



**Hewlett Packard**  
Enterprise

# **HPE Reference Architecture for Microsoft SQL Server 2017 on Linux with HPE 3PAR StoreServ Storage using compression**

HPE 3PAR StoreServ 9450 Storage

# Contents

Executive summary.....	3
Solution overview.....	3
Solution components.....	4
Hardware.....	4
Software.....	7
Best practices and configuration guidance for the solution.....	8
Storage.....	8
SAN.....	9
Compute.....	9
Capacity and sizing.....	11
Workload description.....	11
Analysis and recommendations.....	12
Summary.....	13
Appendix A: Bill of materials.....	14
Resources and additional links.....	15

## Executive summary

This Reference Architecture provides best practices for the compression technology that is a key component of HPE Adaptive Data Reduction. Adaptive Data Reduction is a collection of technologies that come standard with HPE 3PAR StoreServ Storage which are designed to reduce your data footprint. This Reference Architecture features HPE 3PAR StoreServ Storage with Microsoft® SQL Server 2017 on Linux®. Hewlett Packard Enterprise engineers compared database workloads on both uncompressed and compression-enabled HPE 3PAR volumes. The test scenarios explore data warehouse (DW) and online transaction processing (OLTP) use cases. All testing was done using an all-flash storage array.

### HPE 3PAR StoreServ Storage compression advantages

- **Substantial cost savings:** HPE 3PAR compression is included with the HPE 3PAR StoreServ Storage license at no additional charge.
- **Ease of use:** HPE 3PAR compression is enabled on an HPE 3PAR virtual volume simply by setting compression to “Yes.” Compression can be enabled on volume creation, or it can be applied to a volume already containing data and the data will be compressed within a reasonably short time.
- **Improved compression ratio:** When enabled, HPE 3PAR StoreServ Storage can provide a greater than 2:1 compression ratio. (Compressibility of data will vary with each database environment.)
- **Saves disk space:** HPE 3PAR StoreServ Storage has the ability to compress an entire volume. It compresses all data—not just tables or tablespaces selected for SQL Server. Overall, you save disk space by using HPE 3PAR compression in a Microsoft SQL Server 2017 environment.

**Target audience:** This document is intended for IT administrators, storage administrators, and solution architects planning a Microsoft SQL Server 2017 deployment on Linux operating systems with HPE 3PAR StoreServ 9450 Storage. Readers should be familiar with the fundamentals of Microsoft SQL Server and HPE 3PAR StoreServ Storage best practices.

**Document purpose:** This document is for HPE customers seeking a compression-enabled storage solution for their Microsoft SQL OLTP and DW workloads—especially those customers who require high performance and scalability offered by the HPE 3PAR StoreServ 9450 Storage all-flash array.

This Reference Architecture describes solution testing performed in August 2017.

**Disclaimer:** Products sold prior to the separation of Hewlett-Packard Company into Hewlett Packard Enterprise Company and HP Inc. on November 1, 2015 may have a product name and model number that differ from current models.

## Solution overview

Microsoft's latest SQL Server 2017 is now available on Linux. Microsoft SQL Server 2017 is the first version of the Microsoft SQL Server database that runs on Linux, MacOS, and Windows Server®. Microsoft SQL Server 2017 provides customers industry-leading performance, real-time intelligence, and least vulnerability via encryption at rest and in motion.

Volume-level compression is now available in HPE 3PAR StoreServ 9450 Storage, running 3PAR OS v3.3.1. The solution described in this paper is built using these advancements on an industry-standard HPE ProLiant DL380 Gen9 Server running Red Hat® Enterprise Linux (RHEL) 7.3 and using a Fibre Channel-connected SAN.

This solution explores data warehouse (DW) and online transaction processing (OLTP) use cases on an HPE 3PAR all-flash array, and compares database workloads on both uncompressed and compression-enabled HPE 3PAR volumes.

Microsoft SQL Server 2017 on Linux is managed by the well-known SQL Server Management Studio, as well as a new `mssql-tools` package that can be installed on various releases of Linux.

Configuring and exporting a volume with compression enabled is very simple to accomplish via the HPE 3PAR StoreServ Management Console (SSMC) UI.

Figure 1 shows a diagram of the components involved. Regarding the SAN, the server has a single connection to each switch, and each of the four HPE 3PAR nodes has a single connection to each switch. SAN zoning is then performed on the HPE SN6000B Fibre Channel Switch. The SQL Server Management Studio and HPE 3PAR SSMC are hosted and installed on a Windows® VM in the data center. Only network connectivity between this VM, the RHEL 7.3 server, and the HPE 3PAR StoreServ 9450 Storage is needed. This same VM can also be used as the SSH client for accessing the HPE 3PAR StoreServ Storage and the RHEL machine for configuration.

## SQL Server 2017 on Linux with HPE 3PAR StoreServ 9450 Storage

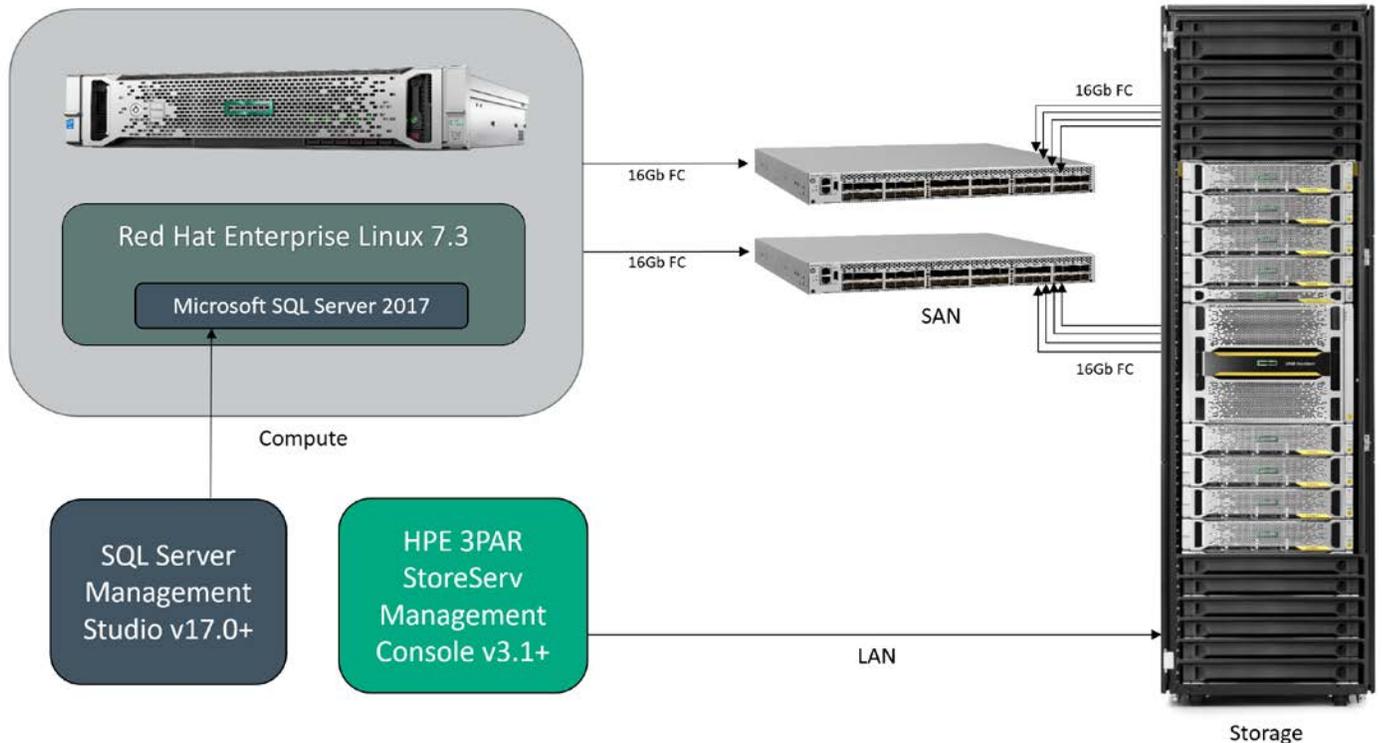


Figure 1. Solution overview

### Solution components

#### Hardware

##### HPE 3PAR StoreServ 9450 Storage

HPE 3PAR StoreServ 9000 Storage is an enterprise-class flash array that helps you consolidate primary storage workloads—for file, block, and object—without compromising performance, scalability, data services, or resiliency. This HPE 3PAR model based on the proven HPE 3PAR architecture is purpose built for all-flash consolidation, delivering the performance, simplicity, and agility needed to support your hybrid IT environment. HPE 3PAR StoreServ 9000 Storage is available in a single all-flash model—the 9450—that offers rich Tier-1 data services, quad-node resiliency, fine-grained Quality of Service (QoS), seamless data mobility between systems, high availability through a complete set of persistent technologies, and simple and efficient data protection with a flat backup to HPE StoreOnce Backup appliances.

Figure 2 shows the HPE 3PAR StoreServ 9000 Storage hardware. See the *HPE 3PAR StoreServ 9000 Storage QuickSpecs* for more details at <http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=a00005876enw>



**Figure 2.** HPE 3PAR StoreServ 9450 Storage Array with 16 Gb Fibre Channel modules

### **HPE ProLiant DL380 Gen9 Server**

The HPE ProLiant DL380 Generation 9 (Gen9) Server is designed to address the needs of large enterprises to remote office and branch office environments. Reliability, serviceability, and near-continuous availability, backed by a comprehensive warranty, make it ideal for any environment.

Key features:

- The HPE ProLiant DL380 Gen9 Server supports industry-standard Intel® Xeon® E5-2600 v3 and E5-2600 v4 processors with up to 22 cores, 12 Gb/s SAS, and 40 Gb NICs with a broad range of compute options, and leverages 2400 MHz DDR4 HPE Smart Memory supporting up to 3.0 TB.
- It has redundant HPE Flexible Slot power supplies, an HPE Flexible Slot Battery Backup module, and support for HPE Power Discovery Services offering.
- The redesigned chassis can accommodate 8 to 24 SFF drives or 4 to 12 LFF drives, also NVMe options, and a choice of embedded 4x1 GbE, HPE FlexibleLOM, or PCIe standup 1 GbE to 40 GbE adapters. In conjunction with the embedded SATA HPE Dynamic Smart Array B140i controller for boot, data, and media needs, the redesigned HPE Flexible Smart Array and HPE Smart SAS HBA controllers allow the flexibility to choose the optimal 12 Gb/s controller for your needs.
- The HPE ProLiant DL380 Gen9 Server can be managed and monitored by HPE iLO, HPE OneView, or HPE Smart Update Manager (SUM), with each tool providing a variety of management functions.

Figure 3 shows the HPE ProLiant DL380 Gen9 Server hardware. See the *HPE ProLiant DL380 Generation 9 QuickSpecs* for more details at <http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=c04375627>



**Figure 3.** HPE ProLiant DL380 Gen9 Server (front facing, no bezel, SFF drives)

### **Fibre Channel Switches – 16Gb SAN – HPE SN6000B Switches and HPE SN1000Q HBAs**

The HPE SN6000B Fibre Channel Switch meets the demands of hyper-scale, private cloud storage environments by delivering market-leading 16 Gb Fibre Channel technology and capabilities that support highly virtualized environments. Designed to enable maximum flexibility and investment protection, the HPE SN6000B Switch is configurable in 24, 36, or 48 ports and supports 4, 8, 10, or 16 Gbps speeds in an efficiently designed 1U package. It also provides a simplified deployment process and a point-and-click user interface making it both powerful and easy to use. The HPE SN6000B Switch offers low-cost access to industry-leading Storage Area Network (SAN) technology while providing "pay-as-you-grow" scalability to meet the needs of an evolving storage environment.

Key features and benefits:

- Delivers 16 Gbps performance with up to 48 ports in an energy-efficient, 1U form factor, providing maximum flexibility for diverse deployment and cooling strategies
- Features Ports on Demand (PoD) capabilities for fast, easy, and cost-effective scaling from 24 to 48 ports in 12-port increments
- Provides a flexible, simple, and easy-to-use SAN solution with industry-leading technology
- Supports highly virtualized, private cloud storage with multitenancy and non-stop operations
- Offers best-in-class port density and scalability for midrange-enterprise SAN switches, along with redundant, hot-pluggable components, and non-disruptive software upgrades
- Metro cloud connectivity features, including integrated DWDM and dark fiber support (optional license)
- In-flight compression and encryption included, providing efficient link utilization and security.
- Yields exceptional price/performance value, exceeding comparable Ethernet storage-based alternatives
- Unique bundle model: The HPE StoreFabric SN6000B 16Gb Bundled Fibre Channel Switch model facilitates ease of ordering by offering the HPE SN6000B 48/24 Fibre Channel Switch along with 24 16 Gb optical transceivers in one package. The HPE StoreFabric SN6000B 16Gb Bundled Fibre Channel Switch model is offered at a lower price than the sum of the list prices of the HPE SN6000B 48/24 Fibre Channel Switch and the 16Gb optical transceivers.

Figure 4 shows the rear view of the HPE SN6000B Fibre Channel Switch. For more details, see the *HPE SN6000B Fibre Channel Switch QuickSpecs* at <http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=c04111492>.



**Figure 4.** HPE SN6000B Fibre Channel Switch (rear view)

## Software

### Operating system

The following operating system was used for the solution.

#### Red Hat Enterprise Linux (RHEL) 7.3

RHEL 7.3 is the latest version of the world's leading enterprise Linux platform. New and enhanced for version 7.3:

- **Linux Containers:** Latest Docker engine and a new container-signing implementation
- **Security:** Updates to SELinux allowing for faster improved ease of use, support for MACSec v2 for encryption between OS and VMs
- **Identity management:** Supports smart card authentication with Active Directory
- **File systems and storage:** NVMe device support, additional parallel Network File System support (Block SCSI and Flex Files)
- **Networking:** High performance via fast data-path channels and stream-control transmission protocol
- **High Availability:** 'pacemaker' utility to manage multisite and stretch clusters with enhanced alerts via SNMP and SMTP
- **IoT:** Support for Bluetooth Low Energy, controller area network (CAN) bus, modem manager updates for cellular

For more details of RHEL 7.3, see [RHEL Technology Overview](#).

### Application software

The following application software was used for the solution.

#### Microsoft SQL Server 2017 for Linux

Microsoft SQL Server 2017 for Linux is supported on the following Linux platforms:

- Red Hat Enterprise Linux 7.3
- Ubuntu Linux 16.04
- SUSE Linux Enterprise Server v12 SP2

Features of Microsoft SQL Server 2017 for Linux:

- Observe up to 30x transaction speed with in-memory OLTP
- Enhanced security with encryption at rest and in motion
- Real-time intelligence: Gain transformative insights with up to 1 million predictions per second using built-in Python and R integration

### Management and support software

The following management and support software was used for the solution.

- [Microsoft SQL Server Management Studio \(SSMS\)](#)  
SSMS 17.2 was used for this paper; 17.0+ is required.
- [HPE 3PAR StoreServ Management Console \(SSMC\)](#) – Web-based management tool for the HPE 3PAR array  
SSMC v3.2 overall enhancements:
  - View alert notifications for all HPE 3PAR StoreServ Storage systems at any time.
  - Create a custom login banner for all users.
  - Customize filters by changing the object's attributes.
  - Customize the columns displayed in each object's table views.
  - Improved performance when managing multiple HPE 3PAR StoreServ Storage systems.
- SSH Terminal such as PuTTY or MobaXterm

### Best practices and configuration guidance for the solution

The environment can be separated into three main sections:

- Storage – HPE 3PAR StoreServ 9450 Storage
- SAN – Fibre Channel with HPE SN6000B Switches and HPE SN1000Q Server HBAs
- Compute – HPE ProLiant DL380 Gen9 Server running RHEL 7.3 and Microsoft SQL Server 2017 for Linux

### Storage

The HPE 3PAR StoreServ 9450 Storage must be running 3PAR OS 3.3.1, or later. This version is required for the compression feature that is used.

HPE 3PAR StoreServ Management Console (SSMC) 3.1 or later is also required. A separate physical or virtual machine (Windows or Linux) is needed to host this web-based management software.

This solution takes advantage of the automatically created common provisioning groups (CPGs).

Create a pair of HPE 3PAR Virtual Volumes – one for SQL Server 'data' and one for SQL Server 'log' storage, enabling compression. This is shown in Figure 5.

Figure 5. Creating a compression-enabled HPE 3PAR Virtual Volume

## Note

If your database includes in-memory tables for performance, HPE recommends creating the `log` volume on a performance-oriented CPG. The intention is to ensure the transaction log volume can keep up with the in-memory transaction activity. Though this testing did not observe a bottleneck with HPE 3PAR compression enabled, you might want to disable it when seeking the highest possible performance from the storage.

## SAN

Two HPE SN6000B SAN switches were used to provide redundant fabrics. Fibre cabling and zoning were done according to the [HPE 3PAR best practices guide](#), 'FC hosts zoning' section. Each HPE 3PAR node has a 16 Gb fibre connection to each of the two SAN zones for a total of 8 connections.

## Compute

Aspects of the HPE ProLiant DL380 Gen9 Server are:

- 128 GB of RAM, 4x 32 GB modules
- OS volume is RAID1 consisting of 2x 800 GB SSDs
- Fresh install of RHEL 7.3 Server
- Configured network and proxy to allow for yum updates (RHN) and SQL Server installation
- Firmware on SN1000Q Fibre Channel cards updated to 06.01.14 and the latest driver installed
- Utilized [HPE 3PAR Host Explorer 3.2.0](#) to import host/server to the HPE 3PAR StoreServ Storage array, [per documentation](#)
- DM-Multipath installed and configured (see [RHEL 7.3 multipath manual](#) for details)

The following steps are an example of mounting and configuring the volumes exported by the storage.

1. Mount and format volumes:

- a. After storage is exported from the array, rescan the bus to find the new LUNs:

```
# ls /sys/class/scsi_host
host0 host1 host2

# echo "- - -" > /sys/class/scsi_host/host0/scan
# echo "- - -" > /sys/class/scsi_host/host1/scan
# echo "- - -" > /sys/class/scsi_host/host2/scan
```

- b. Determine the multipath device for the volume. For example, WWN ending in A0007E1CC:

```
# multipath -ll | grep -i "A0007E1CC"
mpathi [360002ac0000000000000000a0007e1cc] dm-11 3PARdata,VV
```

- c. Format the volume, XFS in this case:

```
# mkfs.xfs /dev/mapper/mpathi
meta-data=/dev/mapper/mpathi      isize=512    agcount=32, agsize=4015196 blks
=                               sectsz=512   attr=2, projid32bit=1
=                               crc=1       finobt=0, sparse=0
data      =                       bsize=4096  blocks=157286272, imaxpct=25
=                               sunit=4     swidth=4096 blks
naming    =version 2               bsize=4096  ascii-ci=0 ftype=1
log       =internal log          bsize=4096  blocks=76800, version=2
=                               sectsz=512   sunit=4 blks, lazy-count=1
realtime  =none                  extsz=4096  blocks=0, rtextents=0
```

- d. Mount at desired location (that is, /mnt/log-compressed) by modifying /etc/fstab or otherwise.

2. Install and configure SQL Server 2017 for Linux ([main doc](#))

3. Configure the default log directory for SQL Server to point to the volumes exported by the HPE 3PAR StoreServ Storage, [per documentation](#):

```
# sudo chown mssql /mnt/log-compressed
# sudo chgrp mssql /mnt/log-compressed
# sudo /opt/mssql/bin/mssql-conf set filelocation.defaultlogdir /mnt/log-compressed
```

4. Configure the default data directory for SQL Server to point to the volumes exported by the HPE 3PAR StoreServ Storage, [per documentation](#):

```
# sudo chown mssql /mnt/data-compressed
# sudo chgrp mssql /mnt/data-compressed
# sudo /opt/mssql/bin/mssql-conf set filelocation.defaultdatadir /mnt/data-compressed
```

5. Restart SQL Server after any configuration changes:

```
# sudo systemctl restart mssql-server
```

6. If not already installed in your environment, install Microsoft SQL Server Management Studio (Linux and Windows versions available) – <https://docs.microsoft.com/en-us/sql/linux/sql-server-linux-setup-tools>

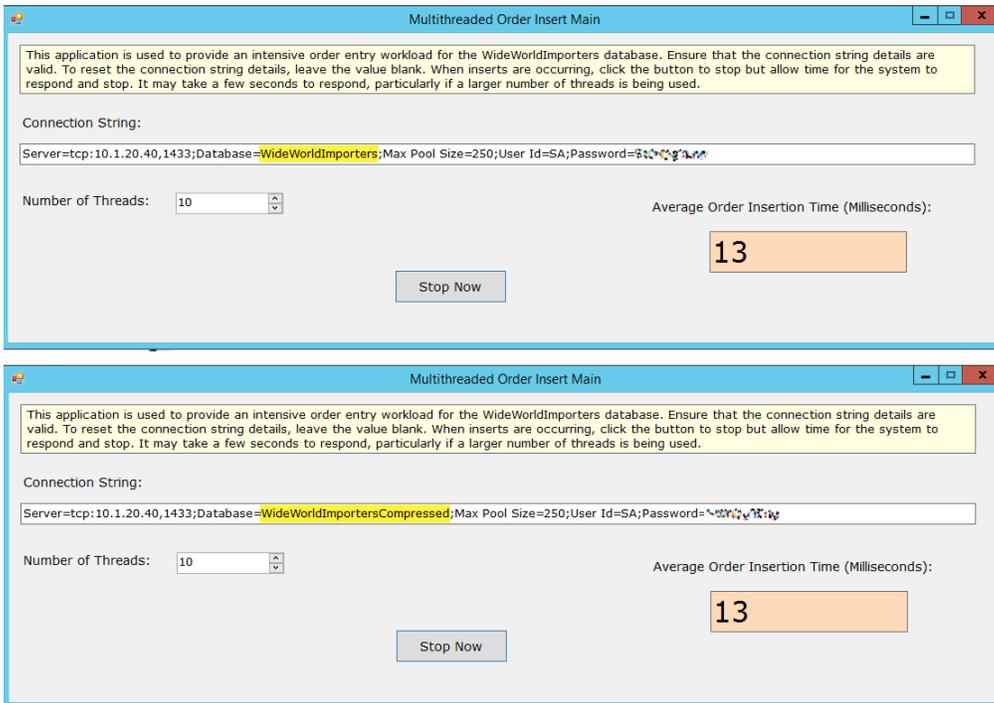
For documentation about using Windows SSMS with Linux, see <https://docs.microsoft.com/en-us/sql/linux/sql-server-linux-manage-ssms>.

## Capacity and sizing

HPE 3PAR StoreServ Storage volumes should be sized appropriately given your database application needs. Though this solution did not focus on in-memory tables, it is important to note that transaction speed can be dramatically increased (up to 40 times) with in-memory tables and natively compiled stored procedures. Therefore, make sure the transaction log volume is able to keep up. When seeking the highest possible performance, the administrator can decide to not use an HPE 3PAR compressed volume for the transaction log.

## Workload description

The Microsoft WideWorldImporters database, associated sample scripts, and workload drivers provided the workload. This is a sample database created by Microsoft for evaluating Microsoft SQL Server 2016 and later. In particular, the 'Multi-threaded Order Insert' executable was used on databases stored on compressed versus standard volumes. In Figure 6, the WideWorldImporters database had data and log contents stored on separate volumes. The 'WideWorldImportersCompressed' database had data and log contents stored on HPE 3PAR compression-enabled volumes.



**Figure 6.** Multithreaded-insert workload without/with HPE 3PAR compression

The example databases, WideWorldImporters and WideWorldImportersDW, and executable workload drivers are available on GitHub for testing in your environment. See the WideWorldImporters GitHub repository for source.

Figure 7 shows the table layout of the example databases and initial contents of the Sales.orders table.

The screenshot displays the Microsoft SQL Server Management Studio (SSMS) interface. On the left, the Object Explorer shows the database structure for 'WideWorldImporters' and 'WideWorldImportersDW'. The central pane shows a SQL query window with the following query:

```

/***** Script for SelectTopNRows command from SSMS *****/
SELECT TOP (1000) [OrderID]
, [CustomerID]
, [SalespersonPersonID]
, [PickedByPersonID]
, [ContactPersonID]
, [BackorderOrderID]
, [OrderDate]
, [ExpectedDeliveryDate]
, [CustomerPurchaseOrderNumber]
, [IsUnderSupplyBackordered]
, [Comments]
, [DeliveryInstructions]
, [InternalComments]
, [PickingCompletedWhen]
, [LastEditedBy]
, [LastEditedWhen]
FROM [WideWorldImporters].[Sales].[Orders]

```

Below the query, the Results pane shows the first six rows of data from the Sales.Orders table:

OrderID	CustomerID	SalespersonPersonID	PickedByPersonID	ContactPersonID	BackorderOrderID	OrderDate	ExpectedDeliveryDate	CustomerPurchaseC
1	832	2	NULL	3032	45	2013-01-01	2013-01-02	12126
2	803	8	NULL	3003	46	2013-01-01	2013-01-02	15342
3	105	7	NULL	1209	47	2013-01-01	2013-01-02	12211
4	57	16	3	1113	NULL	2013-01-01	2013-01-02	17129
5	905	3	NULL	3105	48	2013-01-01	2013-01-02	10269
6	976	13	3	3176	NULL	2013-01-01	2013-01-02	13383

On the right, the Properties window shows connection parameters for the current connection, including connection name, connection details, and server version (14.0.900.75).

Figure 7. SSMS view of Microsoft-provided example OLTP and DW databases.

## Analysis and recommendations

Key results from the OLTP workload and DW database testing are:

- SQL Server 'Transaction Log' volumes were observed to exhibit between 2.7:1 and 3:1 compression.
- SQL Server 'Data' volumes were observed to exhibit between 2.1:1 and 2.4:1 compression.
- SQL transaction speed was the same, whether or not the database was stored on standard **or** compressed HPE 3PAR volumes.

Similar compression ratios were observed when a standard volume was converted to a compressed volume, making volume compression applicable for existing environments as well. This operation can be performed via the HPE SSMC UI or via the HPE 3PAR CLI with the `tunevv` command.

---

### Note

If desired, the volume can be converted back to uncompressed. Both of these conversions are online operations.

---

Because no transaction performance impact was noticed on any of the workloads, HPE recommends enabling HPE 3PAR compression on the volumes containing SQL Server 'log' and 'data' content. Transactional database workloads are good candidates for taking advantage of compression. The customer also benefits from volume-based compression as the entire data and log contents are candidates for compression—as opposed to specifying compression on a per-table basis.

If multiple in-memory tables are present in the environment, further evaluation is required to ensure there is not a bottleneck at the storage array or SAN. If more array performance is deemed necessary, disable compression on the volume hosting the transaction log.

Deduplication, a sibling feature to compression in the HPE Adaptive Data Reduction suite, is also easily enabled or disabled. Typically, deduplication is not applied to transactional databases. However, certain situations may warrant closer inspection; one example being a DevOps environment with multiple Docker instances of the same SQL Server for Linux environment.

PolyBase support was not available at time of writing (SQL Server 2017 for Linux version RC2, v14.0.900.75) but is scheduled for a later release.

## Summary

This solution explored the versatility of HPE 3PAR StoreServ 9450 Storage and Microsoft SQL Server 2017 on Linux. The HPE 3PAR StoreServ Storage compression feature is easily enabled or disabled on a per-virtual-volume basis. Toggling the volume compression feature is an online operation making future adjustments possible.

This solution utilized the example WideWorldImporters and WideWorldImportersDW databases and associated workloads provided in Microsoft's GitHub repository. Both OLTP and DW scenarios with compression functionality on the HPE 3PAR StoreServ 9450 Storage have provided at least a 2:1 reduction for the SQL Server data volume and up to a 3:1 compression for the Microsoft SQL Server transaction log volume. Transaction speed was unaffected by HPE 3PAR volume compression.

Compression should not be enabled on volumes requiring the utmost performance—such as the volume holding the transaction log for in-memory tables. On-site evaluation with an applicable workload would need to be performed to make this determination.

## Appendix A: Bill of materials

### Note

Part numbers are at time of testing and subject to change. The bill of materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult with your HPE Reseller or HPE Sales Representative. (For more detail, see: [hpe.com/us/en/services/consulting.html](http://hpe.com/us/en/services/consulting.html).)

**Table 1.** Bill of materials

Qty	Part number	Description
<b>HPE ProLiant Server</b>		
1	719064-B21	HPE DL380 Gen9 8-SFF CTO Server
1	719051-L21	HPE DL380 Gen9 Intel Xeon E5-2620v3 (2.4GHz/6-core/15MB/85W) FIO Processor Kit
1	719051-B21	HPE DL380 Gen9 Intel Xeon E5-2620v3 (2.4GHz/6-core/15MB/85W) Processor Kit
4	726722-B21	HPE 32GB 4Rx4 PC4-2133P-L Kit
2	872359-B21	HPE 800GB SATA 6G WI SFF SC DS SSD
1	QW972A	HPE StoreFabric SN1000Q 16Gb 2-port PCIe Fibre Channel Host Bus Adapter
1	719073-B21	HPE DL380 G9 3Slot PCIe Riser Kit
2	720478-B21	HPE 500W Flex Slot Platinum Hot Plug Power Supply Kit
1	733660-B21	HPE 2U SFF Easy Install Rail Kit
<b>HPE 3PAR StoreServ 9450 Storage</b>		
1	QOE92A	HPE 3PAR 9450 2N+SW Storage Base
2	Q7F41A	HPE 3PAR 9450+SW Storage Node
4	QOE97A	HPE 3PAR 9000 4pt 16Gb FC HBA
8	QOE95A	HPE 3PAR 9000 2U SFF Drive Enclosure
48	QOF40A	HPE 3PAR 9000 400GB+SW SFF SSD
16	716197-B21	HPE Ext 2.0m MiniSAS HD to MiniSAS HD Cbl
1	P9M30A	HPE 3PAR Direct Connect Cabling Option
1	QOF86A	HPE 3PAR StoreServ RPS Service Processor
1	L7F04A	HPE 3PAR 9450 All-inclusive Multi-system Software LTU
<b>HPE SAN</b>		
2	QR481B	HPE SN6000B Fibre Channel Switch
12	QK724A	HPE B-series 16Gb SFP+ Short Wave Transceiver
10	AJ835A	2M LC/LC FC cable
<b>Software</b>		
1		Red Hat Enterprise Linux 7.3
1		Microsoft SQL Server 2017 for Linux

## Resources and additional links

HPE 3PAR StoreServ Architecture

<http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA3-3516ENW>

HPE Adaptive Data Reduction

<http://h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA4-9573ENW>

HPE Adaptive Data Reduction Chalk Talk

*ADR: Zero-detect, Deduplication, Compression and Data Packing*

<youtube.com/watch?v=SisUMy26q5U>

Red Hat Enterprise Linux 7.3 Release Notes

[https://access.redhat.com/documentation/en-US/Red\\_Hat\\_Enterprise\\_Linux/7/html/7.3\\_Release\\_Notes/index.html](https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/7.3_Release_Notes/index.html)

RHEL 7.3 Technology Overview

<redhat.com/cms/managed-files/li-rhel-73-whats-new-technology-overview-inc0464295lw-201611-en.pdf>

Microsoft SQL Server 2017

<microsoft.com/en-us/sql-server/sql-server-2017>

Microsoft SQL Server 2017 for Linux Release Notes

<https://docs.microsoft.com/en-us/sql/linux/sql-server-linux-release-notes>

Microsoft Wide World Importers Sample Databases v1.0

<https://github.com/Microsoft/sql-server-samples/releases/tag/wide-world-importers-v1.0>

HPE Reference Architectures

<hpe.com/info/ra>

HPE Servers

<hpe.com/servers>

HPE Storage

<hpe.com/storage>

HPE Networking

<hpe.com/networking>

HPE Technology Consulting Services

<hpe.com/us/en/services/consulting.html>

To help us improve our documents, please provide feedback at <hpe.com/contact/feedback>.



**Sign up for updates**



---

© Copyright 2017–2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows Server, and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries.

a00026670enw, April 2018, Rev. 1