

# Dell EMC Ready Solution for Microsoft SQL Server 2017 with Storage Spaces Direct Ready Node Cluster

## Hyper-converged Implementation

August 2018

H17293

## Deployment Guide

### Abstract

This deployment guide describes how to deploy Microsoft SQL Server 2017 on a Storage Spaces Direct Ready Node cluster.

Dell EMC Solutions

## Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA 08/18 Deployment Guide H17293.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.



# Contents

- Introduction ..... 4**
- About this guide ..... 4**
- Deployment overview ..... 5**
- Deploying SQL Server ..... 5**
- Deploying the Microsoft application agent for Data Domain Boost with SQL Server ..... 13**

## Introduction

Dell EMC Ready Solution for Microsoft SQL Server 2017 with Microsoft Storage Spaces Direct (S2D) is a preconfigured solution with preconfigured components. At the base of the solution are the versatile Dell EMC PowerEdge 14<sup>th</sup> Generation servers. Each server is a building block in the architecture's modular design, enabling a scalable and high-performance database infrastructure.

The solution is implemented on a hyper-converged infrastructure (HCI) that uses Dell EMC Ready Nodes for Microsoft Windows Server 2016 with Microsoft Hyper-V virtualization. The virtualized infrastructure solution can be deployed in either of the following ways:

- **Manual operating system deployment**—A manual installation starting from operating system deployment through cluster creation
- **Factory operating system deployment**—A factory installation where the Dell EMC Ready Nodes come from the factory with Windows Server 2016 already installed

## About this guide

### Purpose and scope

This guide provides information about deploying SQL Server 2017 on an HCI implementation of the Ready Solution. The HCI implementation uses Dell EMC Ready Nodes for Windows Server 2016 with Hyper-V and S2D. The guide focuses on SQL Server 2017 deployment and its optimal configuration.

For details about how to deploy, set up, and configure an HCI using S2D, see [Dell EMC Microsoft Storage Spaces Direct Ready Node: Deployment guide for scalable hyper-converged infrastructure with RX740xd and R640 Storage Spaces Direct Ready Nodes](#).

This guide is applicable only to infrastructure that is based on the validated and preconfigured Dell EMC Ready Nodes for Windows Server 2016 with Hyper-V.

### Audience

The target audience for this guide includes systems engineers, field consultants, partner engineering team members, and customers with some knowledge about how to deploy HCIs with Windows Server 2016 Hyper-V, S2D, and SQL Server 2017. Deployment personnel must have knowledge of the following technologies, concepts, and activities:

- Dell EMC PowerEdge servers, and deployment and configuration of the PowerEdge BIOS and iDRAC
- Dell EMC Networking switches, and concepts such as Data Center Bridging and Virtual Link Trunking
- Deployment and configuration of a Windows Server 2016 Hyper-V infrastructure
- Deployment and configuration of Windows System Center products such as SCOM if the solution deployment involves onboarding an S2D cluster into the existing System Center

## We value your feedback

Dell EMC and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the [Dell EMC Solutions team](#) with your comments.

Authors: Vaani Kaur, Sanjeev Ranjan

## Deployment overview

Dell EMC S2D Ready Nodes are preconfigured with server and switch deployments; therefore, the information in this guide is based on the assumption that the Ready Nodes have been deployed according to the procedures in [Dell EMC Microsoft Storage Spaces Direct Ready Node: Deployment guide for scalable hyper-converged infrastructure with RX740xd and R640 Storage Spaces Direct Ready Nodes](#).

This guide details the procedures for deploying and configuring the S2D Ready Node cluster up to virtual machine (VM) deployment. The guide also recaps some procedures that are relevant for SQL Server deployments. To support high-availability (HA) needs for any enterprise SQL Server deployment, multiple options are available at different deployment levels. For HA on a host level, we recommend Windows Host Clustering because it takes advantage of the hardware resiliency that is built into the S2D Ready Node. However, to further safeguard against operating system corruption and VM-level failures, we recommend using SQL Server clustering in conjunction with Windows Host Clustering. This guide addresses both types of clustering.

## Deployment workflow

Perform the procedures in the order presented in this guide to deploy SQL Server 2017 on a S2D Ready Node cluster.

## Deploying SQL Server

### Changing the host name

By default, the operating system deployment assigns a random name as the host computer name. For easier identification and uniform configuration, change the host name to something that is easily identifiable and relevant. You can change the host name by using the `Rename-Computer` cmdlet:

```
Rename-Computer -NewName S2DNode01 -Restart
```

---

**Note:** This command induces an automatic restart at the end of the rename operation.

---

### Joining cluster nodes to an Active Directory domain

The cluster nodes must be part of a Microsoft Active Directory domain before you can create a cluster. You can perform this task by using the `Add-Computer` cmdlet.

---

**Note:** Connecting to Active Directory Domain Services by using the host management network might require routing to the Active Directory network. Ensure that all the hosts in the S2D cluster can communicate with Domain Services before proceeding to join the domain.

---

```
$credential = Get-Credential
Add-Computer -DomainName S2dlab.local -Credential $credential -
Restart
```

---

**Note:** This command induces an automatic restart at the end of the domain join operation. Run the command on each host that will be part of the S2D cluster.

---

**Note:** Optionally, you might want to add all newly created computer objects from the HCI cluster deployment to a different Organizational Unit (OU) in the Active Directory domain. In this case, you can use the `-OUPath` parameter along with the `Add-Computer` cmdlet.

---

## Creating a host cluster

To prepare the nodes for cluster creation and create the host cluster:

1. Run the `Get-PhysicalDisk` command on all cluster nodes and review the output to ensure that:
  - All disks are in healthy state.
  - Each node has an equal number of disks.
  - The nodes have homogeneous hardware configurations.

2. Run the following cmdlet and review the output:

```
Test-Cluster -Node
S2DSYSTEM01,S2DSYSTEM02,S2DSYSTEM03,S2DSYSTEM04 -Include
'Storage Spaces Direct', 'Inventory', 'Network', 'System
Configuration'
```

The `Test-Cluster` cmdlet generates an HTML report of all validations performed and includes a summary of the validation.

3. Run the following cmdlet to create the new cluster and enable S2D, and review the output before proceeding with the deployment:

```
New-Cluster -Name S2DSYSTEM -Node
S2DSYSTEM01,S2DSYSTEM02,S2DSYSTEM03,S2DSYSTEM04 -
StaticAddress 172.16.102.55 -NoStorage -IgnoreNetwork
172.16.103.0/24, 172.16.104.0/24 -Verbose
```

---

**Note:** The `StaticAddress` parameter is used to specify a cluster IP address in the same IP subnet as the host management network. The `NoStorage` switch parameter specifies that the cluster will not include shared storage.

---

The `New-Cluster` cmdlet generates an HTML report of all configurations performed and includes a summary of the configurations.

## Creating and managing virtual disks

Creating the host cluster and enabling S2D on the cluster creates only a storage pool and does not provision any virtual disks in the storage pool. Run the `New-Volume` cmdlet to provision new virtual disks and, in turn, provision them as the cluster shared volumes for the S2D HCI cluster.

---

**Note:** You can use CSVs to deploy instances for SQL Server clustering.

---

The infrastructure and your requirements determine the number, size, and tiers of each of the volumes to be created. For general guidance on S2D volume creation, see [Creating volumes in Storage Spaces Direct](#) on the Microsoft website.

## Creating and configuring a VM

Create and configure a VM on the newly created virtual disk. Run the PowerShell commands in this procedure as an administrator.

For more details about any of these steps, see [Dell EMC Microsoft Storage Spaces Direct Ready Node: Deployment guide for scalable hyper-converged infrastructure with RX740xd and R640 Storage Spaces Direct Ready Nodes](#).

1. Create the VM:

```
$VMPath="C:\ClusteredStorage\volume1\VMs\" New-VM -
Name SP -Path $VMPath -MemoryStartupBytes <Memory> -
BootDevice CD -SwitchName <Switch Name> -VHDPATH $VHDPATH
```

2. Configure the VM:

```
Set-VMProcessor -VMName <VMname> -Count <vCPU Count> Add-
VMNetworkAdapter -VMName <VM Name> -Name <Virtual Network
Adapter> Set-VMNetworkAdapterVlan -VMName <VM Name> -
VMNetworkAdapterName <Virtual Network Adapter> -Access -
VlanId <vLAN ID>
```

3. Configure the virtual disk for operating system deployment:

```
Set-VMdvdDrive -VMName <VMName> -ControllerNumber
<ControllerNumber> -Path <ISOPATH>
```

4. Perform the following post-deployment activities:

- a. Post-deployment verification
- b. Day 0 operating system updates
- c. Operating system license activation

## Sample specifications for database VMs

The following table lists a possible set of specifications for the database VMs.

**Table 1. Sample of database VM specifications**

Component	Specification
Virtual processors	12 vCPUs
Memory	224 GB
OS VHDX	1 x 200 GB

## Initializing and formatting the storage volumes

Inside the VM created for SQL Server, the storage volumes are formatted with an allocation unit size of 64 KB to maximize the disk performance. Run the following Windows PowerShell commands to initialize and format the storage volumes on all the database servers.

---

**Note:** For easier manageability, ensure that all the database servers are assigned the same drive letter for a LUN.

---

1. Get the newly added storage volumes:  

```
Get-Disk | where-object Isoffline -eq $True
```
2. Initialize the disk:  

```
Initialize-Disk -Number <1> -PartitionStyle MBR
```
3. Partition the disk:  

```
New-Partition -DiskNumber <1> -DriveLetter 'E' -  
UseMaximumSize
```
4. Format the volume:  

```
Format-Volume -DriveLetter <'E'> -FileSystem REFS -  
AllocationUnitSize 65536 -Confirm:$false
```
5. Repeat the preceding steps for all the storage volumes to be assigned to SQL Server VMs.

### Verifying the cluster deployment

Verify the cluster deployment as follows:

1. Create VMs.
2. Migrate VMs across the cluster and verify that the migration is progressing without any interruption.
3. Shut down one of the hosts and verify that the VMs located on that host are migrating within the cluster.
4. Turn off one of the S4048-ON network switches and ensure that the VMs are running without any network loss.

### Configuring a SQL Server 2017 failover cluster

To configure database servers, perform the following tasks:

1. Create a VM.  
For more information, see [Creating and configuring a VM](#).
2. Install the operating system.
3. Rename and add a guest Windows Server to the Active Directory domain.  
For more information, see [Changing the host name](#) and [Joining cluster nodes to an Active Directory domain](#).
4. Install and configure Windows failover clustering.  
For more information, see [Creating a host cluster](#).



## Installing the SQL Server 2017 failover cluster

Install a SQL Server cluster instance in the following sequence.

---

**Note:** Before you begin, ensure that you have access to the installation media.

---

1. Run the setup in the first node.
2. Run the setup in all the other nodes where you want to install the SQL Server clustered instance.

### Installing the first node of the failover cluster

Install the first node of the SQL Server failover cluster as follows:

1. Initiate the installation:
  - a. Run `Setup.exe` as an administrator.  
The **SQL Server Installation Center** window is displayed.
  - b. In the left pane, click **Installation**, and then select **New SQL Server failover cluster installation**.
  - c. On the **Product key** page, enter the product key, and then click **Next**.
  - d. On the **License Terms** page, select **I Accept the license terms**, and then click **Next**.  
The installation wizard looks for rule errors and reports any errors in the **Global Rules** window.
  - e. On the **Microsoft Update** page, select **Use Microsoft Update to check for updates (recommended)**, and then click **Next**.
2. Click **Next** to proceed with the installation.
3. On the **Install Failover Cluster Rules** page, review and resolve the reported errors, if any, and then click **Next**.
4. On the **Setup Role** page, select **SQL Server Feature** Installation, and then click **Next**.
5. On the **Feature Selection** page, select **Database Engine Services and Management Tools - Basic**, and then click **Next**.
6. On the **Feature Rules** page, ensure that all the rules have passed, and then click **Next**.
7. On the **Instance Configuration** page, at **SQL Server Network Name**, enter the SQL Server instance network name, and then click **Next**.
8. On the **Cluster Resource Group** page, at **SQL Server Cluster Resource Group Name**, specify the SQL Server cluster resource group name, and then click **Next**.
9. On the **Cluster Disk Selection** page, select the default cluster disk selection, and then click **Next**.
10. On the **Cluster Network Configuration** page, select the IPv4 network and configure the IP address.

11. On the **Server Configuration** page:
  - a. Enter the respective domain accounts and passwords for the SQL Server agent and SQL Server engine.
  - b. On the **Collation** tab, click **Customize**, and then select **Latin1\_General** for **Collation Designation**.
  - c. On the **Collation** tab, select **Accent-Sensitive**, **Kana-Sensitive**, and **Width-Sensitive**, and then click **OK**.
12. On the **Database Engine Configuration** page:
  - a. On the **Server Configuration** tab, select **Windows authentication mode**, and then click **Add** to enter the SQL Server administrators.
  - b. On the **Data Directories** tab, define **Data root directory** (the path for disks) or **System database directory**, **User database directory**, **User database log directory**, and **Backup directory**.
  - c. On the **TempDB** tab, specify a location for the SQL database tempdb files and the number of tempdb datafiles to be created.
  - d. Click **Next**.
13. On the **Feature Configuration Rules** page, click **Next**.
14. On the **Ready to Install** page, review all the displayed configurations, and then click **Install**.

When the installation is complete, the **Complete** window is displayed.
15. Verify that all the items are successfully installed, and then click **Close**.

### Installing additional failover cluster nodes

Install additional failover cluster nodes as follows:

1. Initiate the installation:
  - a. Run the SQL Server setup as an administrator.
  - b. In the left pane, click **Installation**, and then select **Add node to a SQL Server failover cluster**.
  - c. On the **Product key** page, verify the product key, and then click **Next**.
  - d. On the **License Terms** page, accept the license terms, and then click **Next**.
  - e. On the **Microsoft Update** page, select **Use Microsoft Update to check for updates (recommended)**, and then click **Next**.
2. On the **Node Rules** page, review and resolve any reported errors, and then click **Next**.
3. On the **Cluster Node Configuration** page, select the SQL Server instance that you want to add as a node to the cluster.
4. On the **Cluster Network Configuration** page, select the IPv4 network and configure the IP address.
5. On the **Service Accounts** page, enter and confirm the SQL Server agent and SQL Server engine passwords.

6. On the **Feature Rules** page, review all the displayed rules, and then click **Next**.
7. On the **Ready to Add Node** page, verify the settings, and then click **Install**.
8. When the **Complete** window appears, click **Close** to exit the installation wizard.

## Configuring SQL Server for high performance

Configure SQL Server for high performance as follows:

1. Open SQL Server Management Studio, and click **New Query**.
2. Configure SQL Server memory settings.  
For more information, see [Configuring SQL Server memory settings](#).
3. Configure SQL Server MAXDOP settings.  
For more information, see [Configuring SQL Server MAXDOP settings](#).
4. Configure SQL Server LPIM settings.  
For more information, see [Configuring SQL Server LPIM settings](#).
5. Move tempdb to another volume.  
For more information, see [Moving the tempdb database to another volume](#).

### Configuring SQL Server memory settings

Run the following SQL query to set the maximum memory:

```
sp_configure 'show advanced options', 1;
GO
RECONFIGURE;
GO
sp_configure 'max server memory', <Memory in MB>;
GO
RECONFIGURE;
GO
```

### Configuring SQL Server MAXDOP settings

Run the following SQL query to set the MAXDOP settings:

```
USE Master;
GO
EXEC sp_configure 'show advanced options', 1;
GO
RECONFIGURE WITH OVERRIDE;
GO
EXEC sp_configure 'max degree of parallelism', <MAXDOP value>;
GO
RECONFIGURE WITH OVERRIDE;
GO
```

---

**Note:** Choose the MAXDOP setting that is most suitable for your SQL Server configuration.

---

## Configuring SQL Server LPIM settings

Set the **Lock pages in memory** (LPIM) option as follows:

1. In the Start menu, click **Run** and, in the **Open** box, type:  
`gpedit.msc`
2. In the **Local Group Policy Editor** console, expand **Computer Configuration**, and then expand **Windows Settings**.
3. Expand **Security Settings**, and then expand **Local Policies**.
4. Select the **User Rights Assignment** folder.  
The policies are displayed in the details pane.
5. Double-click **Lock pages in memory**.
6. In the **Local Security Setting – Lock pages in memory** dialog box, click **Add User or Group**.
7. In the **Select Users, Service Accounts, or Groups** dialog box, add an account that has privileges to run `sqlservr.exe`.
8. Log out and then log back in to implement the change.

## Moving the tempdb database to another volume

To isolate the tempdb I/O from other databases, move the tempdb database to a different volume as follows:

1. Open SQL Server Management Studio and click **New Query**.
2. To move the tempdb files, run the following SQL query:

```
USE master;
GO
ALTER DATABASE tempdb MODIFY FILE (NAME = <temp>, FILENAME =
<'E:\SQLData\tempdb.mdf'>);
GO
ALTER DATABASE tempdb MODIFY FILE (NAME = <temp2>, FILENAME
= <'E:\SQLData\tempdb_mssql_2.ndf'>);
GO
ALTER DATABASE tempdb MODIFY FILE (NAME = <temp3>, FILENAME
= <'E:\SQLData\tempdb_mssql_3.ndf'>);
GO
ALTER DATABASE tempdb MODIFY FILE (NAME = <temp4>, FILENAME
= <'E:\SQLData\tempdb_mssql_4.ndf'>);
GO
ALTER DATABASE tempdb MODIFY FILE (NAME = <temp5>, FILENAME
= <'E:\SQLData\tempdb_mssql_5.ndf'>);
GO
ALTER DATABASE tempdb MODIFY FILE (NAME = <temp6>, FILENAME
= <'E:\SQLData\tempdb_mssql_6.ndf'>);
GO
ALTER DATABASE tempdb MODIFY FILE (NAME = <temp7>, FILENAME
= <'E:\SQLData\tempdb_mssql_7.ndf'>);
```

```

GO
ALTER DATABASE tempdb MODIFY FILE (NAME = <temp8>, FILENAME
= <'E:\SQLData\tempdb_mssql_8.ndf'>);
GO
ALTER DATABASE tempdb MODIFY FILE (NAME = <templog>,
FILENAME = <'E:\SQLLog\templog.ldf'>);
GO

```

3. Restart the SQL Server service.

### Verifying Hyper-V deployment

After the deployment is complete, follow the steps in [Verifying the cluster deployment](#) to verify the deployment of Hyper-V virtualization.

## Deploying the Microsoft application agent for Data Domain Boost with SQL Server

The Microsoft application agent for Dell EMC Data Domain Boost (DD Boost) enables database and application administrators to efficiently back up and restore applications by using the applications' native tools.

Microsoft application agent backups are application-consistent, and the Microsoft application agent can completely restore the backups.

This section describes how to deploy the Microsoft application agent for DD Boost as follows:

1. Configure the Data Domain system.
2. Install and configure the Microsoft application agent on the SQL Server host.

### Configuring the Data Domain system

Configure the Data Domain system as follows.

---

**Note:** For more information, see the [Dell EMC Microsoft Application Agent Installation and Administration Guide](#) on Dell EMC Online Support.

---

1. Log in to the Data Domain system as the sysadmin user:

```

ssh sysadmin@172.16.191.250
Enter password: *****

```

2. Create a Data Domain file system:

```

filesystem create

```

3. Enable the file system:

```

filesystem enable

```

4. Create a DD Boost user and password:

```

User add ms_agent_user password m*****
ddboost set user-name ms_agent_user

```

5. Enable the DD Boost software:  

```
ddboost enable
```
6. Create a Data Domain storage unit:  

```
ddboost storage-unit create <MS_BACKUP> user <MS_AGENT_USER>
```
7. Enable application-optimized compression:  

```
mtree option set app-optimized-compression <sqlserver1>  
mtree </data/coll/ms_backup>
```
8. Enable distributed segment processing:  

```
ddboost option set distributed-segment-processing enabled
```

## Installing and configuring the Microsoft application agent on the SQL Server host

Install and configure the Microsoft application agent on the SQL Server host.

### Installing the Microsoft application agent

1. Download the Microsoft application agent 4.5.x software package from [Dell EMC Online Support](#), and copy or move the agent software to the database host.
2. Open and extract `msappagent45_win_x64.zip` with WinZip.
3. Start the installation wizard by double-clicking the `emcmsappagent.exe` file.
4. On the **Welcome** page, select **I agree to the license term and agreements**, and then click **Next**.
5. On the **Change Install Location** page, perform one of the following tasks, and then click **Next**:
  - In the **Folder name** field, leave the default installation location, which is **C:\Program Files\DPSAPPS\MSAPPAGENT**.
  - To specify a different installation location, click **Change**, specify the installation location, and then click **OK**.
6. On the **Configure Installation Options** page, select **SSMS Plug-in** to install the SQL server Management Studio plug-in, which you can use to perform SQL Server backup and restore operations.
7. Click **Install**.

### Configuring the application agent

Configure the Microsoft application agent for DD Boost with SQL Server.

---

**Note:** For more information, see the [Dell EMC Microsoft Application Agent Installation and Administration Guide](#) on Dell EMC Online Support.

---

1. Create the `C:\ddconfig.cfg` file with the following parameters:

```
DDBOOST_USER=<ms_agent_user>  
DEVICE_HOST=<l**.dell EMC.com>  
DEVICE_PATH=</ms_backup>  
LOCKBOX_PATH=<"C:\Program Files\DPSAPPS\common\lockbox">
```

```
CLIENT=<m****.appagentdev.com>
DEBUG_LEVEL=0
```

2. Create a lockbox—an encrypted file that the Microsoft application agent uses to store confidential data, such as login credentials, and protect that data from unauthorized access—by typing:

```
msagentadmin administration --createLB --lockbox
<C:\Lockboxes>
```

where `C:\Lockboxes` is the location of the lockbox.

If you do not specify a folder, the lockbox is created in the following default directory: `C:\Program Files\DPSAPPS\common\lockbox`

3. Create a configuration file that contains the required credentials.

For example, type the following command to edit `config.cfg`:

```
c:\Lockboxes\config.cfg
```

4. Add credentials to the configuration file.

For example, for a Data Domain user, add the following credentials:

```
LOCKBOX_PATH=<C:\Lockboxes>
DDBOOST_USER=<ms_agent_user>
DEVICE_HOST=<l***.dell EMC.com>
DEVICE_PATH=</ddbdatest/mattp/ms_backup>
```

5. Register the credentials with the lockbox:

```
msagentadmin administration --registerLB -config
"<config_file_path>"
```

For example:

```
msagentadmin.exe administration --registerLB --config
"C:\lockbox-config-details.cfg"
```

You are prompted for any required passwords for the configured user accounts.

## Next steps

After the deployment is complete, you can perform many post-deployment tasks such as adding and removing disks, adding and removing servers, and various other management tasks. For more instructions, see [Dell EMC Microsoft Storage Spaces Direct Ready Node: Deployment guide for scalable hyper-converged infrastructure with RX740xd and R640 Storage Spaces Direct Ready Nodes](#).