Implementing SQL Server 2016 with Microsoft Storage Spaces Direct on Dell EMC PowerEdge R730xd

Performance Study

Dell EMC Engineering
October 2017

A Dell EMC Performance Study

Revisions

Date	Description
October 2017	Initial release

Table of contents

Executive summary

Traditional data centers were built around discrete groups of infrastructure: storage, networking, and compute. These infrastructure products were purpose-built, with control and programmability embedded in the hardware. Traditional deployments have enjoyed a lot of success. However, they lack agility, scalability, ease of management, as well as cost advantage compared to hyper-converged infrastructure.

Hyper-converged infrastructure (HCI) based on software-defined technology is evolving rapidly to meet the requirements of modern data centers. Storage Spaces Direct (S2D) is a relatively new technology from Microsoft as part of the Windows Server 2016. It uses industry-standard servers with local-attached drives to create highly available, highly scalable software-defined storage.

To simplify the S2D deployment process for customers, Dell EMC has introduced the Microsoft Storage Spaces Direct Ready Node solution. With improved IO performance, S2D can now comfortably support mission critical workloads such as SQL Server. To aid IT professionals in their SQL implementation, the Dell EMC Global Solutions team has conducted OLTP and OLAP testing to evaluate the performance and scalability of the S2D Ready Node solution.

Our testing results demonstrate that the S2D Ready Node is able to scale storage and compute on demand while providing resilience at drive, server, and VM levels. A 4-node configuration including HA can deliver up to 10,850 transactions per second (TPS) and 20,861 queries per hour (QphH) making it a compelling software-defined infrastructure to deploy SQL Server workloads.

1 Introduction

Dell EMC Microsoft Storage Spaces Direct Ready Nodes (S2D RN) are pre-configured with certified components. At the base of their deployment is the versatile Dell EMC PowerEdge R730xd server. The architecture follows a modular approach treating each R730xd server as a building block. This helps to manage deployments in terms of scale and performance.

Windows Storage Spaces Direct is a software-defined storage offering, part of Windows Server 2016. S2D can have either hyper-converged or converged implementations. This study focuses on hyper-converged implementations.

1.1 Scope

This study examines the performance and scalability of SQL Server 2016 running in an S2D environment built upon PowerEdge R730xd.

1.2 Audience

This document is intended for IT professionals, database administrators, and consultants interested in deploying SQL Server 2016 with Storage Spaces Direct software-defined storage on PowerEdge R730xd servers.

2 Dell EMC Ready Node for S2D

Dell EMC offers Ready Nodes for S2D, which are tested and validated configurations, to provide the best possible performance. The Ready Node described in this study is a validated, pre-configured building block for a software-defined storage solution with S2D.

S2D RN consists of pre-configured R730xd servers. PowerEdge R730xd is a versatile server supporting a broad range of deployment options in terms of NICs, drives, and processor capabilities.

2.1 Ready Node configurations

Dell EMC supports three configurations for S2D Ready Node deployments—Hybrid Capacity, Hybrid Balanced, and All-Flash configurations. Due to the performance requirements for SQL Server, Hybrid Balanced and All-Flash configurations are preferred for such workloads. The details for the configurations are listed in Table 1.

Table 1 Ready Node configurations for S2D

Parameter	Hybrid Balanced	All Flash
Chassis configuration	R730xd (3.5” x 12)	R730xd (2.5” x 24)
Cache drive options	• 800 GB Intel S3710 WI SATA SSD • 1920 GB Samsung SM863a Mixed-Use SATA SSD	• 800 GB Samsung PM1725 Mixed-Use PCIe SSD • 1600 GB Samsung PM1725a Mixed-Use PCIe SSD
Capacity drive options	4/6/8 TB 3.5” SATA 6 Gbps HDD	1920 GB Samsung SM863a Mixed-Use SATA SSD

3 Solution overview

The RN configuration used in this study is part of the S2D Hybrid Balanced (HYB-32) configuration. The processor, memory, and storage were chosen to optimize performance. The hardware and software components used in the solution are shown in Figure 1.

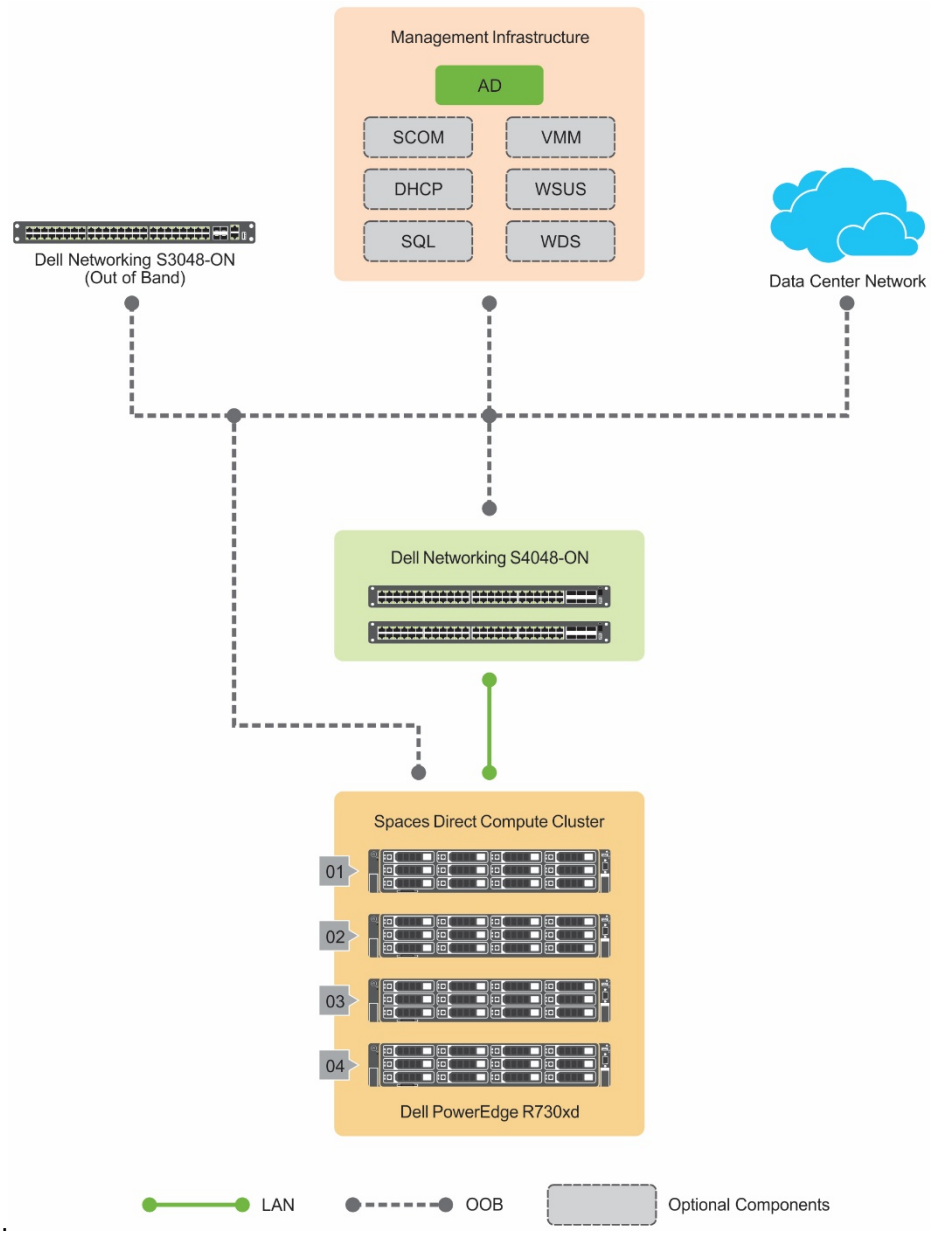

Figure 1 Solution overview

3.1 Configuration details

The hardware specifications used in this study are listed in Table 2.

Table 2 Solution component details

Component type	Details	Count
Server	Dell EMC PowerEdge R730xd server	4
Processor	Intel Xeon E5-2660 v4 2.0 GHz with 14 cores	8
Memory	512 GB (16 x 32 GB 2400 MT/s DIMM) – per server	NA
Storage controller	Internal HBA330	4
Storage – SSDs (cache)	800 GB Intel Write Intensive Solid State Drive, 2.5 inch, SATA	16
Storage – HDDs (capacity)	4 TB Toshiba 7.2K RPM, 6 Gbps, 3.5 inch, SATA	32
Network card	**NDC**: Intel i350 QP 1 GbE (default) **Add-On**: Mellanox Connect-X Pro 3 DP (10 GbE SFP+)	4+4
Network switch – backend	Dell EMC Networking S4048–ON	2
Network switch – management	Dell EMC Networking S3048–ON	2

3.2 PowerEdge R730xd

Dell EMC PowerEdge R730xd, a Dell EMC 13th generation server, is optimized for hyper-converged infrastructure. With two sockets and a wide range of processor options, R730xd also provides flexible computing capabilities based on customer needs.

This 2U rack mounted server provides high storage density with up to 16 x 3.5” drives: 12 x 3.5” (front bay) and 4 x 3.5” (internal bay). In addition, the 2 x 2.5” flex bay drive slots help achieve storage needs in HCI. PowerEdge R730xd can also be configured by using a 2.5’’ chassis form factor that supports up to 24 drives in the chassis.

The different HDD, SSD, and NVMe drive options available with PowerEdge R730xd server enable multiple hybrid and all-flash configurations for the Dell EMC Ready Nodes.

PowerEdge R730xd provides the storage density and compute power to maximize the benefits of S2D and make the best use of the advanced features in Windows Server 2016.

3.3 Networking 4048-ON

Dell EMC Networking S4048-ON is a 10 GbE, L2 and L3 capable network switch for storage, cluster, and client traffic. Some of the features supported by S4048-ON are:

- Routable RDMA over Converged Ethernet (RoCE)
- Converged network support for Data Center Bridging (DCB), with Priority Flow Control (802.1Qbb) and Enhanced Transmission Selection (802.1Qaz)

These features make S4048-ON the preferred choice for networking in an S2D environment.

3.4 Networking 3048-ON

Dell EMC Networking S3048-ON is a 1000 BASE-T, L2 and L3 capable switch that provides 48 ports supporting 10 MbE/100 MbE/1 GbE and four 10 GbE SFP+ uplinks. In the Ready Node configuration, S3048-ON is deployed to support the Out Of Band (OOB) connectivity between the PowerEdge R730xd servers and the Dell Networking S4048-ON switches.

4 Performance benchmarking results and analysis

This section outlines the performance results of PowerEdge R730xd with Storage Spaces Direct on SQL Server 2016. Tests were performed to evaluate OLTP and OLAP industry standard benchmarks, TPC-E and TPC-H, respectively.

4.1 OLTP test methodology and results

A 4-node R730xd server S2D Ready Node cluster was benchmarked. Within this cluster, three nodes were configured for performance, and one node was reserved for High Availability (HA). Each S2D node was configured with two VMs. Each VM supported a TPC-E schema database of 600 GB. Benchmarking of this S2D-RN cluster was conducted started with one node, then scaled to two and three.

Table 3 lists the configuration details of the benchmarking.

Table 3 OLTP test configuration details

Component	Component details per VM
Logical CPUs	26 Cores
Memory for OS	192 GB
Database Size	600 GB
TempDB Drives	50 GB VD
Database Drive	600 GB VD
Log Drive	100 GB VD
Virtual Network Interface	1
Test Workload	Quest Benchmark Factory TPC-E
Scale Factor	41

The results of the study are shown in Figure 2.

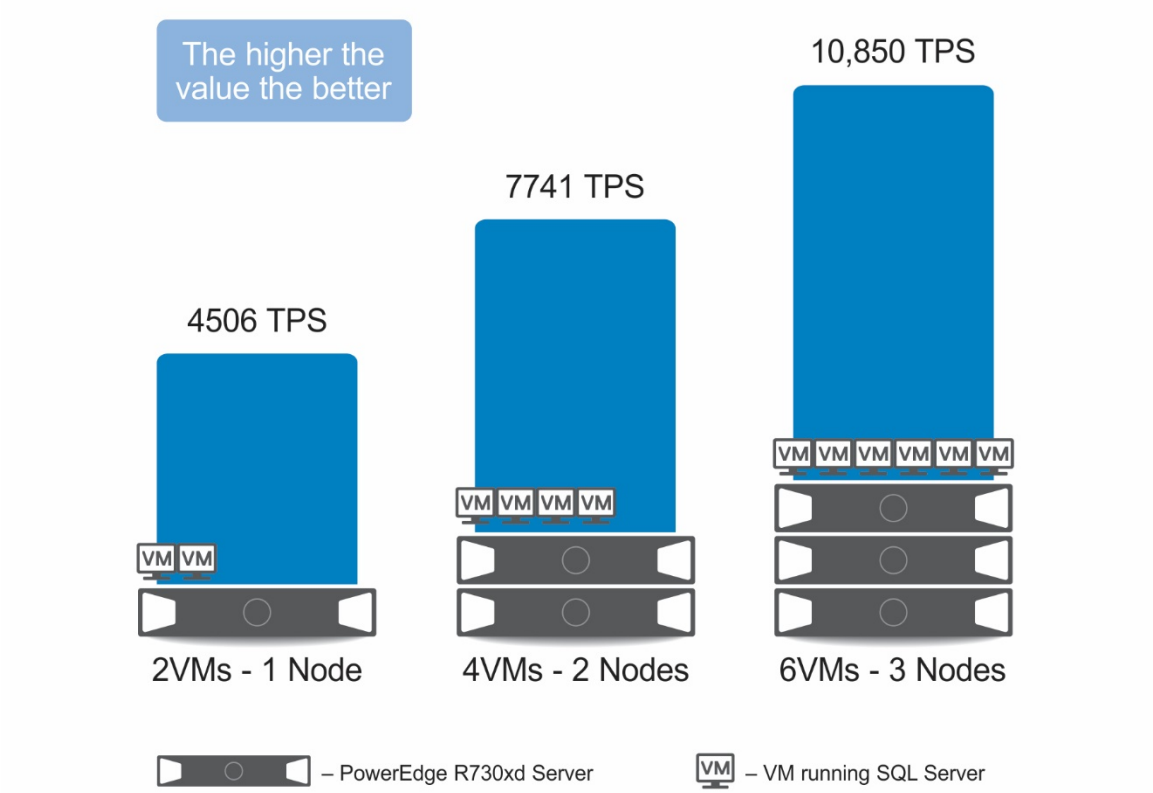

Figure 2 OLTP test results

4.2 OLAP test methodology and results

Online Analytical Processing (OLAP) provides enterprises with the capability to create actionable business intelligence. To demonstrate the OLAP capabilities of SQL Server 2016 on S2D infrastructure, the TPC-H benchmark was used in our testing.

TPC-H is a decision support benchmark. To measure performance, a number of queries are processed along with concurrent data modifications. The TPC-H benchmark processes queries and measures metrics using the Power Test and Throughput Test.

The TPC-H Power Test measures the query execution capability of a system when it is connected to a single user. The queries generated by a single user are referred to as a stream. This test simulates the queries submitted as a single stream, and measures the elapsed time to complete the same.

The TPC-H Throughput Test measures the highest number of queries that can be processed by a system submitted by multiple streams.

The primary metric used to measure the performance of a system under test for a TPC-H benchmark is the Composite Query per Hour or QphH. It is the geometric mean of the power and throughput measured for a system during the tests. The QphH that is measured is relative to its particular dataset. For example, in our testing, the dataset used is 500 GB, hence the QphH is listed as QphH@500 GB. The formula for calculating QphH is as follows:

$$\text{QphH@size} = \sqrt{Power@size \; x \; Throughput@size}$$

Table 4 lists the configuration details of the benchmarking.

Table 4 OLAP test configuration details

Component	Component details per VM
Logical CPUs	24 vCPUs (2 VMs per host)
	54 vCPUs (1 VM per host)
Memory for OS	160 GB
Database Size	500 GB
TempDB Drives	200 GB VD
Database Drive	4 x 300GB VD
Log Drive	200 GB VD
Virtual Network Interface	1
Test Workload	Quest Benchmark Factory TPC-H
Scale Factor	300

To demonstrate scaling, the 500 GB TPC-H dataset was replicated in VMs during the test. Figure 3 shows the VMs and the total QphH attained.

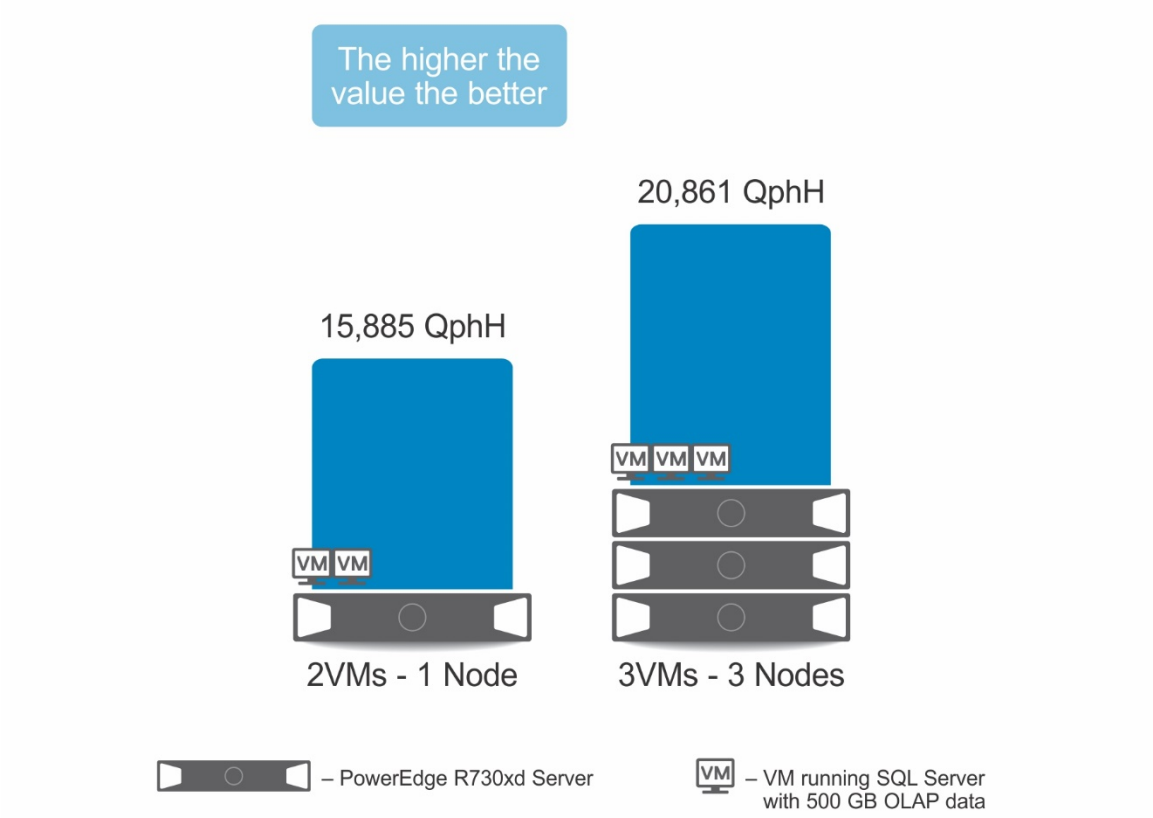

Figure 3 OLAP test results

5 Conclusion

Storage Spaces Direct (S2D) on Windows Server 2016 introduces many improvements that enable highly available and highly scalable software-defined storage while simplifying deployment and management. S2D provides new capabilities to pool locally attached SSDs, HDDs, and NVMe devices to deliver performance and scalability. The Dell EMC S2D RN provides customers customized and validated solutions. Dell EMC Storage Spaces Direct Ready Node (S2D RN) configurations enable customers to build high-performance SQL Server databases for different kinds of workloads. These include fast-paced, low-latency OLTP workloads, or OLAP workloads that can serve a high volume and variety of data.

As demonstrated by the results of this study, the Ready Node configuration used in this study is able to:

- Scale storage and compute on demand while providing resilience at drive, server, and VM levels.
- Deliver 10,850 TPS and 20,861 QphH making it a compelling software-defined infrastructure to implement SQL Server workloads.

6 References

See the following referenced or recommended resources related to this study.

- Design Guide for implementing Microsoft Storage Spaces Direct on PowerEdge R730xd with SQL Server 2016
- SQL Server 2016 Overview
- Hyper-converged solution using Storage Spaces Direct in Windows Server 2016
- Windows Server 2016 Overview