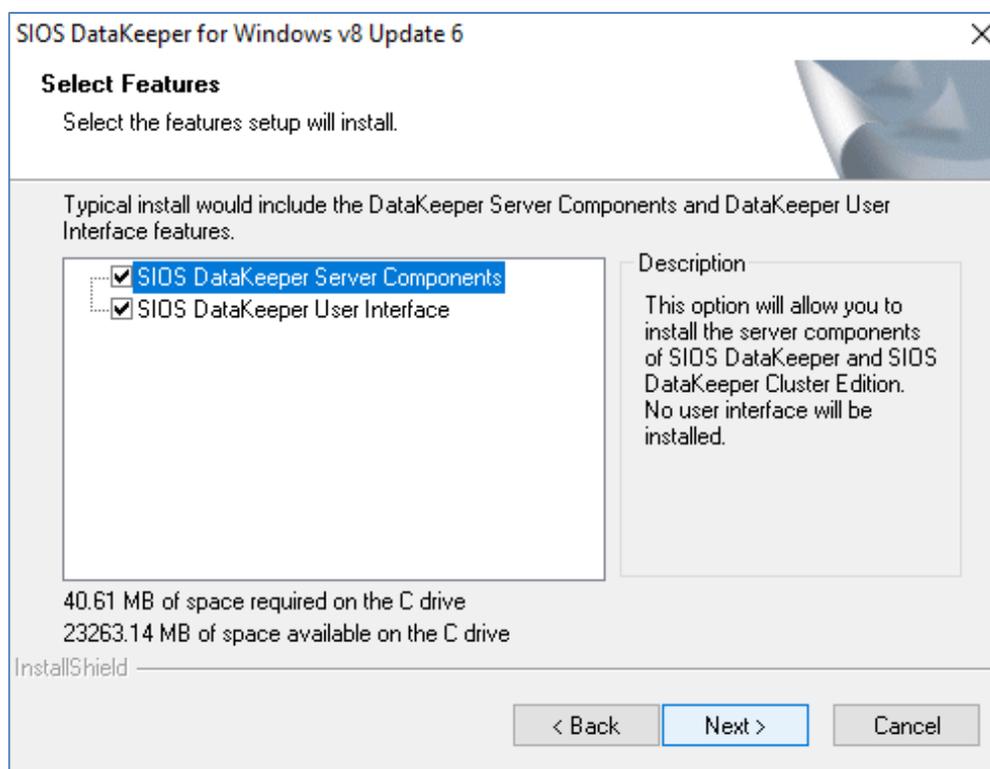
SIOS DataKeeper requires .Net Framework 3.5 to be enabled on each cluster node. Enable that now on both cluster nodes.



Now it is time to install SIOS DataKeeper Cluster Edition. Run the setup on each node, license, and reboot.



Install both components.



For the service account, use a domain account that is in the local Administrator's group on each server.

SIOS DataKeeper for Windows v8 Update 6

DataKeeper Service Logon Account Setup

Specify the user account for this service. (Format: Domain\UserID -or- Server\UserID)

User ID:
DATAKEEPER\dave

Password:
●●●●●●●●●●

Password Confirmation:
●●●●●●●●●●

InstallShield

< Back Next >

Launch the DataKeeper interface on one of the nodes and create your first Job

Connect to both servers.

SIOS DataKeeper

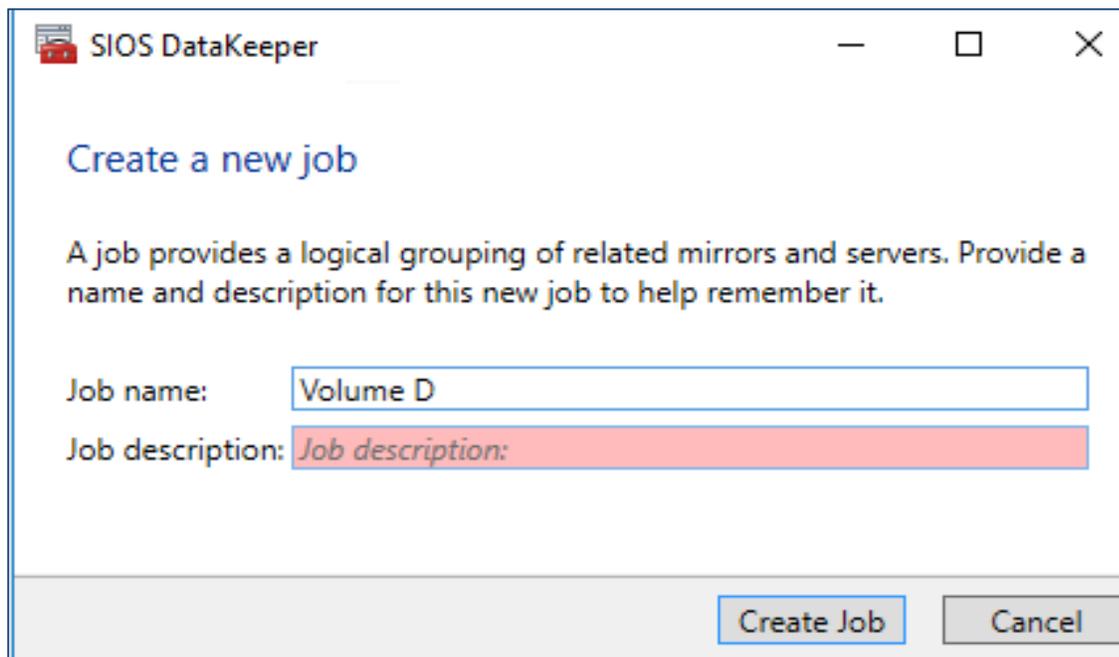
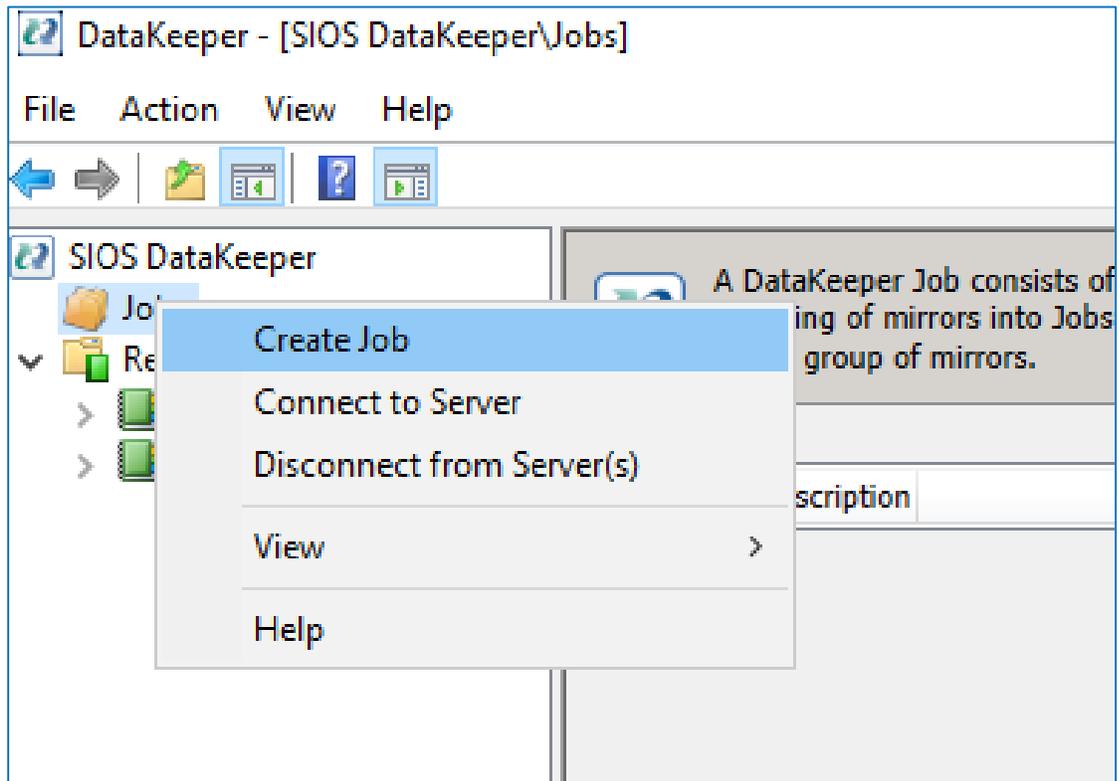
Enter the server to connect to

Provide the name or IP address for the server you would like to connect to.

Server:

Connect Cancel

Create a New Job



New Mirror

Choose a Source

Choose a Source

Choose a Target

Configure Details

Choose the server with the source volume.

Server: [Connect to Server](#)

Choose the IP address to use on the server.

IP address:

Choose the volume on the selected server.

Volume:

Next Cancel

New Mirror

Choose a Target

Choose a Source
 Choose a Target
 Configure Details

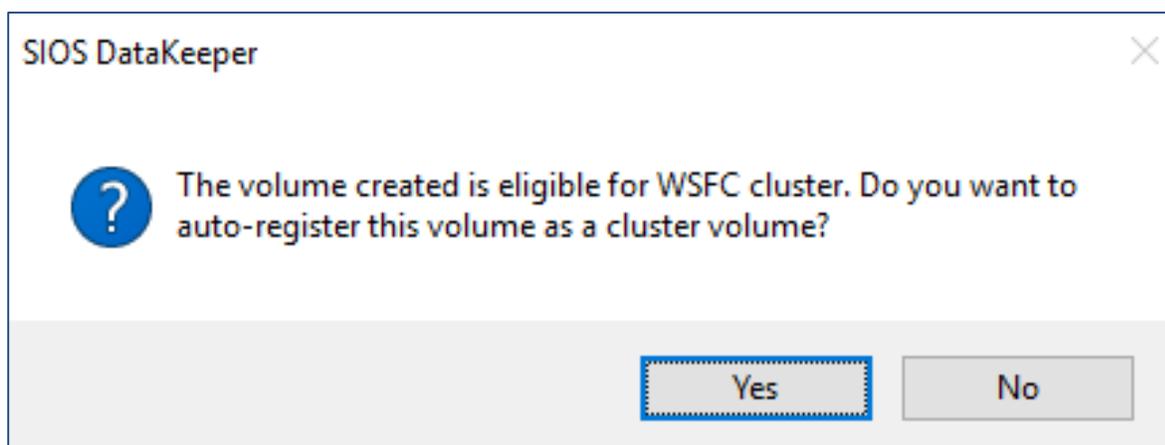
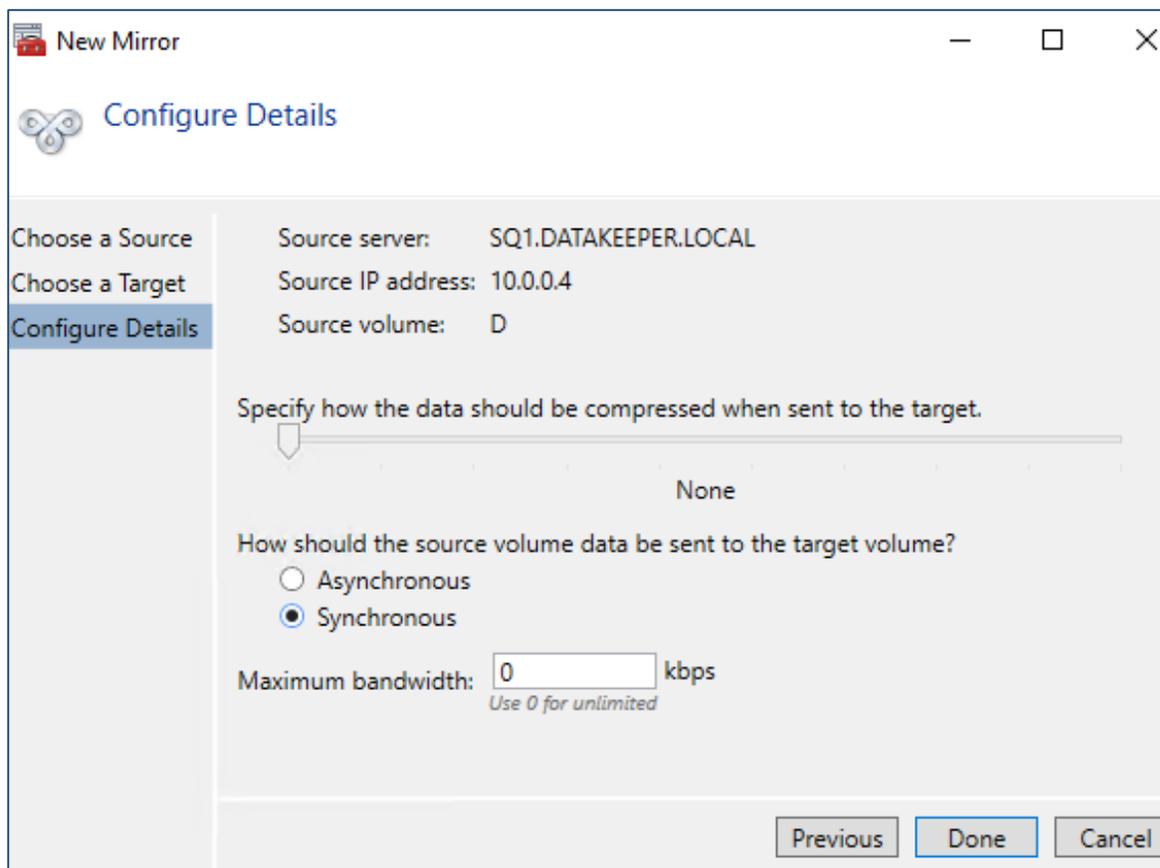
Source server: SQ1.DATAKEEPER.LOCAL
 Source IP address: 10.0.0.4
 Source volume: D

Choose the server with the target volume.
 Server: [Connect to Server](#)

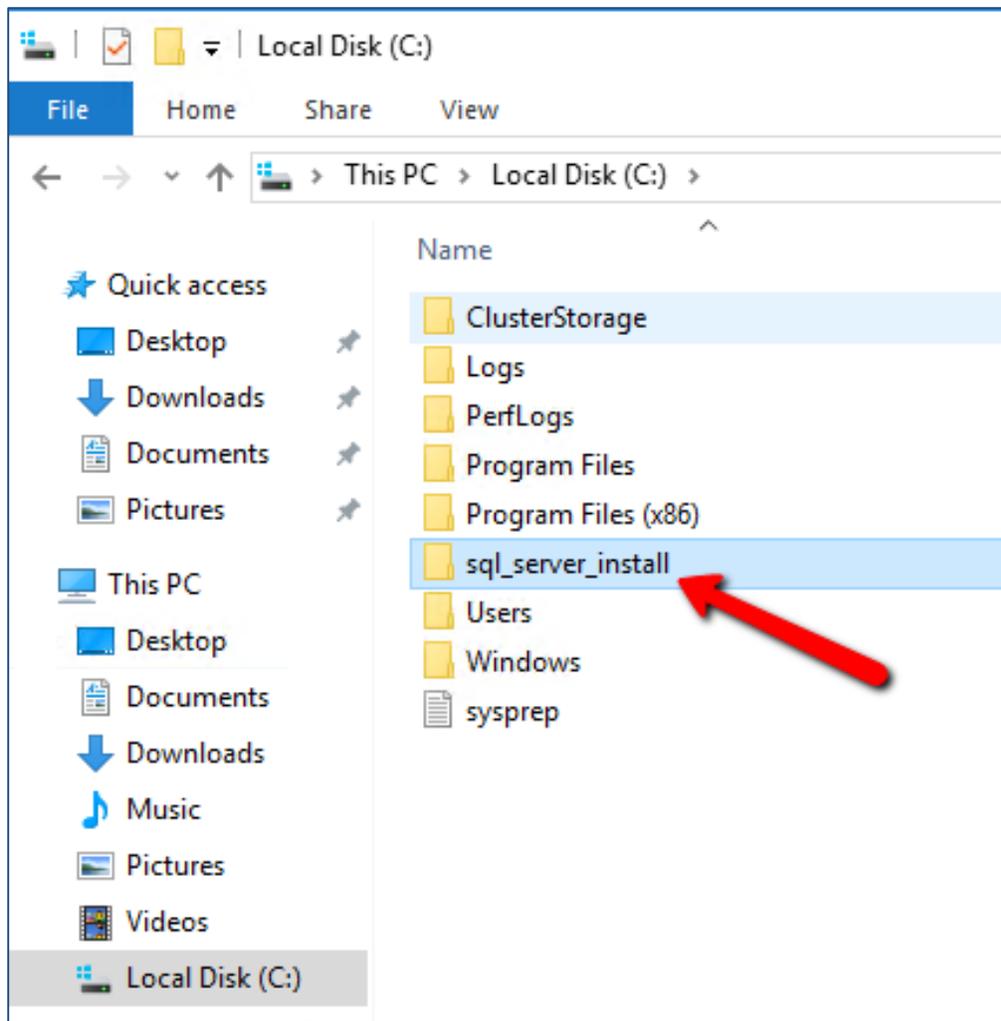
Choose the IP address to use on the server.
 IP address:

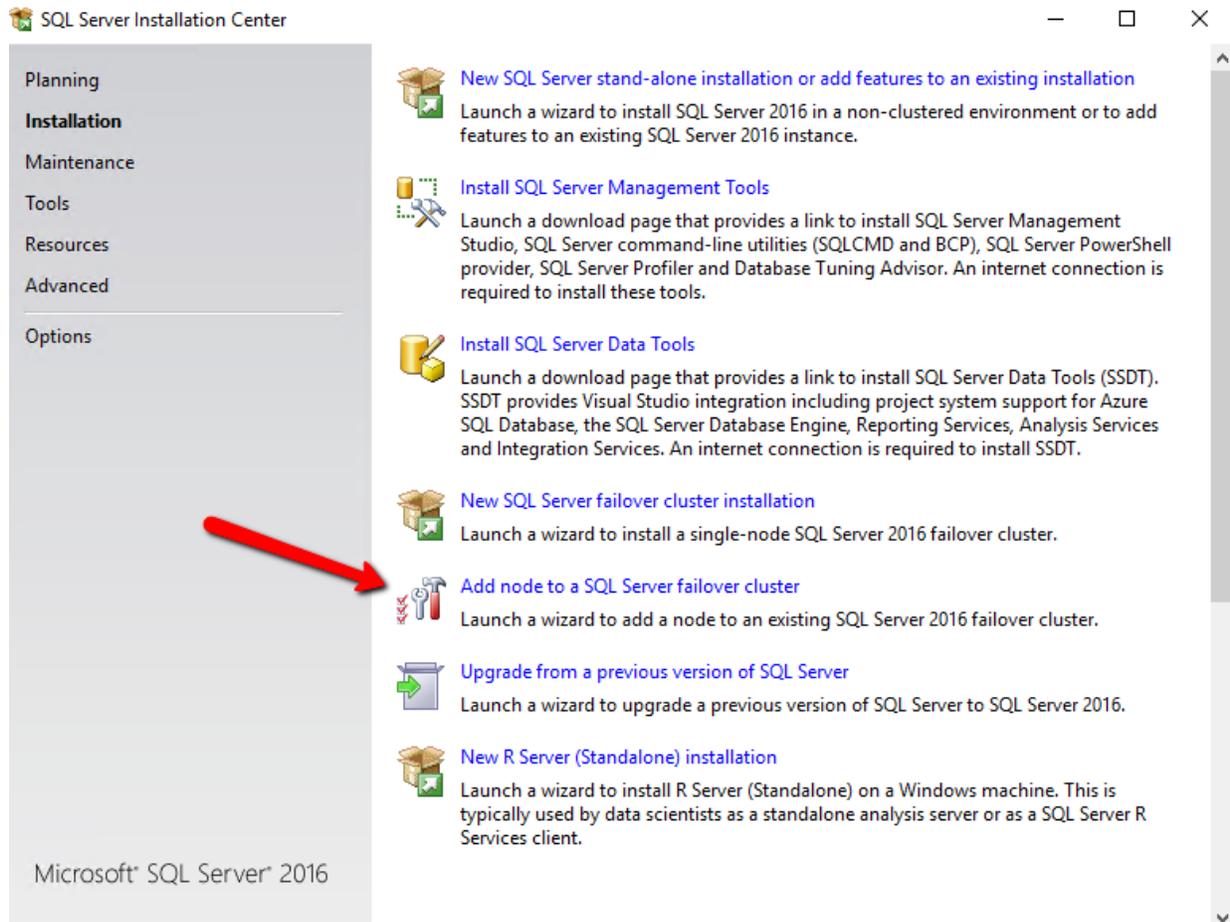
Choose the volume on the selected server.
 Volume:

For mirrors within the same region, choose synchronous mirroring. For replication between regions choose asynchronous replication.



Now that we have a DataKeeper Volume in Available Storage we are ready to install SQL Server into the cluster. However, if you have provisioned an instance of Windows with SQL already installed, the first step is to uninstall the existing standalone instance of SQL and reinstall it into the cluster. This needs to be done on both cluster nodes. Once you have uninstalled SQL, you will see a SQL install directory on the C drive.





You can ignore the one warning about cluster validation. We know that the cluster only has one network, but network redundancy is provided by the platform.

Install a SQL Server Failover Cluster

Install Failover Cluster Rules

Setup rules identify potential problems that might occur while running Setup. Failures must be corrected before Setup can continue.

Global Rules
 Microsoft Update
 Product Updates
 Install Setup Files
Install Failover Cluster Rules
 Product Key
 License Terms
 Feature Selection
 Feature Rules
 Feature Configuration Rules
 Ready to Install
 Installation Progress
 Complete

Operation completed. Passed: 20. Failed 0. Warning 1. Skipped 0.

Hide details << Re-run

[View detailed report](#)

Rule	Status
Distributed Transaction Coordinator (MSDTC) service	Passed
Microsoft Cluster Service (MSCS) cluster verification errors	Passed
Microsoft Cluster Service (MSCS) cluster verification warnings	Warning
Remote registry service (SQ1)	Passed
Domain controller	Passed
Microsoft .NET Application Security	Passed
Windows Firewall	Passed
DNS settings (SQ1)	Passed
WOW64 setup	Passed
Windows Management Instrumentation (WMI) service (SQL2)	Passed
Cluster Remote Access (SQL2)	Passed

< Back Next > Cancel

Install a SQL Server Failover Cluster

Feature Selection

Select the Standard features to install.

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Install Failover Cluster Rules
- Product Key
- License Terms
- 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

Features:

Instance Features

- Database Engine Services
- SQL Server Replication
- Full-Text and Semantic Extractions for Sea
- Data Quality Services
- PolyBase Query Service for External Data
- Analysis Services
- Reporting Services - Native

Shared Features

- Reporting Services - SharePoint
- Reporting Services Add-in for SharePoint Proc
- Data Quality Client
- Client Tools Connectivity

Feature description:

The configuration and operation of each instance feature of a SQL Server instance is isolated from other SQL Server instances. SQL Server instances can operate side-by-side on

Prerequisites for selected features:

Already installed:

- Windows PowerShell 3.0 or higher
- Microsoft Visual Studio 2010 Redistributable

Disk Space Requirements

Drive C: 1482 MB required, 23899 MB available

Select All
Unselect All

Instance root directory:

Shared feature directory:

Shared feature directory (x86):

< Back
Next >
Cancel

Instance Configuration

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

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Install Failover Cluster Rules
- Product Key
- License Terms
- 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\MSSQL13.MSSQLSERVER

Detected SQL Server instances and features on this computer:

Instance	Cluster Network Name	Features	Edition	Version	Inst
<Shared Comp...		Conn, BC, SDK		13.0.14500.10	
<Shared Comp...		DQC		13.0.1601.5	
<Shared Comp...		IS		13.1.4001.0	

< Back Next > Cancel

Install a SQL Server Failover Cluster

Cluster Resource Group

Create a new cluster resource group for your SQL Server failover cluster.

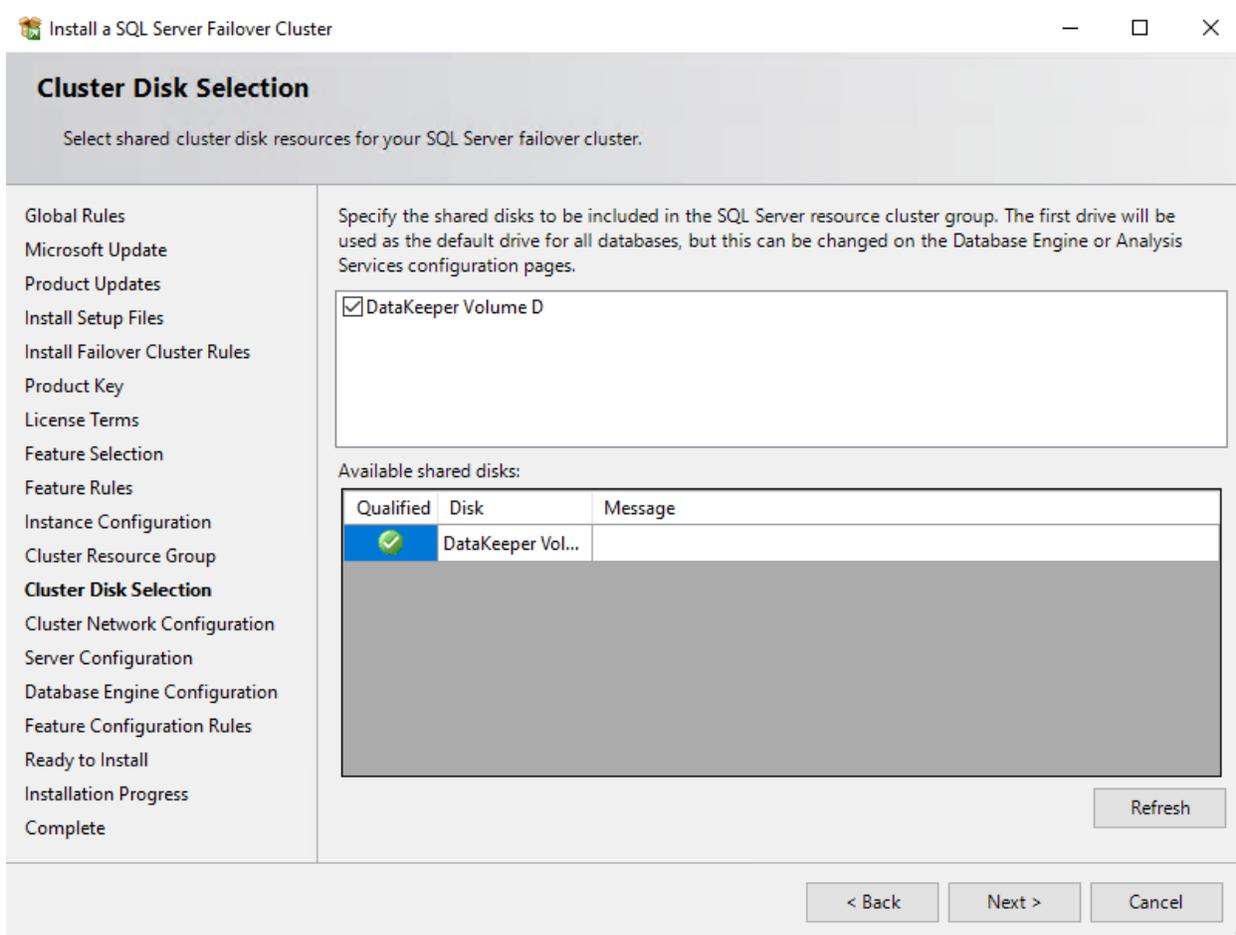
Specify a name for the SQL Server cluster resource group. The cluster resource group is where SQL Server failover cluster resources will be placed. You can choose to use an existing cluster resource group name or enter a new cluster resource group name to be created.

SQL Server cluster resource group name:

Qualified	Name	Message
	Available Storage	The cluster group 'Available Storage' is reserved by Windows Fai...
	Cluster Group	The cluster group 'Cluster Group' is reserved by Windows Failov...

Refresh

< Back Next > Cancel



Be sure to specify the same address we used when we created the custom routes earlier.

Cluster Network Configuration

Select network resources for your SQL Server failover cluster.

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Install Failover Cluster Rules
- Product Key
- License Terms
- 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 the network settings for this failover cluster:

<input checked="" type="checkbox"/>	IP Ty...	DHCP	Address	Subnet Mask	Subnet(s)	Network
<input type="checkbox"/>	IPv4	<input type="checkbox"/>		255.255.0.0	10.1.0.0/16	Cluster Network 1
<input checked="" type="checkbox"/>	IPv4	<input type="checkbox"/>	10.0.1.5	255.255.0.0	10.0.0.0/16	Cluster Network 2

This must be the same address we used earlier when we created the custom routes for the listener.



Refresh

< Back Next > Cancel

Server Configuration

Specify the service accounts and collation configuration.

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Install Failover Cluster Rules
- Product Key
- License Terms
- 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

Service Accounts Collation

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

Service	Account Name	Password	Startup Type
SQL Server Agent	DATAKEEPER\dave	●●●●●●●●	Manual ▾
SQL Server Database Engine	DATAKEEPER\dave	●●●●●●●●	Manual ▾
SQL Full-text Filter Daemon Launc...	NT Service\MSSQLFDLa...		Manual
SQL Server Browser	NT AUTHORITY\LOCAL ...		Automatic ▾

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

Database Engine Configuration

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

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Install Failover Cluster Rules
- Product Key
- License Terms
- 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

Specify the authentication mode and administrators for the Database Engine.

Authentication Mode

- Windows authentication mode
- Mixed Mode (SQL Server authentication and Windows authentication)

Specify the password for the SQL Server system administrator (sa) account.

Enter password:

Confirm password:

Specify SQL Server administrators

DATAKEEPER\dave (dave)

SQL Server administrators have unrestricted access to the Database Engine.

Add Current User | Add... | Remove

< Back | Next > | Cancel

Ready to Install

Verify the SQL Server 2016 features to be installed.

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Install Failover Cluster Rules
- Product Key
- License Terms
- 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

Ready to install the SQL Server 2016 failover cluster:

- [-] Summary
 - ... Edition: Standard
 - ... Action: InstallFailoverCluster (Product Update)
- [-] Prerequisites
 - [-] Already installed:
 - ... Windows PowerShell 3.0 or higher
 - ... Microsoft Visual Studio 2010 Redistributables
 - ... Microsoft .NET Framework 4.6
 - [-] 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 Database Engine: MSSQL 12 MSSQL SERVER

Configuration file path:

C:\Program Files\Microsoft SQL Server\130\Setup Bootstrap\Log\20180102_150049\ConfigurationFile.ini

< Back Install Cancel

Install a SQL Server Failover Cluster

Complete

Your SQL Server 2016 failover cluster installation is complete with product updates.

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Install Failover Cluster Rules
- Product Key
- License Terms
- 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**

Information about the Setup operation or possible next steps:

Feature	Status
Database Engine Services	Succeeded
Data Quality Services	Succeeded
Full-Text and Semantic Extractions for Search	Succeeded
SQL Server Replication	Succeeded
SQL Browser	Succeeded
SQL Writer	Succeeded

Details:

Product Update:
 Product Update has successfully applied KB 3182545 <<http://support.microsoft.com/?id=3182545>>. These updates have set the patch level of the Setup operation to 13.1.4001.0.

Summary log file has been saved to the following location:
C:\Program Files\Microsoft SQL Server\130\Setup Bootstrap\Log\20180102_150049\Summary_sq1_20180102_150049.txt

Close

We are now ready to install SQL on 2nd cluster node

SQL Server Installation Center

Planning
Installation
Maintenance
Tools
Resources
Advanced
Options

Microsoft® SQL Server® 2016

-  [New SQL Server stand-alone installation or add features to an existing installation](#)
Launch a wizard to install SQL Server 2016 in a non-clustered environment or to add features to an existing SQL Server 2016 instance.
-  [Install SQL Server Management Tools](#)
Launch a download page that provides a link to install SQL Server Management Studio, SQL Server command-line utilities (SQLCMD and BCP), SQL Server PowerShell provider, SQL Server Profiler and Database Tuning Advisor. An internet connection is required to install these tools.
-  [Install SQL Server Data Tools](#)
Launch a download page that provides a link to install SQL Server Data Tools (SSDT). SSDT provides Visual Studio integration including project system support for Azure SQL Database, the SQL Server Database Engine, Reporting Services, Analysis Services and Integration Services. An internet connection is required to install SSDT.
-  [New SQL Server failover cluster installation](#)
Launch a wizard to install a single-node SQL Server 2016 failover cluster.
-  [Add node to a SQL Server failover cluster](#)
Launch a wizard to add a node to an existing SQL Server 2016 failover cluster.
-  [Upgrade from a previous version of SQL Server](#)
Launch a wizard to upgrade a previous version of SQL Server to SQL Server 2016.
-  [New R Server \(Standalone\) installation](#)
Launch a wizard to install R Server (Standalone) on a Windows machine. This is typically used by data scientists as a standalone analysis server or as a SQL Server R Services client.

Add Node Rules

Setup rules identify potential problems that might occur while running Setup. Failures must be corrected before Setup can continue.

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Add Node Rules**
- Product Key
- License Terms
- Cluster Node Configuration
- Feature Rules
- Ready to Add Node
- Add Node Progress
- Complete

Operation completed. Passed: 20. Failed 0. Warning 1. Skipped 0.



Hide details <<

Re-run

[View detailed report](#)

Rule	Status
Microsoft Cluster Service (MSCS) cluster verification warnings	Warning
Remote registry service (SQL2)	Passed
Domain controller	Passed
Microsoft .NET Application Security	Passed
Windows Firewall	Passed
DNS settings (SQL2)	Passed
WOW64 setup	Passed
Windows Management Instrumentation (WMI) service (SQ1)	Passed
Cluster Remote Access (SQ1)	Passed
Distributed Transaction Coordinator (MSDTC) installed (SQ1)	Passed
Remote registry service (SQ1)	Passed

< Back

Next >

Cancel

Add a Failover Cluster Node

Cluster Node Configuration

Add a node to an existing SQL Server failover cluster.

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Add Node Rules
- Product Key
- License Terms
- Cluster Node Configuration**
- Cluster Network Configuration
- Service Accounts
- Feature Rules
- Ready to Add Node
- Add Node Progress
- Complete

SQL Server instance name: MSSQLSERVER

Name of this node: SQL2

Disk Space Requirements: Drive C: 1482 MB required, 27251 MB available

Instance Name	Cluster Network Name	Features	Nodes
MSSQLSERVER	SQLCLUSTER	SQLEngine, SQ...	SQL1

< Back Next > Cancel

Once again, be sure to use the same IP address we used when we create the custom routes earlier.

Add a Failover Cluster Node

Cluster Network Configuration

Specify additional IP addresses that are available and valid on the current node and subnet (previously-configured SQL Server failover cluster IP addresses are shown read-only and dimmed).

Specify the network settings for this failover cluster:

<input checked="" type="checkbox"/>	IP Ty...	DHCP	Address	Subnet Mask	Subnet(s)	Network
<input checked="" type="checkbox"/>	IPv4	<input type="checkbox"/>	10.1.1.5	255.255.0.0	10.1.0.0/16	Cluster Network 1
<input checked="" type="checkbox"/>	IPv4	<input type="checkbox"/>	10.0.1.5	255.255.0.0	10.0.0.0/16	Cluster Network 2

Remember, this is the same address we used when we created the custom routes earlier

Refresh

< Back Next > Cancel

Add a Failover Cluster Node

 SQL Server Setup detected that there are multiple subnets. Because of this, Setup sets the IP address resource dependency using an OR relationship for SQL Server multi-subnet failover clustering, so failover to other nodes does not happen until all the network cards fail on the node that owns the failover cluster. This may impact multi-homed cluster configurations on a subnet when client connections become unavailable. Do you want to proceed with SQL Server multi-subnet failover cluster configuration?

 Yes No

Service Accounts

Specify the service accounts and collation configuration.

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Add Node Rules
- Product Key
- License Terms
- Cluster Node Configuration
- Cluster Network Configuration
- Service Accounts**
- Feature Rules
- Ready to Add Node
- Add Node Progress
- Complete

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	DATAKEEPER\dave	●●●●●●●●	Manual
SQL Server Browser	NT AUTHORITY\LOCAL SE...		Automatic ▾
SQL Server Agent	DATAKEEPER\dave	●●●●●●●●	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

Ready to Add Node

Verify the SQL Server 2016 features to be installed as part of the add node operation.

- Global Rules
- Microsoft Update
- Product Updates
- Install Setup Files
- Add Node Rules
- Product Key
- License Terms
- Cluster Node Configuration
- Cluster Network Configuration
- Service Accounts
- Feature Rules
- Ready to Add Node**
- Add Node Progress
- Complete

Ready to add this node to the SQL Server 2016 failover cluster:

- [-] Summary
 - ... Edition: Standard
 - ... Action: AddNode (Product Update)
- [-] Prerequisites
 - [-] Already installed:
 - ... Windows PowerShell 3.0 or higher
 - ... Microsoft Visual Studio 2010 Redistributables
 - ... Microsoft .NET Framework 4.6
 - [-] 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 Database Engine: MSSQL 13 MSSQLSERVER

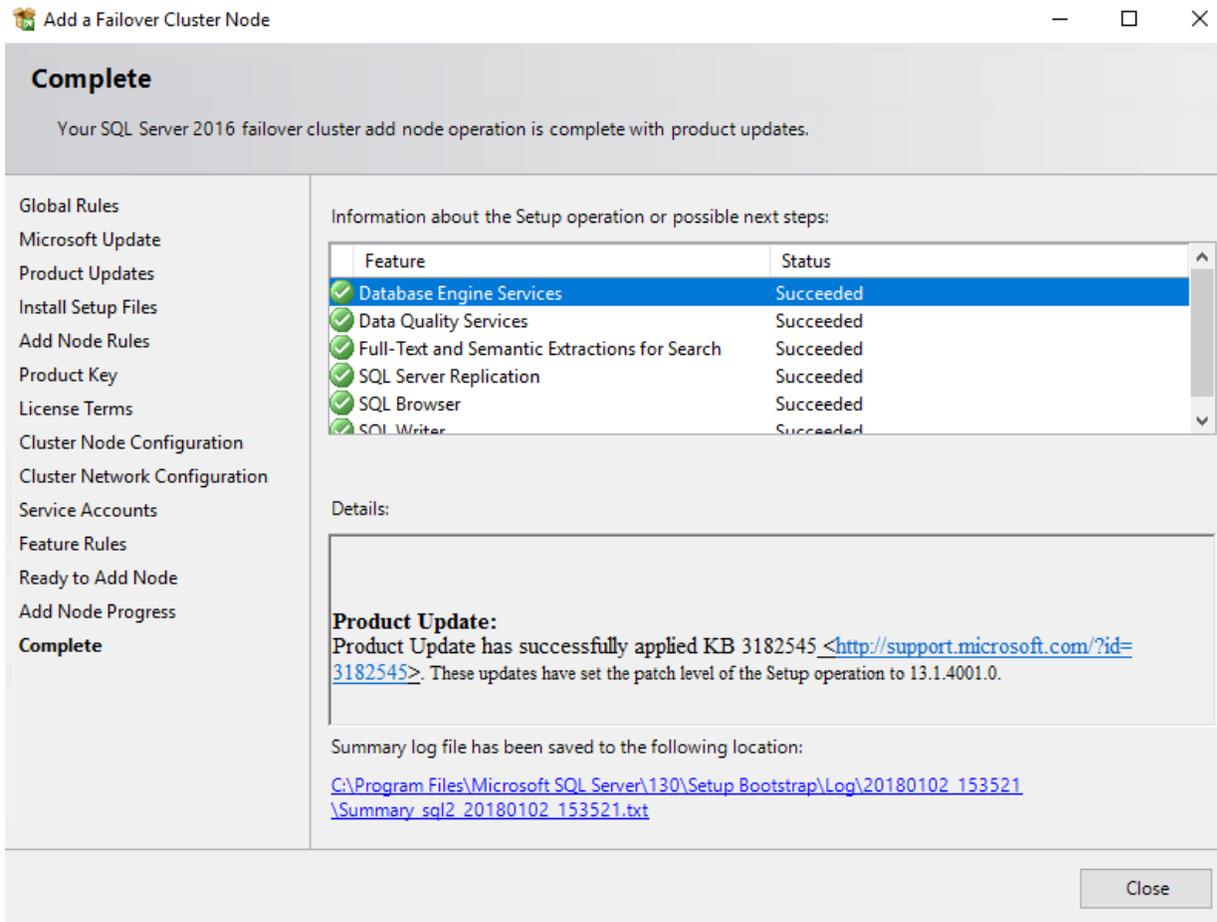
Configuration file path:

C:\Program Files\Microsoft SQL Server\130\Setup Bootstrap\Log\20180102_153521\ConfigurationFile.ini

< Back

Install

Cancel



Once the cluster is installed, try to connect to the cluster remotely from SQL Server Management Studio. If you can connect then you have done everything correctly. However, if you cannot connect, then chances are your firewall is not configured properly, or you have not configured the routing properly. If you can ping the hostname but not the cluster name there is a good chance the routing is not configured properly.

Configure Clients

As with any multisite cluster, you will need to configure your clients to connect properly. Most modern clients support the multisubnetfailover=true property in the connection string. For more information on multisite clusters, please consult the Microsoft documentation at <https://docs.microsoft.com/en-us/sql/sql-server/failover-clusters/windows/sql-server-multi-subnet-clustering-sql-server>

In Summary

Configuring a SANless multisite SQL Server Failover Cluster in GCP is an excellent solution for SQL Server high availability in the cloud. SIOS DataKeeper Cluster Edition provides replication and cluster integration software that allows you to build SANless clusters not only in the GCP but also on all other clouds, as well as on virtual and physical servers as well. Once you understand the networking requirements of the GCP you will have a very familiar user experience which leverages Windows Server Failover Clustering as the high availability solution. Because Windows Server Failover Clustering providing the high availability, this solution runs on both Windows Server 2012 R2 and Windows Server 2016 and supports SQL Server 2012 through 2017, both Standard and Enterprise Editions.

SIOS Technology Inc.
Tel: 650.645.7000.
<https://us.sios.com>

© 2023 SIOS Technology Corp. All rights reserved. SIOS, SIOS Technology, SIOS DataKeeper, SIOS LifeKeeper, SIOS Protection Suite and associated logos are registered trademarks or trademarks of SIOS Technology Corp. and/or its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. QS-1018-A

