

Microsoft SQL Server with Dell EMC and VMware Technology

Introducing SQL Server 2019 and Dell EMC Unity XT Midrange All-Flash Storage

September 2019

H17911

White Paper

Abstract

This white paper introduces SQL Server 2019 and Unity XT all-flash storage. It includes best practices for virtualizing SQL Server with VMware technology and configuring and managing the underlying Dell EMC infrastructure components.

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 © 2019 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 09/19 White Paper H17911.

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

Contents

- Executive summary.....4
- New features in SQL Server 2019.....6
- Dell EMC Unity XT midrange storage.....8
- VMware virtualization.....18
- Summary.....22
- References.....23

Executive summary

SQL Server evolution

In 2017, Dell EMC and VMware released the results of a survey about trends in SQL Server data environments. The report, [SQL Server Transformation: Toward Agility and Resiliency](#), drew on the experience of the Professional Association of SQL Server (PASS) membership community. The findings show that SQL Server database environments are growing in both size and complexity, fueled by rising data volumes and new business demands. SQL Server databases now sit at the heart of many existing enterprises, powering mission-critical applications, and are being positioned to power new “digital-native” businesses.

In the two years since our survey was conducted, Microsoft has been working to develop the next generation of the SQL Server data platform—SQL Server 2019. In addition to continuing the 30-year-old trend of constant improvement of the core relational engine features and data warehousing, Microsoft is adding new services and features. For example, SQL Server 2019 includes support for Big Data workloads with Apache Spark and Hadoop Distributed File System (HDFS).

SQL Server 2019 with Dell EMC infrastructure

Dell EMC and Microsoft have been partnering on solutions for SQL Server for decades. Successful implementation of a comprehensive database platform such as Microsoft SQL Server requires design coordination between the implemented software features and the underlying infrastructure. That infrastructure includes CPU processing power, memory resources, storage design, and networking service. For any type of application workload, Dell EMC is a single source for the essential infrastructure of your SQL Server platform.

The Dell EMC PowerEdge server portfolio offers many combinations of CPU processor and memory configurations. The configurations are suitable for everything from small departmental applications to the largest mission-critical systems like enterprise resource planning (ERP), data warehousing, advanced analytics, eCommerce, and the full spectrum of custom applications. Our storage portfolio provides high-capacity, high-performance, and reliable storage for both unstructured and structured data. Our comprehensive storage product line complements the innovative SQL Server transition into a complete data hub.

Customers who deploy SQL Server 2019 with Dell EMC infrastructure can deliver transformational insights from both structured and unstructured data with the combined power of SQL Server and Apache Spark. SQL Server also requires a mixture of technologies for client access, server-to-server networking, and server-to-storage networking. Dell EMC’s vision for the network is based on a disaggregated model offering an open ecosystem in which organizations can select from a wide range of innovative, industry-standard network applications, operating systems, and hardware platforms. This approach gives customers maximum control over the technologies that they choose and the architectures that they adopt, resulting in measurable cost savings and increased service agility.

SQL Server and VMware virtualization

VMware provides virtualization for all the critical infrastructure components that SQL Server requires for high performance and consistent operations. In addition to providing low-cost and consistent private cloud services for workloads, VMware is also now pioneering hybrid cloud operational models that span private and public cloud architectures. Many organizations have adopted a “virtualize first” standard to lower infrastructure costs, standardize skill sets, provide consistently high availability, and simplify disaster recovery.

As reported in [SQL Server Transformation: Toward Agility and Resiliency](#), 94 percent of SQL Server professionals who responded to our survey reported having some level of virtualization in their environment. Seventy percent of those using virtualization reported that VMware was their primary virtualization technology provider. More than half of survey respondents—60 percent—reported rates of SQL Server virtualization of 75 percent or more. Further, survey results provide strong evidence that having high availability and disaster recovery implemented in the virtualization layer were important factors in the decision to virtualize SQL Server databases.

Document purpose

This paper provides an overview of SQL Server 2019 and Dell EMC infrastructure in a VMware virtualized environment. It introduces key features of SQL Server 2019 and the newest Dell EMC series of midrange all-flash storage arrays—Dell EMC Unity XT. It includes best practices for configuring and managing the infrastructure components and virtualizing SQL Server.

Audience

This white paper is for SQL Server administrators, IT architects, and others who are interested in learning about the benefits of deploying SQL Server 2019 in a VMware virtualized environment with Dell EMC infrastructure components.

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 by [email](#) or provide your comments by completing our [documentation survey](#).

Authors: Phil Hummel, Anil Papisetty

Note: For links to additional documentation about Dell EMC solutions for SQL Server, see [Microsoft SQL Info Hub for Ready Solutions](#).

New features in SQL Server 2019

The SQL Server 2019 database platform includes a broad range of technologies, features, and services, supporting mission-critical applications such as analytics, in-memory databases, business intelligence (BI), and reliable and scalable online transaction processing (OLTP). The SQL Server platform has acquired capabilities to handle data integration, data warehousing, reporting, high-speed advanced analytics, data replication and programmability features including hosting in-database common language runtimes, service broker hosting, and semi-structured datatype management. Microsoft realizes that with such a breadth of services available in the platform, not all customers or applications need every feature enabled on every SQL Server instance. In many cases, it is preferable to separate SQL Server services either through virtualization or by managing service-specific implementation.

Key new features and updates in SQL Server 2019 include:

- SQL Server on Linux
 - Supports participation in transactional, merge, and snapshot replication topologies in the publisher, distributor, or subscriber roles
 - Is configurable for user authentication through Microsoft Active Directory and for replication and distributed queries
 - Supports participation in Availability Groups that are authenticated by Active Directory, in addition to the previously supported certificate-based authentication
 - Can initiate and participate in distributed transactions through the Linux version of the Microsoft Distributed Transaction Coordinator (MSDTC), which can also participate in distributed transactions, including developer support, with other third-party transaction coordinators
- Big Data clusters (new feature)
 - Support Spark Linux containers on Kubernetes
 - Integrate with data stored on HDFS
 - Enable advanced analytics and machine learning on Spark
 - Use Spark streaming to ingest data into persistent SQL Server storage—SQL data pools
 - Use query books that provide a notebook experience in Microsoft Azure Data Studio
 - Use Sparklyr R interface
 - Use Jupyter notebooks to assist with deployment and discovery, diagnosis, and troubleshooting for components in a SQL Server Big Data cluster
- Always-On Availability Groups
 - Support up to five synchronous replica pairs—one primary and up to four secondary replicas—with automatic failover between replicas
 - Enable high-availability configurations for SQL Server running in containers, using Kubernetes as an orchestration layer

- Azure Data Studio
 - Provides a lightweight, open source, cross-platform desktop tool for the most common tasks in data development and administration
 - Connects to SQL Server on premises and in the cloud from Windows, macOS, and Linux
- Polybase
 - Supports using external table column names for querying SQL Server, Oracle, Teradata, MongoDB, and ODBC data sources
- Machine Learning Services
 - Enables running R and Python scripts on Linux
 - Can be installed on a Windows Server failover cluster for high availability
 - Uses existing R or Python scripts without modification to process data at the table partition level to train a model for each table partition and parallelize model training per partition
- Power BI Report Server
 - Enables scaling to thousands of users
 - Enables reports that are designed in Power BI Desktop to be deployed on an on-premises server rather than from the Power BI cloud service

For more details about Microsoft SQL Server 2019, see the [SQL Server 2019 CTP announcement archive](#) and the [Microsoft SQL Server 2019 technical white paper](#).

Dell EMC Unity XT midrange storage

Unity XT overview

Dell EMC Unity XT midrange platforms are shared storage solutions that provide low latency, high throughput, and low management overhead for SQL Server workloads. All Unity XT systems employ a dual-active architecture that uses dual storage processors (SPs) to serve host I/O and run data operations in an active/active manner. Unity XT dual SPs use full 12 Gb SAS back-end connectivity and a patented multicore architected operating environment to deliver high performance and efficiency. Disk array enclosures provide additional storage capacity.

The Unity XT family includes Unity XT All-Flash, Unity XT Hybrid, UnityVSA, and Unity Cloud Edition.

Unity XT Hybrid and All-Flash unified storage systems

Intel processor-based Unity XT Hybrid and Unity XT All-Flash storage systems implement an integrated architecture for block, file, and VMware VVols, with concurrent support for native network-attached storage (NAS), iSCSI, and Fibre Channel (FC) protocols. Unity XT Hybrid and Unity XT All-Flash platforms are NVMe-ready.

Unity XT Hybrid systems are designed for multicloud environments. They provide:

- Scalability up to 16 PB raw capacity
- Inline data reduction for all flash pools
- Fast installation and configuration (average 25 minutes)

Unity XT All-Flash systems deliver speed, efficiency, and multicloud support, providing:

- Up to two times more performance¹
- Up to 7:1 inline data reduction²
- Fast installation and configuration (less than 30 minutes)

[Unity XT All-Flash unified storage for SQL Server](#) on page 9 provides more details about Unity XT All-Flash systems.

UnityVSA

The UnityVSA system is software-defined virtual storage for VMware ESXi environments, using server-based, shared, or cloud storage capacity. UnityVSA HA, a dual-storage-processor deployment of UnityVSA, provides additional resiliency against disaster. UnityVSA storage provides:

- Up to 50 TB of fully functional unified storage
- Compatibility with Unity XT systems and features
- Support for high availability (UnityVSA HA)

¹ Based on Dell EMC internal testing in March 2019 of Unity XT 880F versus Unity 650F with an 8K block size on 100 percent reads, 100 percent writes, and 70:30 mixed workload.

² [Dell EMC Unity XT Rises to the Top in Independent Storage Performance and Efficiency Testing](#) (Aug 2019).

- NAS and iSCSI connectivity
- Data replication to and from other Unity XT platforms

Unity Cloud Edition

For file synchronization and disaster recovery operations in the cloud, the Unity XT family includes Unity Cloud Edition, which provides:

- Full-featured software-defined storage capabilities deployed in the cloud
- Easy deployment of block and file storage with VMware Cloud on AWS
- Support for disaster recovery and as-needed operational demands including test/dev and data analytics

Unity XT All-Flash unified storage for SQL Server

In the 2017 Unisphere Research report [SQL Server Transformation: Toward Agility and Resiliency](#), respondents reported their use of flash storage technology as follows:

- Twenty-two percent had all-flash storage arrays in production (16 percent) or were planning for them (6 percent).
- Thirty percent had hybrid arrays that included flash storage.
- Thirteen percent used direct-attached flash technology.
- Thirteen percent backed up SQL Server databases to flash storage arrays.

This rapid adoption of all-flash storage for use with SQL Server means that the Unity XT All-Flash arrays are particularly well suited for SQL Server developers and administrators. Unity XT All-Flash systems provide SQL Server developers and administrators with features and performance that go beyond what typical storage area networks (SANs) offer.

The Unity XT All-Flash systems, which are NVMe-ready to deliver high performance and low latency, have an efficient 2U form factor with a modern architecture. They are designed for all-flash performance that supports dual-socket CPUs, dual-active controllers, and multicore optimization. The following table provides an overview of each of the four Unity XT All-Flash models.

Table 1. Unity XT All-Flash models

Unity XT model	Processors (per SP)	Memory (GB per SP)	Maximum number of drives	Maximum raw capacity (PB)
380F	1 Intel E5-2603 v4 6c/1.7 GHz	64	500	2.4
480F	2 Intel Xeon Silver 4108 8c/1.8 GHz	96	750	4.0
680F	2 Intel Xeon Silver 4116 12c/2.1 GHz	192	1,000	8.0
880F	2 Intel Xeon Gold 6130 16c/2.1 GHz	384	1,500	16.0

For more information about the full line of Unity XT Series All-Flash systems, see the [Dell EMC Unity XT Storage Series Specification Sheet](#).

Storage pools

Many SQL Server professionals are aware that all modern storage arrays provide the ability to group disks into larger units of storage with a fixed level of RAID protection. We refer to discrete groups of drives with RAID protection as traditional pools. While Unity XT Hybrid systems support only traditional pools, Unity XT All-Flash arrays also support dynamic storage pools. With dynamic storage pools, RAID protection is applied to drive extents, which are storage units that are smaller than a full disk. Dynamic pools enable greater flexibility in managing and expanding the disk pools.

[Dell EMC Unity Storage with Microsoft SQL Server](#) provides best practices for managing storage pools to achieve the best performance with the least complexity. For example, Dell EMC recommends minimizing the number of Unity XT storage pools to reduce complexity and increase flexibility. However, configuring additional storage pools might be appropriate in some circumstances, including those where you need to:

- Separate workloads with different I/O profiles.
- Dedicate resources to meet specific performance goals.
- Separate resources for multitenancy.
- Create smaller failure domains.

Storage volumes (LUNs)

The concept of thin provisioning (just-in-time allocation) has become ubiquitous in the virtualization and storage technology arena. Although Unity XT arrays also support thick provisioning, few reasons remain for using that technology, especially on all-flash storage where new allocations can be performed quickly. By default, the Unity XT system creates thin storage objects, which are virtually provisioned and space efficient. For a more detailed discussion of storage object types, see the [Dell EMC Unity: Best Practices Guide](#).

It is more difficult to determine the best tradeoff between management and flexibility when deciding how many volumes to create on an array. A best practice for Unity with SQL Server is to create a volume for each user database file for ultimate flexibility. In practice, most organizations adopt a multitier approach where the most mission-critical databases get maximum flexibility and less critical databases group files on fewer, larger volumes. We recommend that you examine all the requirements for the databases and any related applications because many data protection and monitoring technologies depend on file placement isolation.

Many volumes can be difficult to manage, especially in virtualized environments. Virtualized SQL Server environments are a good example of where placing multiple file types on a single volume might make sense. The DBA or storage administrator, or both, should find the right balance between flexibility and maintainability when determining the number of volumes to create.

File storage

NAS servers host file systems on the Unity XT storage system. File systems can be accessed via SMB or NFS protocols and can be shared to both protocols simultaneously by using a multiprotocol file system. NAS servers use virtual interfaces to enable host

connectivity to SMB, NFS, and multiprotocol file systems, as well as VMware NFS datastores and VMware Virtual Volumes. File systems and virtual interfaces are isolated to a single NAS server, allowing for multitenancy over multiple NAS servers. NAS servers are hosted on a storage processor and automatically fail over if the storage processor becomes faulted. Any associated file systems fail over as well.

SQL Server 2012 (11.x) and later versions support the Server Message Block (SMB) 3.0 protocol, enabling network file sharing for storage. For both stand-alone and failover cluster installations, you can install system databases (master, model, msdb, and tempdb) and Database Engine user databases with the SMB storage option. Using SMB storage is a good option when using Always On Availability Groups because the file share witnesses require access to a highly available network-accessible resource.

Creating SMB file shares for SQL Server deployment with Unity XT storage is a simple, three-step process: Create a NAS server, create the file system, and create the SMB share. Dell EMC Unisphere storage management provides setup wizards that guide you through the process. However, there are some important considerations to keep in mind when hosting SQL Server workloads on SMB file shares that do not necessarily apply to using SMB file shares as backup files. Microsoft has compiled a list of installation and security considerations together with currently known issues; for details, see [Install SQL Server with SMB fileshare storage](#) in Microsoft SQL Docs.

Storage snapshots

Data has become a critical strategic business initiative, and today's mission-critical environments demand more than redundancy. They require that applications always be online, with nondisruptive operations and upgrades. They also require high performance and the increased data availability that local snapshot replication and remote replication provide.

The Unity XT storage array offers snapshot capabilities for block and file storage resources that use common workflows, operations, and architecture. The Unity snapshot methodology provides a simple and effective way to protect data. Snapshots facilitate data recovery by enabling rollback to an earlier snapshot, or you can copy select data from the snapshot to protect against a storage system outage. The following table shows snapshot retention timeframes for Unity XT systems.

Table 2. Local and remote snapshot retention

Snapshot type	CLI		UI		REST	
	Manual	Scheduled	Manual	Scheduled	Manual	Scheduled
Local	1 year	1 year	5 years	4 weeks	100 years	No limit
Remote	5 years	255 weeks	5 years	255 weeks	5 years	255 weeks

Snapshots are not a direct replacement for other data protection methods such as native backup. Snapshots can augment traditional backup as a first line of defense for extremely low RTO scenarios.

The Dell EMC Unity snapshots feature incorporates data reduction and advanced deduplication. Snapshots also benefit from the space savings that are achieved on the source storage resource. When you take a snapshot of a data-reduction-enabled storage resource, the data on the source can be compressed or deduplicated.

Snapshot considerations

Important considerations to ensure a successful database recovery when using snapshots with SQL Server databases include:

- All components of a SQL Server database must be protected as a set. When data and log files are located on separate LUNs, the LUNs must be part of a consistency group. The consistency group ensures that the snapshot is taken at the same time on all LUNs in the group. When data and log files are located on multiple SMB file shares, the shares must reside on the same file system.
- When performing a SQL Server database recovery from a block-based snapshot, if the SQL Server instance is to remain online, use the Unisphere attach-to-host action. For file-based recovery, an additional SMB share is created using the snapshot as the source. After the volumes are attached, the database can be attached under a different name or you can replace the existing database with the recovered database.
- When performing a recovery using the Snapshot Restore method within Unisphere storage management, take the SQL Server instance offline. SQL Server is unaware of the underlying recovery operations. Taking the instance offline ensures that the volumes are not corrupted by database write activity before doing a recovery. Once the instance is restarted, SQL Server crash recovery will run and bring the databases to a consistent state.
- Enable snapshots on a few storage objects at a time, and then monitor the system to be sure it is still within the recommended operating ranges before enabling more snapshots.

Snapshot automation and scheduling

You can automate snapshots on Unity XT systems. The following default snapshot settings are available in Unisphere storage management: Default Protection, Protection with Shorter Retention, and Protection with Longer Retention. Each setting takes daily snapshots and retains them for various lengths of time.

You can choose one or both of two scheduling options—every x number of hours (between 1 and 24) and daily/weekly. Daily/weekly snapshot scheduling enables you to specify certain times on specific days for snapshots to occur. For each selected option, you must also set a retention policy, which can be set to pool automatic deletion or time-based retention.

For more information about Unity snapshots, see the [Dell EMC Unity: Snapshots and Thin Clones White Paper](#).

Thin clones

A thin clone is a read-write copy of a thin block storage resource, such as a volume, consistency group, or VMware VMFS datastore, that shares blocks with the parent resource. Thin clones are a great way to present SQL Server database copies in a fast and space-efficient manner that cannot be achieved with traditional SQL Server tools. After the thin clone is presented to the host, the volumes can be brought online, and the database is attached using the attach database method in SQL Server.

Note: When using the refresh feature with thin clones, take all databases that are on the thin clone offline before performing the refresh operation. SQL Server is unaware of when the thin-clone-based LUNs are being reverted. Failure to take the databases offline before performing a refresh can result in data inconsistency errors or incorrect data results from SQL Server, or both.

Replication

Replication is a software feature that synchronizes data to a remote system within the same site or at a different location. Unity replication and configuration options provide an effective way to meet the recovery time objective (RTO) and recovery point objective (RPO) requirements of SQL Server databases while maintaining a good balance of performance and bandwidth efficiency.

When using Dell EMC Unity replication to protect SQL Server databases on multiple volumes, contain all data and log volumes for a database within a single consistency group or file system. Replication is then configured on the consistency group or file system and can contain volumes or shares for more than one database. Databases that require different replication settings must be on separate LUNs, consistency groups, or file systems.

Thin clones are compatible with both synchronous and asynchronous replication. When a thin clone is replicated to the destination, it becomes a full copy of the volume, consistency group, or VMFS datastore. After being replicated, the thin clone is a fully independent volume with its own settings. The following figure shows the replication process of a thin clone volume:

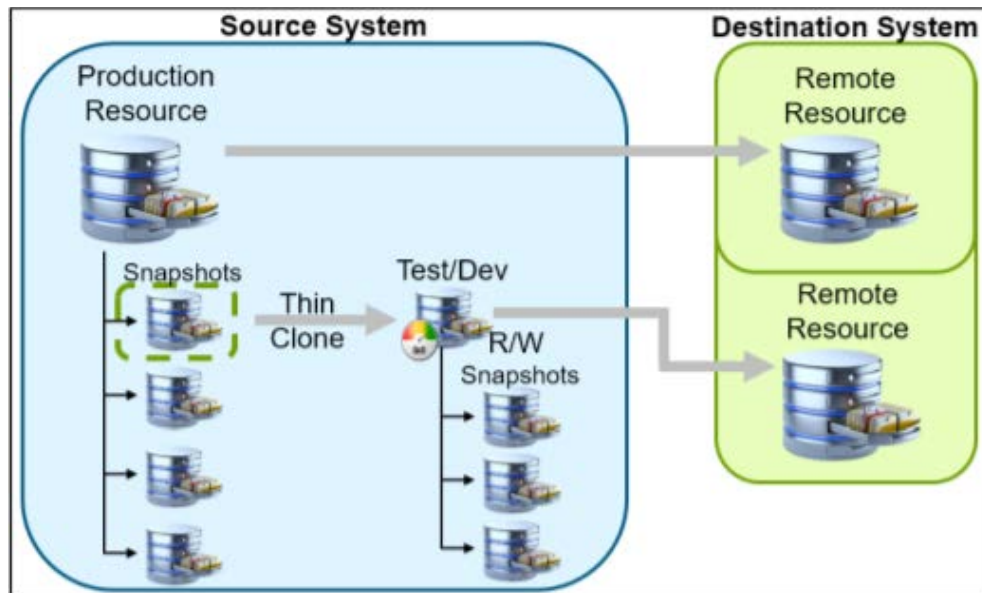


Figure 1. Replication of a thin clone volume

Replication of tempdb is never required because the file is rebuilt when SQL Server is restarted and so the metadata does not match that of another SQL Server instance. Carefully selecting the volumes to replicate and the contents of those volumes eliminates unnecessary replication traffic.

Integrated Copy Data Management for Microsoft SQL Server

SQL Server professionals have been extremely cautious when considering the use of storage-based snapshots and replication. Although different vendors might use different marketing terms to describe product features, most modern storage products (including all Dell EMC products) can produce “operating-system consistent” copies of any type of file by:

- Preserving operating-system write order consistency through all the levels from the host to the disk
- Offering a way to group volumes so that the write order is preserved across multiple files on different volumes

Unfortunately, the database community refers to the copies that are produced by these features as “crash-consistent copies.” The term came into use because recovering a database from this type of copy is the same mechanism that is used when the operating system shut downs unexpectedly and SQL is unable to resolve the changes that are stored only in memory with what is on disk. The use of term “crash-consistent copies” mostly fuels fears and limits legitimate use cases of storage-based copies and replication.

With the widespread adoption of highly scalable storage appliances, the Windows and SQL teams at Microsoft have developed an API for use by storage vendors. The API enables storage vendors to coordinate with the SQL Server database software to create “application-consistent copies” by using the Volume Shadow Copy Service (VSS). These copies mimic how SQL Server and the operating system cooperate during a planned and orderly shutdown of SQL Server. All write buffers are flushed and transactions are suspended until all the disks are up to date and consistent at a known point in time that is recorded in the SQL log.

Dell EMC AppSync software integrated with Unity XT snapshots simplifies and automates the process of generating, consuming, and managing application-consistent copies of production data. The software addresses copy management use cases for database recovery and database repurposing and is especially useful in a consolidated SQL database environment. AppSync software automatically discovers application databases, learns the database structure, and maps the file structure through the hardware or virtualization layers, or both, to the underlying Unity XT storage. It orchestrates all required activities, from copy creation and validation through mounting the snapshots at the target host and launching or recovering the database. The AppSync solution supports and simplifies SQL Server workflows that include refreshing, expiring, and restoring the production database.

Data reduction and advanced deduplication

The Dell EMC Unity family of storage systems offers feature-rich, easy-to-use data reduction services. Data reduction savings are achieved not only on the configured primary storage resources but also on the snapshots and thin clones of those resources. Snapshots and thin clones inherit the data reduction setting of the source storage resource, which increases the space savings provided.

The data reduction feature includes deduplication, compression, and zero-block detection activities that potentially increase the amount of useable storage space for user objects and internal uses. The Unity XT data reduction feature replaces the compression feature

in Unity OE 4.3 and later versions. Compression is an algorithmic data reduction method that can reduce the physical allocation of storage required to save a dataset.

Unity XT systems also provide an advanced deduplication feature, which can be enabled if data reduction is enabled. Advanced deduplication reduces the amount of required storage for user data by keeping only a small number of copies (often just one copy) of a Unity data block. The deduplication scope is a single LUN, so consider this fact when choosing the storage layout. Fewer LUNs result in better deduplication, but more LUNs provide increased performance. The space savings from advanced deduplication can provide the greatest return in most environments but also requires the most CPUs in Unity storage. In OE 5.0, advanced deduplication, if enabled, deduplicates any block (compressed or uncompressed). For more information, see the [Dell EMC Unity: Data Reduction White Paper](#) and the [Dell EMC Unity: Best Practices Guide](#).

The following table shows the supported configurations for data reduction and advanced deduplication:

Table 3. Unity (all models) data reduction and advanced deduplication support

Unity OE version	Technology	Supported pool type	Supported models
4.3 / 4.4	Data reduction	All-flash pool— traditional or dynamic	300, 400, 500, 600, 300F, 400F, 500F, 600F, 350F, 450F, 550F, 650F
4.5	Data reduction		300, 400, 500, 600, 300F, 400F, 500F, 600F, 350F, 450F, 550F, 650F
	Data reduction and advanced deduplication*		450F, 550F, 650F
5	Data reduction		300, 400, 500, 600, 300F, 400F, 500F, 600F, 350F, 450F, 550F, 650F, 380, 480, 680, 880, 380F, 480F, 680F, 880F
	Data reduction and advanced deduplication	450F, 550F, 650F, 380, 480, 680, 880, 380F, 480F, 680F, 880F	

* Data reduction is disabled by default and must be enabled before advanced deduplication is an available option. After data reduction is enabled, advanced deduplication is available but is disabled by default.

Unity data reduction versus SQL Server compression

SQL Server 2008 Enterprise Edition was the first version and edition to include native data compression capability. SQL Server 2008 row- and page-level compression uses knowledge of the internal database table format of SQL Server to reduce the space that database objects consume. The space reduction enables more rows to be stored on a page and more pages to be stored in the buffer pool. Since any data not stored in the 8k data page format—for example, off-row data such as NVARCHAR(MAX)—will not benefit from the row or page compression techniques, Microsoft introduced the Transact-SQL COMPRESS and DECOMPRESS functions. These functions use a traditional data

compression approach (GZIP algorithm) that must be invoked for each data section to compress or decompress data.

Unity XT compression, which is not specific to SQL Server only, uses a software algorithm to analyze and compress data within a storage resource. Unity data compression has been available for block storage volumes and VMFS datastores in an all-flash pool since the release of Unity OE 4.1. Compression is also available for file systems and NFS datastores in all-flash storage pools starting with Unity OE 4.2.

The best choice of data compression for SQL Server data depends on multiple factors. These factors include the type of database content, the amount of available CPU on both the storage and the database servers, and the amount of I/O resources that are required to maintain service-level agreements. In general, you can expect additional space savings with Unity XT arrays for data that is compressed by SQL Server when compression relies on product-specific data page construction as opposed to a general-purpose data compression technology. Data compressed with the TSQL Compress function using the GZIP algorithm is unlikely to gain significant additional reduction from Unity XT compression features because most of the advantages are gained from the first general-purpose algorithm applied.

Unity compression provides space savings if the data on the storage object is at least 25 percent compressible. Before enabling compression on a storage object, determine if it contains data that can be compressed. Do not enable compression on a storage object if space savings will not result. Contact your Dell EMC representative for tools that can analyze the data compressibility. For more information about compression, see the [Dell EMC Unity: Compression White Paper](#).

When deciding whether to use Unity data reduction, SQL Server database-level compression, or both, consider the following information:

- Data that is written to a Unity system is acknowledged to the host after it is stored in system cache. However, the compression process is not invoked until the cache is cleaned to provide the fastest response to the host.
- Compression savings are not only achieved on the compression-enabled Unity XT storage resource but on snapshots and thin clones of the resource as well.
- During the compression process, multiple blocks are aggregated and sent through a sampling algorithm to determine if the data should be compressed. If the sampling algorithm determines that only minimal savings can be achieved, compression is skipped, and the data is written to the pool as if compression is disabled.
- When data is compressed before it is written to disk, drive operations are greatly reduced. Compression also helps to reduce flash wear by reducing the physical amount of data being written to disk.

For more information about SQL Server row and page compression for rowstore tables and indexes, and columnstore and columnstore archival compression for columnstore tables and indexes, see [Data Compression](#) in Microsoft SQL Docs.

Note: All compression requires CPU resources. At high throughput levels, compression can have a detectable impact on performance. The heavy write ratios of OLAP workloads can also reduce the benefits of compression for a SQL Server database.

The Dell EMC Unity engineering team explored the potential savings using real-world data reduction ratios on a Unity array. The team gathered data from VMware virtual machines (VMs), file-share data, SQL Server databases, Microsoft Hyper-V VMs, and more.

The study results showed that the SQL Server log file reduction was nearly 10 times that of the data file:

- Database volume = 1.49:1 (32.96 percent)
- Logs volume = 12.9:1 (92.25 percent)

The SQL Server database was provisioned with two volumes. The database files were stored on one volume and the transaction logs on another. Using data reduction technology with database volumes can produce storage savings; however, you must consider the performance impacts when deciding whether to enable deduplication on database volumes. While actual database reduction levels can vary based on the data being stored, study results showed that storage space for SQL Server transaction logs can be reduced significantly.

For more information about the Unity engineering team's study and findings, see [Dell EMC Unity: Data Reduction Analysis](#).

Data reduction best practices

Before enabling data reduction on a storage object, consider these best practices:

- Monitor the storage system to ensure that it has available resources to support data reduction.
- Enable data reduction on a few storage objects at a time. Monitor the system to be sure it is still within the recommended operating ranges before enabling data reduction on additional storage objects.
- With Unity XT x80F models, consider that data reduction will provide space savings if the data on the storage block is at least 1 percent compressible.

Note: Data reduction on previous Unity x80F models running OE 5.0 provided space savings if the storage block data was at least 25 percent compressible.

- Before enabling data reduction on a storage object, determine if the object contains compressible data. Certain types of data—such as video, audio, image, and binary data—usually get little benefit from compression. Do not enable data reduction on a storage object if no space savings will result.
- Consider selective volume compression for file data, which usually compresses well.

Contact your Dell EMC representative for tools that can analyze the potential for data compressibility.

For more information about compression, see the [Dell EMC Unity: Data Reduction White Paper](#).

VMware virtualization

Introduction

VMware vSphere is an efficient and secure platform for virtualization and the cloud. vSphere core components are VMware vCenter Server and the VMware ESXi hypervisor.

vCenter Server is a unified platform for managing vSphere environments. It features ease of deployment, extensibility onto the public cloud, and proactive optimization of resources for maximum efficiency.

ESXi is a bare-metal hypervisor that is installed directly onto physical servers. ESXi is more efficient than hosted architectures because it has direct access to and control of underlying resources. The ESXi hypervisor has a small footprint of 150 MB, minimizing space requirements and security threats. It provides reliable performance for different application workloads and supports VM configurations of up to 128 vCPUs, 6 TB of RAM, and 120 devices.

SQL Server and VMware virtualization best practices

This section provides a summary of best practices for virtualizing SQL Server with VMware. This information is based on the latest VMware guidance. For more information, see [Architecting Microsoft SQL Server On VMware vSphere—Best Practices Guide](#).

Processors (CPU)

For SQL Server to run efficiently on modern hardware, the SQL Server Operating System (SQLOS) must have a full understanding of hardware layout. With the advent of multicore and multinode non-uniform memory access (NUMA) systems, understanding relationships among cores, logical CPUs, and physical CPUs has become important.

VMware uses the following terms for physical hardware:

- **Physical CPU or physical socket**—A physical CPU installed in the server hardware.
- **Physical core**—An independent processing unit residing on a single processor or physical CPU
- **Logical core**—A logical processor on a physical core with own processor architectural state. Intel Hyper-Threading technology is the most widely deployed example of logical cores.

The corresponding VMware virtualization terminology is:

- **Virtual socket**—Represents a virtualized physical CPU and can be configured with one or more virtual cores.
- **Virtual core**—Equal to a CPU and viewed by an operating system as a separate processor unit.
- **Virtual CPU (vCPU)**—Virtualized central processor unit assigned to a VM. Total number of assigned vCPUs to a VM is calculated as:

Total vCPU = (Number of virtual sockets) * (Number of virtual cores per socket)

Consider the following information about vCPUs when using VMware virtualization with SQL Server:

- When consistent performance is high priority, VMware recommends that the total number of vCPUs that are assigned to all the VMs not exceed the total number of physical cores available on the ESXi host machine.
 - You might be able to increase the number of allocated vCPUs if monitoring suggests that unused CPU resources are available.
 - On systems with Intel Hyper-Threading enabled, the number of logical cores (vCPUs) is twice the number of physical cores. In such case, do not assign the total number of vCPUs.
- Lower-tier SQL Server workloads are, by definition, less impacted by latency variability. Therefore, these workloads can be run on hosts with greater ratios of vCPUs to physical CPUs.
 - The vSphere CPU scheduler policy is designed to balance maximum throughput and fairness between VMs.
 - Reasonable levels of CPU overcommitment can increase overall system throughput, maximize license savings, and continue to maintain adequate performance.
- Intel Hyper-Threading generally improves the overall host throughput by 10 to 30 percent, which suggests a 1.1 to 1.3 ratio of vCPUs to physical CPUs.
 - VMware recommends enabling Hyper-Threading, when available, in the UEFI BIOS so that ESXi can take advantage of the technology.
 - ESXi CPU management is aware of when Hyper-Threading is enabled; therefore, for example, a VM with two vCPUs is mapped to different physical cores instead of to two logical threads on the same physical core.
 - VMware recommends that you perform thorough testing and monitoring when you are using Hyper-Threading for SQL Server workloads.

Memory

Almost all modern servers use a non-uniform memory access (NUMA) design for communication between main memory and the CPUs. NUMA is a hardware architecture for shared memory, implementing subdivision of physical memory banks between physical CPUs. A NUMA node is defined by one or more CPU sockets together with a block of dedicated memory. A single CPU socket and connected memory is a NUMA node.

During the past decade, NUMA has been a much-discussed topic. The relative complexity of NUMA is driven, in part, by different vendor implementations. In virtualized environments, the complexity of NUMA is also driven by the number of configuration options and layers that must be traversed—from the hardware through the hypervisor to a guest operating system, and, finally, to the SQL Server application. A good understanding of NUMA hardware architecture is a requirement for any infrastructure architect or SQL Server DBA working with a virtualized SQL Server.

To achieve better efficiency on servers with high core counts, Microsoft introduced Soft-NUMA. Soft-NUMA enables the partitioning of available CPU resources within one NUMA

node into multiple “Soft-NUMA” nodes. According to VMware, Soft-NUMA is compatible with the VMware virtual NUMA (vNUMA) topology and might further optimize scalability and performance of the Database Engine for most of the workload.

Consider the following information about memory resources when using VMware virtualization with SQL Server:

- Monitor VMs to detect a lack of memory resources for the SQL Server Database Engine. This issue induces Windows Server to page memory to disk, resulting in increased disk I/O activities and poor performance.
- When designing for performance, prevent memory contention between VMs by avoiding overcommitment of memory at the ESXi host level.
- Consider checking the hardware physical NUMA memory allocation to identify the maximum amount of memory that can be assigned to a VM without crossing the physical NUMA boundaries.
- When achieving adequate performance is the primary goal, consider setting the memory reservation equal to the provisioned memory. This setting adjustment eliminates the possibility of ballooning or swapping and guarantees that the VM gets only physical memory.

Virtualized storage

Configuring storage in the virtualized environment requires knowledge of the storage infrastructure. As with NUMA, it also requires an understanding of the coordination between each level in the I/O path—in this case, from the application running in a VM to the physical reading and writing of information on a persistent storage medium.

vSphere provides the following options for configuring storage, all of which have useful applications in a SQL Server implementation with the Unity XT array:

- VMware Virtual Machine File System (VMFS)
 - VMFS is the most widely used datastore on block storage systems such as Unity XT.
 - The Unity XT array is the bottom layer, consisting of physical disks presented as logical disks (volumes) to vSphere.
 - The Unity XT volumes are formatted as VMFS volumes by the ESXi hypervisor.
 - VMware administrators create one or more virtual disks (VMDKs) that are presented to the guest operating system.
- Raw Device Mapping (RDM)
 - RDM enables a VM to directly access a block storage volume (FC or iSCSI) on a Unity XT array without VMFS formatting.
 - Both VMFS and RDM volumes can provide similar transaction throughput. For more information, see [Performance Characterization of Microsoft SQL Server on VMware vSphere 6.5](#).
- NFS
 - For NFS-based storage for ESXi, Dell EMC recommends VMware NFS datastores instead of general-purpose NFS file systems.

- A VM running SQL Server and using VMDKs on an NFS datastore is not aware of the underlying NFS layer. The guest operating system views the VM as a physical server running Windows Server and SQL Server.
- Shared disks for Failover Cluster Instance configurations are not supported using NFS datastores.
- VMware vSphere Virtual Volumes (VVols)
 - The best practices and guidelines for using a [VMFS datastore](#) also apply when VVols are used as the underlying technology.
 - VVols offer finer control at the VM level, independent of the underlying physical storage representation (volumes or file systems, for example).
 - Storage operations are available with individual VM granularity, providing native array-based data services such as compression, snapshots, deduplication, encryption, and replication.
 - Array-based replication with VVols is supported beginning with VVol 2.0 (vSphere 6.5).
 - A VVol disk can be used instead of an RDM disk to provide a disk resource for a SQL Failover Cluster Instance beginning with vSphere 6.7 with support for SCSI-3 persistent reservation.

Virtualized networking

Networking in the virtual world follows the same logical concepts as in the physical world but uses software rather than physical cables and switches. The impact of network latency on SQL Server workloads can vary greatly. Monitoring networking performance metrics on either the existing workload or a well-implemented test system for a representative period helps with the virtual network design.

Consider the following networking information when using VMware virtualization with SQL Server:

- Both standard and distributed virtual switch types provide the functionality required by SQL Server.
- Use VLAN tagging and virtual switch port groups to logically separate management, vSphere vMotion, and network-based storage traffic.
- VMware highly recommends enabling jumbo frames on the virtual switches where you have enabled vSphere vMotion traffic or iSCSI traffic.
- In general, follow the networking guidelines for guest operating systems and hardware.

Summary

SQL Server database environments continue to grow in both size and complexity. With SQL Server 2019, Microsoft improved SQL Server core features and added new ones such as support for Big Data workloads with Apache Spark and HDFS. Dell EMC, in partnership with Microsoft, continues to provide the essential infrastructure components—servers, storage, and networking—for your SQL Server environment.

The adoption of storage appliance technology by SQL Server DBA professionals has varied greatly across organizations. We have seen significant increases in uptime and reductions in the total cost of ownership (TCO) when storage and database professionals collaborate to design infrastructure solutions for SQL Server on shared storage platforms. The Dell EMC Unity XT All-Flash storage array is a midrange solution that is ideal for SQL Server developers and administrators who are seeking high performance and low latency. Designed for all-flash performance, the Unity XT All-Flash system supports dual-socket CPUs, dual-active controllers, and multicore optimization.

Increasingly, organizations are virtualizing their SQL Server environments at some level. Although virtualization adds another layer of design, configuration, and monitoring to the architecture stack, it provides significant benefits. In this paper, we describe some of the commonly used VMware features and benefits in SQL Server environments. We also include best-practice guidelines and links to resources for more detailed information.

References

Dell EMC documentation

The following Dell EMC documentation provides additional and relevant information. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell EMC representative.

- [Microsoft SQL Info Hub for Ready Solutions](#)
- [Dell EMC Unity Storage with Microsoft SQL Server](#)
- [Dell EMC Unity: Best Practices Guide](#)
- [Dell EMC Unity: Compression White Paper](#)
- [Dell EMC Unity: Snapshots and Thin Clones White Paper](#)
- [Dell EMC Unity: Data Reduction Analysis](#)
- [Dell EMC Unity: Data Reduction White Paper](#)
- [Dell EMC Unity XT Storage Series Specification Sheet](#)
- [Dell EMC XT Unity All-Flash Unified Storage](#) (product page)
- [Dell EMC Unity Product Overview and Architecture with Kaushik Ghosh](#) (YouTube)
- [Dell EMC Future-Proof Loyalty Program](#)

VMware documentation

The following VMware documentation provides additional and relevant information:

- [SQL Server Blogs](#)
- [Virtualizing Microsoft SQL on VMware vSphere](#)
- [Architecting Microsoft SQL Server On VMware vSphere—Best Practices Guide](#)
- [Performance Characterization of Microsoft SQL Server on VMware vSphere 6.5](#)

Microsoft documentation

The following Microsoft documentation provides additional and relevant information:

- [Microsoft SQL Server 2019 technical white paper](#)
- [SQL Server Blog](#)
- [Microsoft Data Platform](#)
- [Introducing SQL Server 2019](#)
- [SQL Server 2019 CTP announcement archive](#)