Hitachi Integrated Systems Solution with Microsoft SQL Server

Reference Architecture Guide

MK-SL-355-00
July 2025

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to doc.comments@hitachivantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Thank you!

Revision history

Changes	Date
Initial release	July 2025

Reference Architecture Guide

Executive summary

The data management market is growing at a rapid pace in general. IDC projects worldwide enterprise spending to grow from $23B in 2024 to $32B in 2027 at ~11% YoY.

Statista uses a ranking score to measure the popularity of databases, and as of June 2024, Microsoft SQL Server is the 3rd most popular database in the market.

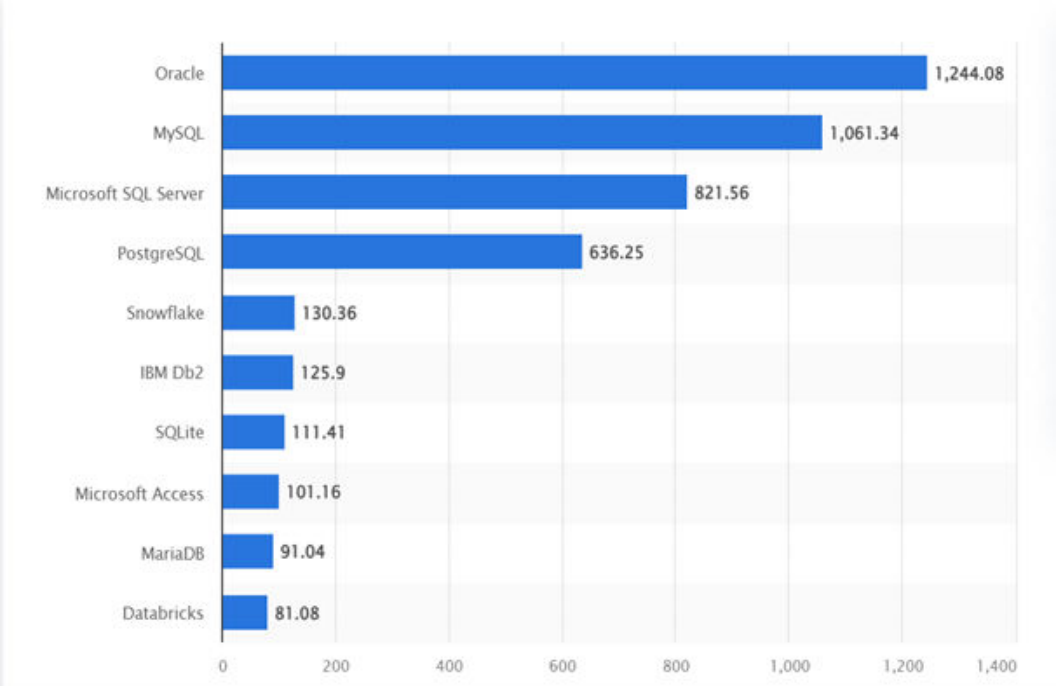

Customers tracked by 6sense show the variety of uses of databases in the industry, with the top usages being related to Software Development, Web Development, and Machine Learning.

Introduction

Microsoft SQL Server is a relational database management system (RDBMS) developed by Microsoft, used to store, retrieve, manage, and manipulate data within a database. It's a software product that allows users to interact with databases using the Structured Query Language (SQL), a programming language, create tables, insert data, query data, and perform other database operations.

Microsoft SQL Server offers features for data integration, reporting, and analysis, including SQL Server Integration Services (SSIS) for data movement, SQL Server Reporting Services (SSRS) for reporting, and SQL Server Analysis Services (SSAS) for data mining and OLAP. These capabilities allow Microsoft SQL Server to be used in various applications, including transaction processing, business intelligence, and data warehousing.

Microsoft offers different editions of SQL Server to suit different needs, including a free Developer Edition, a free Express Edition for small databases, and Standard and Enterprise editions for larger deployments. Microsoft also offers Azure SQL, a cloud-based SQL Server database service.

What makes it unique is that Microsoft SQL Server stands out as a robust and versatile RDBMS with strong integration with other Microsoft products, offering features such as high performance, data security, scalability, and business intelligence capabilities, making it a popular choice for various applications.

Business use cases

Small Medium Business (SMB) customers as well as the large Enterprise customers choose Microsoft SQL Server deployments in their data centers for various reasons: to protect their data (security, privacy, etc.) and to maintain the lowest transaction latency in processing the data.

Depending upon the type of application (banking transactions, healthcare insurance records, ERP, or BI) a different type of infrastructure is needed, and depending upon the business requirements, a certain level of performance is mandated. These requirements are right-sized by infrastructure teams into Servers connected to Storage via Network, forming the basis of configurations that are best suited for each use case and application.

An example of right-sizing such infrastructure, depending upon the customer scenario, transactional processing (OLTP) for a heavily used database by numerous users in a mission-critical scenario might require high-end block storage and high-end, scale-up servers, connected by high-speed networking. Whereas analytical processing (OLAP) of historic data that numerous users do not actively use, but larger in volume, for the sake of report generation (or BI) may require mid-range block storage with mid-range scale-out servers connected by less demanding networking.

Solution benefits

Hitachi Vantara has a history of providing solutions tailored to every type of workload, including various databases, data warehouses, and AI applications. Because databases form the common denominator for all workloads, a solution with databases is designed with the same flexibility as the number of use cases in which databases can be employed.

Hitachi Vantara offers solutions for different sizes, different deployment scenarios, and different use cases. A reference architecture forms the best practice guidelines for using the HitachiIntegrated Systems infrastructure solution stack with customer’s application stack with Microsoft SQL Server.

Here are some of the benefits of using the Hitachi Integrated Systems solution with Microsoft SQL Server:

1. Accelerate application performance with Hitachi Integrated Systems.

   Slow application response times and outdated data can stall decision-making and frustrate users. The Hitachi Integrated Systems solution is purpose-built to accelerate Microsoft SQL Server performance by optimizing the underlying storage infrastructure. Our low-latency, high-throughput storage delivers rapid access to data, enabling applications to respond faster and improving the end-user experience across business-critical systems.

   Siloed data and fragmented systems force employees to toggle between tools, reducing productivity and introducing inconsistencies. Hitachi's infrastructure centralizes data across departments and workloads, ensuring that all teams access a unified, consistent, and real-time view of critical business information. This eliminates the lag and confusion caused by outdated or inaccessible data, empowering employees to make faster, more informed decisions.

   By leveraging Hitachi's high-performance storage architecture, applications built on SQL Server gain dramatically improved I/O performance—minimizing query times, accelerating reporting, and shortening batch processing windows. This enables smoother workflows for users and faster time-to-insight for the business.

   In addition, our solution seamlessly connects SQL Server with other enterprise applications, automating routine tasks like reporting and data consolidation. This ensures consistency across systems, reduces manual work, and supports a truly agile business environment. For example, when integrated with tools like CRM and customer support platforms, real-time access to updated customer data enables more personalized and efficient service.

   With Hitachi's integrated infrastructure for SQL Server, your applications don't just work—they perform. Faster. Smarter. Better.

2. Reduce TCO, balancing infrastructure costs with software license costs.

   Managing multiple disconnected systems often leads to higher costs and duplication of efforts. Our integrated infrastructure solution for Microsoft SQL Server reduces overhead costs by reducing manual errors, streamlining operations, and improving efficiency by automating process workflows, unifying systems, and consolidating resources.

   An infrastructure solution that minimizes CPU/core counts minimizes Microsoft SQL Server licenses. Also, a well-planned infrastructure solution minimizes the premium data center footprint, reduces power/cooling costs, and optimizes resource management.

   Unified infrastructure management automation further reduces inefficiencies leading to significant cost savings over time, as businesses no longer need to maintain disconnected systems.

   At the business apps layer, by automating data integration processes, businesses save time and resources, further reduce operational costs and license/support footprint for the organization. An example of this are ERP integrations that add business value while reducing the incurring of repeated costs by ensuring consistent, real-time data across all functions.

3. Take advantage of the scalability and flexibility of database deployment sizes.

   Outdated legacy systems and disconnected systems struggle to adapt to growing with the business needs. Our integrated infrastructure solution for Microsoft SQL Server provides scalability and flexibility in growing and adapting to a business's changing needs, making them more sustainable for long-term use.

   Modularity in our integrated systems means business apps based on Microsoft SQL Server can grow with the business needs, accommodating increasing data volumes, while offering scalability to manage expanding datasets and new data sources seamlessly with new applications and tools. Our solutions are offered with various t-shirt sizing, for performance or for affordability, for transactions (OLTP) or for analytics (OLAP), for small to large databases and for consolidation of numerous databases, using our immense storage scalability with VSP.

   Likewise, integrated solutions for ERP based on Microsoft SQL Server can grow with the business by supporting additional modules, users, and increased data requirements.

4. Drive resilient, compliant operations with Hitachi's high-performance storage infrastructure.

   Operational inefficiencies often stem from manual processes and fragmented systems. Managing compliance and enforcing security across these disjointed environments creates unnecessary complexity and risk. Hitachi's Integrated Systems solution delivers a unified storage infrastructure that simplifies and secures enterprise operations—powering critical business applications like Microsoft SQL Server with speed, reliability, and control.

   Our high-availability storage platform enables consistent data access across all departments, helping organizations eliminate redundant tools for HR, finance, operations, and beyond. With centralized infrastructure at the core, different systems can seamlessly share data, reducing duplication, improving data integrity, and enhancing workflow efficiency across the enterprise.

   As historical data volumes increase, our storage systems ensure scalable performance and seamless consolidation from disparate sources into a unified data environment. This allows for faster querying, easier access, and streamlined analysis—while ensuring that all stored data remains consistent, accurate, and up to date across applications. Built-in data protection, encryption, and access controls also strengthen your security posture while enabling easier adherence to regulatory frameworks.

   Unlike isolated storage environments that make governance and compliance reporting a challenge, Hitachi's integrated storage infrastructure enables centralized monitoring, auditability, and policy enforcement. With data flowing through a single, resilient platform, your teams can more easily align with industry regulations and security standards—while reducing the risk of breaches and non-compliance penalties.

   When ERP systems, databases, and analytics platforms are underpinned by Hitachi's intelligent storage foundation, business operations become more reliable, secure, and scalable—enabling enterprises to confidently meet today's performance demands and tomorrow's compliance challenges.

Reference Architecture Guide for Hitachi Solution for Microsoft SQL Server

Use this reference architecture guide to understand how the Hitachi Integrated Systems solution with Microsoft SQL Server provides a high-performance, low-latency, integrated, and converged solution for Microsoft SQL Server using Hitachi Virtual Storage Platform One Block. The environment uses VSP One Block and Hitachi Advanced Server HA820 G3 with 3rd Generation Intel Xeon Scalable Processors

With these products, design an SQL Server converged infrastructure to meet your requirements and budget.

This solution uses Hitachi's latest block storage offering high-performing VSP One Block storage system with NVMe SSDs to boost performance and lower I/O latency. Hitachi Advanced Server HA820 G3 is used in this reference architecture to run a dedicated Microsoft SQL Server 2022 database with the Windows Server 2022 clusters option, and it uses Windows 2022 Datacenter edition for the operating system.

This document is useful for the following roles:

- Database administrator
- Storage administrator
- Database performance analyzer
- IT professional with the responsibility of planning and deploying a Microsoft SQL Database solution

To use this reference architecture guide, you need familiarity with the following:

- Hitachi Virtual Storage Platform One Block
- Hitachi Advanced Server HA820 G3 Servers
- Storage area networks
- Microsoft SQL Server 2022
- Windows Server 2022 Datacenter version operating system
- HammerDB

Note: Testing of this configuration was in a lab environment. Many factors affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Solution overview

Hitachi Integrated Systems solution with Microsoft SQL Server database is a highly engineered, pre-tested, and validated system designed to deliver exceptional performance and reliability in dynamic and demanding SQL environments.

This reference architecture integrates the Hitachi Integrated Systems platform, using Hitachi Virtual Storage Platform One Block storage. It meets the high availability, performance, and scalability needs of both OLTP and OLAP workloads. The solution is built on the Hitachi Advanced Server HA820 G3 and VSP One Block storage systems.

To benchmark the Microsoft SQL Server performance with the Hitachi VSP One Block storage, we used the open-source HammerDB tool, which is widely recognized within the industry. Derived from TPC specifications, HammerDB runs two primary workloads: transactional and analytical, to effectively evaluate system performance.

Business benefits

Here are some benefits of this reference architecture:

- Achieve high Microsoft SQL Server Database performance with VSP One Block storage systems.
- Provide a solution for customers who are looking for low I/O latency, high throughput and minimal response time for Microsoft SQL Server database.

High-level infrastructure

Hitachi Solution for Databases with Microsoft SQL server includes the following components:

- Hitachi Advanced Server HA820 G3 servers
- Hitachi Virtual Storage Platform One Block 26
- Brocade G720 32 Gbps SAN infrastructure
- Cisco 10/25 GbE LAN infrastructure

The configuration of the Virtual Storage Platform One Block 26 and Hitachi Advanced Server HA820 G3 features fully redundant hardware.

Note: While this reference architecture uses the VSP One B26, the solution supports the entire Hitachi storage portfolio.

The following figure shows the high-level infrastructure network for this solution.

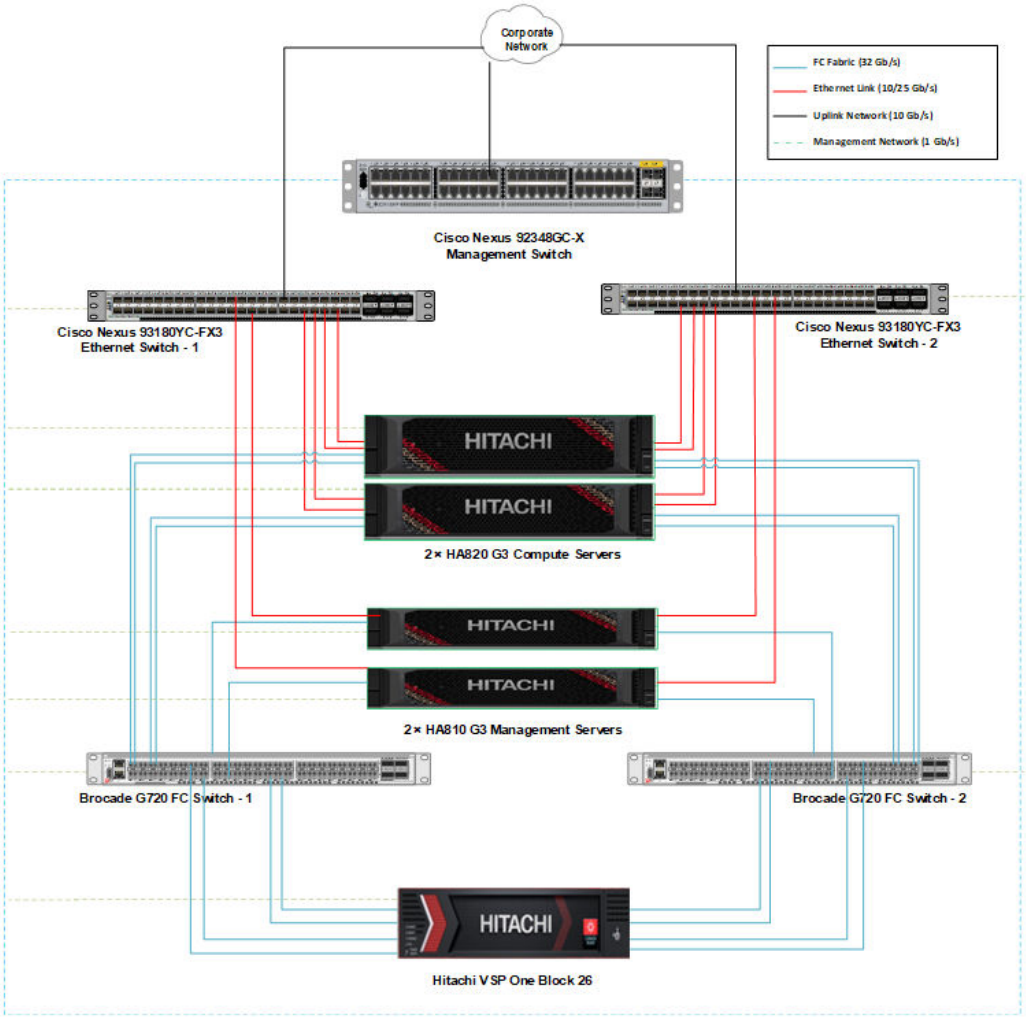

To ensure minimal impact on production database performance, Hitachi Vantara recommends the following configuration:

- Use a dedicated storage system for the production database.
- Implement a separate storage system for data replication at an alternate site to ensure business continuity, if required.

The uplink speed to the corporate network is determined by the customer's environment and specific requirements. For environments requiring higher bandwidth, the Cisco Nexus 93180YC-FX switches support uplink speeds of 40 GbE, with the option to upgrade to 100 GbE.

Key solution components

The key solution components for this solution are listed in the following tables. Detailed component information is provided in Product descriptions (on page 19).

The following table lists the hardware components used in this solution.

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	Hitachi VSP One Block 26	8 × CHA pairs (16 × 32 Gbps Fibre Channel ports)	A3-03-01-40/01 (TLC) A3-03-21-40/00 (QLC)	1
		24× 1 Backend NVMe SSD		
Hitachi Vantara	HA820 G3 server Compute nodes	2 × Intel Emerald Rapids Processor 8568Y+CPU @ 2300 MHz 64 GB× 32 DDR5 (2048 GB) RDIMM	iLO6: 1.58 Mar 22 2024 System ROM: U54 v2.16 (03/01/2024)	2
		2 × SN1700E 64 GB 2p FC HBA	Firmware:14.2.589.5 Driver: lpfc Driver version: 14.0.0.4	
		2 × Dual Port 25 GbE NIC Intel E810 PCIe cards	Driver: ice Driver version: 5.15 Firmware: 4.30	
Hitachi Vantara	Management Server HA810 G3 server	2 × Intel Xeon processors 4310, 12-core, 2.10 GHz, 120W 256 GB (32 GB × 8) DIMM DDR5	System ROM: U54 v2.16 (03/01/2024) iLO6:.58 Mar 22 2024	2
		1 × Dual Port 25 Gb ENIC Intel E810 PCIe card	Driver version: ice0.8.2-k Firmware:2.42	

Vendor	Hardware	Detail Description	Version	Quantity
		2 × SN1610E 32 Gb 2p Fibre Channel HBA	Driver: lpfc Driver version: 14.0.499.31 Firmware: 14.0.499.29	
Brocade	G720 Fibre Channel switches	24 × 48 port Fibre Channel switches 32 Gbps SFPs	Fabric OS: v9.1.1c	2
Cisco	Nexus C93180YC-FX3	48 × 10/25 GbE ports	NXOS: version 9.3.8	2
		6 × 40/100 Gbps Quad SFP (QSFP28) ports 48 × 10/25 GbE ports		
	Nexus 92348 GC-X	1 GE 48-Port Gb Ethernet switch	NXOS: version 9.3.8	1

Note: Customers can choose larger capacity SSDs to fit their business requirements.

Certain components may be optional depending on the existing infrastructure and required interconnect topology. This might include the SAN, IP switches, and the management servers; however, this reference architecture documents the environment tested in the lab to support a full deployment of the architecture including supporting components.

The following table lists the software components used in this solution.

Software	Version	Function
Windows 2022 Datacenter	1129-20348	Operating system for SQL Server
Microsoft SQL Server 2022 Enterprise edition	Server 2022 (RTM) - 16.0.1000.6 (X64)	Database software
SQL Server Management Studio (SSMS)	19.3.4.0	Software application used for configuring, manage and administering all components of SQL Server.
HammerDB	V4.12	Benchmarking kit

Solution design

This section describes the reference architecture for implementing Hitachi VSP One Block 26 storage with Microsoft SQL Server on Windows Server 2022 Datacenter. Each SQL Server environment uses one Hitachi Virtual Storage Platform One Block 26.

The infrastructure configuration includes the following:

- Windows 2022 Datacenter Servers – There are 2 × HA820 G3 servers with Windows 2022 Datacenter Edition as the operating system. We have added these 2-server hostnames to the existing domain.
- SQL Server Clustered environment- 2 × HA820 G3 servers are configured for active-passive cluster environment.
- Storage System –There are DDP Pools mapped to each port that are presented to the server as LUNs.
- SAN Connection –There are SAN connections to connect the Fibre Channel HBA ports to the storage.

Storage architecture

This section describes the storage architecture for this solution.

Storage configuration

The storage configuration takes into consideration Hitachi Virtual Storage Platform One Block (VSP One B26) and Microsoft recommended best practices for the design and deployment of database storage.

Hitachi Virtual Storage Platform One Block

VSP One Block arrives pre-configured, including Dynamic Drive Protection Pool (DDP) groups. DDP replaces traditional RAID groups in VSP One Block storage systems, providing the resilience of RAID 6 with distributed spare space and support for an arbitrary number of drives (from 9-32 per group). Adding drives one (or more) at a time is also supported. DDP improves the resilience of the appliance by dramatically lowering rebuild times.

The following figure shows the high-level storage configuration for this solution.

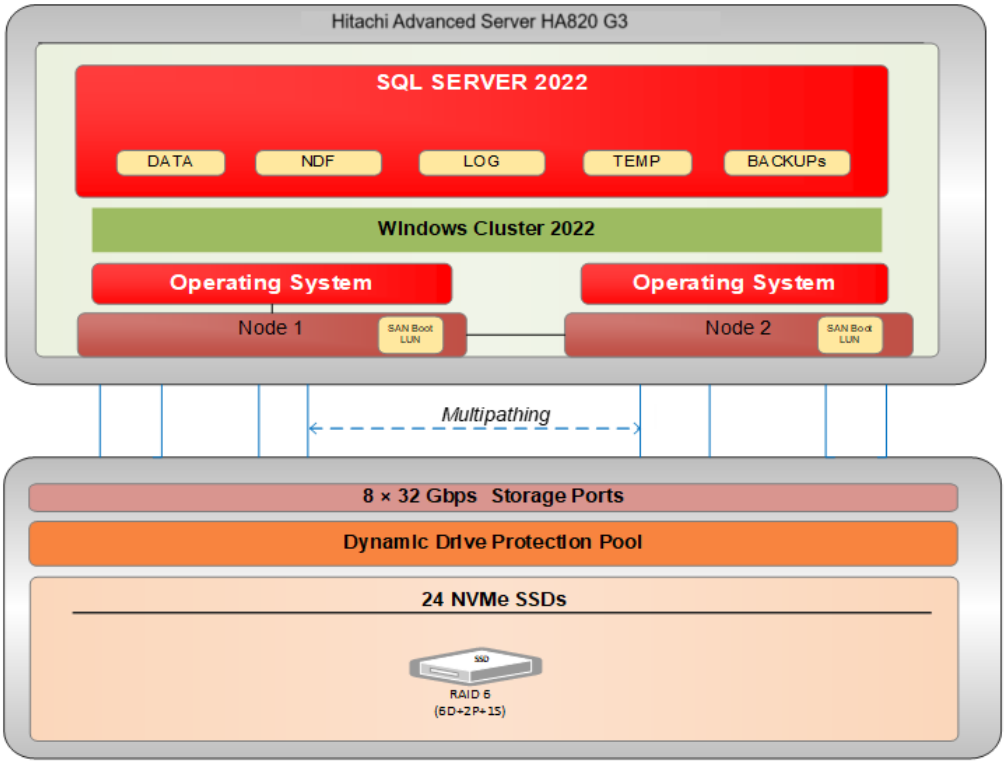

Hitachi VSP One Block 26

On the VSP One Block 26 storage system in this solution, we used a RAID 6 (6D+2P +1Spare) configuration along with Dynamic Drive Protection (DDP) for both TLC and QLC NVMe SSD-based configurations. Except for the SSD type and capacity, all other architectural components, system settings, and SQL Server configurations remained identical.

The following table provides storage pool details.

Pool ID	TLC SSD Configuration	QLC SSD Configuration
Pool Type	Dynamic Drive Protection Pool	Dynamic Drive Protection Pool
RAID Group	1-1	1-1
RAID Protection Type	RAID 6 (6D+2P+1S)	RAID 6 (6D+2P+1S)
Lun Type/Capacity	1.9 TiB NVMe SSD	30 TiB NVMe SSD
Number of Drives	9	24
Number of LDEVs	28	28
Total LDEV Sizes	7.74 TB	7.74 TB
Pool Capacity	10.11 TB	542.30 TB
Volume Type	DRS	DRS

The following table lists the logical storage configuration used in this solution.

LUN Usage	Number of LUNs	LUN Size	Purpose/Usage	Shared Across Cluster Nodes
OS Installation	2	100 GB	One LUN per server	No
Data LUN	20	300 GB	SQL Data files	Yes
Log LUN	2	320 GB	SQL Log files	Yes
Tempdb LUN	2	350 GB	SQL Temp DB files	Yes
SQL Server Binaries	2	100 GB	SQL Server System database files	Yes

Note: Each Windows node is set up with this drive configuration.

Database layout

The database layout design adheres to Hitachi Vantara best practices for the Hitachi Virtual Storage Platform One Block 26, particularly for workloads with small random I/O profiles such as OLTP transactions. It also incorporates Microsoft's best practices when leveraging Hitachi storage. The storage design for the database layout should be based on the specific requirements of the application implementation. The design can vary significantly depending on the RAID configuration and the number of LUNs used during setup. The components in this solution offer flexibility, enabling adaptation to different deployment scenarios, ensuring an optimal balance between performance and manageability for each unique use case.

Server and application architecture

SAN architecture

In this solution, the VSP One Block 26 storage array is configured with a RAID 6 (6D+2P+1 Spare) setup using 9 NVMe SSDs within a single DDP pool. Each SSD has a capacity of 1.9 TB, and with DDP enabled, the total usable capacity is 10.11 TB. This configuration provides a balance of data protection and performance, ensuring high availability and efficient storage usage.

Additionally, the DRS pool ensures that resources (such as CPU, memory, and storage) are utilized efficiently based on real-time demand. The system automatically adjusts resource allocation to optimize performance and ensure that workloads are balanced across the available storage resources, reducing potential bottlenecks and improving overall system efficiency.

The SAN connectivity is detailed as follows.

Map the provisioned LDEVs to multiple ports on each Hitachi Virtual Storage Platform One B26.These LDEV port assignments provide multiple paths between the storage system and the host for high availability. This reference architecture uses two dual port Emulex HBAs per HA820 G3 server.

- 8 SAN switch connections are used for VSP One Block Fibre Channel ports
- 8 SAN switch connections are used for server HBA ports

The following table lists details of the Fibre Channel SAN configuration between the Hitachi Virtual Storage Platform B26 and the HA820 G3 database servers.

Server	HBA	Host Group Name	Host Name	Switch Zone	Storage System	Storage Port	Brocade G720 Switch
HA820 G3 Server	HBA1	CN31	CN31_HBA1_1	CN31_HBA1_1_ASE42_43_1C	VSP One B26	CL1-C	SW1-29
	HBA2	CN31	CN31_HBA1_2	CN31_HBA1_2_ASE42_43_2C		CL2-C	SW2-30
	HBA3	CN31	CN31_HBA2_1	CN31_HBA2_1_ASE42_43_5C		CL5-C	SW1-29
	HBA4	CN31	CN31_HBA2_2	CN31_HBA2_2_ASE42_43_8C		CL8-C	SW2-30
HA820 G3 Server	HBA1	CN31	CN31_HBA1_1	CN31_HBA1_1_ASE42_43_3C		CL3-C	SW1-29
	HBA2	CN31	CN31_HBA1_2	CN31_HBA1_2_ASE42_43_4C		CL4-C	SW2-30
	HBA3	CN31	CN31_HBA2_1	CN31_HBA2_1_ASE42_43_6C		CL6-C	SW1-29
	HBA4	CN31	CN31_HBA2_2	CN31_HBA2_2_ASE42_43_7C		CL7-C	SW2-30

Microsoft Multipath I/O options and settings

To ensure continuous connectivity between the compute server and storage, SAN multipathing is necessary to achieve the following objectives:

- Data Availability: Ensuring maximum uptime for SQL Server.
- Path Failover: Enabling automatic failover in case of a link failure.

This solution was tested using Microsoft Multipath I/O (MPIO). The recommended configuration for SAN multipathing in this reference architecture is to use the Least I/O load-balancing policy, which ensures data redundancy and optimal performance for sequential workloads.

Network architecture

The Windows 2022 Servers network is configured with dual port 25 GbE Ethernet NICs to provide redundant connectivity. NIC teaming was used to allocate the private and public IP addresses to the two Windows Servers. When creating NIC teaming pairs, ports should be used from different cards to avoid single points of failure.

We recommend using pairs of 10 Gbps NICs for the interconnect network and public network.

The following figure shows the IP network switch connection.

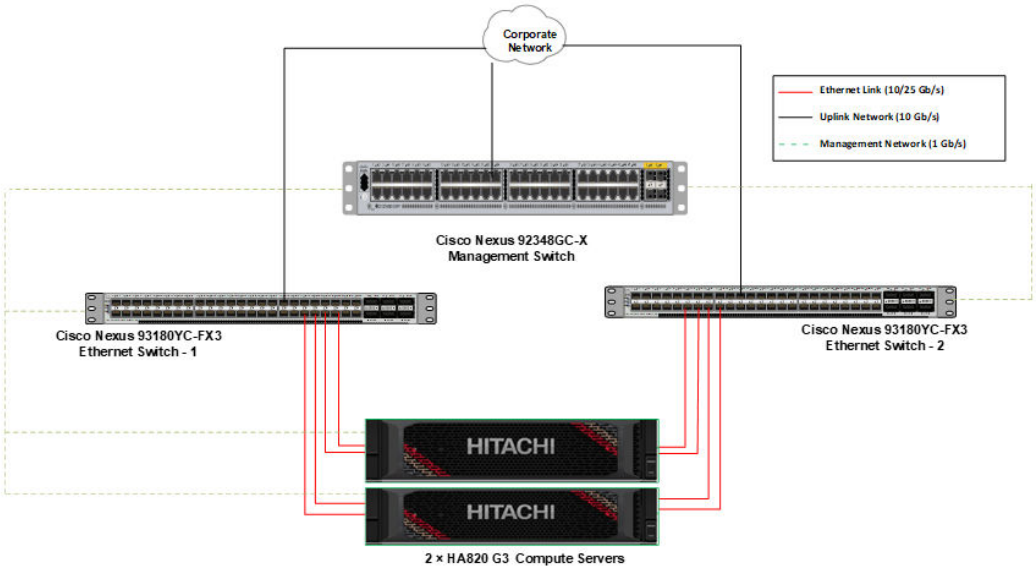

The following table lists the network configuration used in the testing environment.

Server	NIC Port	VLAN/ Subnet	NIC Teaming	IP Address	Network	Bandwidth (Gbps)
SQL Server Database Server1	NIC 1	33	Team0	192.168.1.10	Private	25
	PORT 1					
	NIC1					25
	PORT 2					
	NIC 2	33	Team1	10.76.33.160	Public SQL	25
	PORT 1					
	NIC 2					25
	PORT 2					
SQL Server Database Server2	NIC 1	33	Team0	192.168.1.11	Private	25
	PORT 1					
	NIC 1					25
	PORT 2					
	NIC	33	Team1	10.76.33.161	Public SQL	25
	PORT 1					
	NIC 2					25

Server	NIC Port	VLAN/ Subnet	NIC Teaming	IP Address	Network	Bandwidth (Gbps)
	PORT 2					

Engineering validation

This section highlights the key insights derived from the test results of the deployment, using the Hitachi Advanced Server HA820 G3 and Hitachi Virtual Storage Platform One Block 26.

Database configuration

The following table lists the SQL Server 2022 database parameters used in the benchmark tests.

SQL Server Database Parameter	Value
Version	Microsoft SQL Server 2022 (RTM) - 16.0.1000.6 (X64)
SQL Server Database size	3 TB (For OLAP)
	1 TB (For OLTP)
Trace Flags on SQL Server	OLTP → 652, 661, 3502, 8744 OLAP → 3502
Configuration settings on SQL Server	Lightweight Pooling
	Min Server Memory, Max Server Memory
	Recovery Interval
	Priority Boost
	Optimize for ad hoc workloads
	Cost Threshold for Parallelism
	Max Degree of Parallelism
	Max Worker Threads

Test environment

The following table lists the test environment details for one database instance.

Item	Description	Value
Operating system	-	Windows Server 2022 Datacenter
Workload type	-	OLTP/OLAP
Database size	-	1 TB(OLTP) / 3 TB (OLAP)
Number of physical cores	CPU cores	2 × 76 cores per socket (152 logical CPUs)
Memory	Server memory	768 GB per server
Number of SUT Servers	Profiles	2
Number of Users	Client Connections	1-256
Network	Cluster interconnect	2 × 25 Gbps NIC Teaming

Test methodology

The performance testing was conducted using HammerDB, a widely recognized benchmarking tool that aligns with industry standards. Derived from the TPC specifications, HammerDB can execute both transactional and analytical workloads to assess system performance. Specifically, the TPROC-C (OLTP) workload was used for transactional tests, while TPROC-H (OLAP) was employed for decision support workload evaluation.

The testing methodology serves the following key objectives:

- Performance Validation: Ensuring the solution meets the necessary quality assurance standards for optimal performance.
- Infrastructure Evaluation: Comparing various infrastructure technologies and solutions based on performance and cost efficiency (price/performance ratio).
- Performance Enhancement: Identifying opportunities for improving system efficiency and scalability.

Conclusion

The performance of applications in the database environment is very critical. The storage performance is also critical. In this testing of OLTP workloads, the VSP B26 can achieve 344K IOPS with the HammerDB benchmark kit, and this made the mission critical database environment benefit from the high storage performance.

We performed comprehensive database validation tests on a Hitachi Integrated Systems platform comprising the Hitachi Advanced Server HA820 G3 as the compute node and the Hitachi Virtual Storage Platform One Block 26 as the storage solution. The testing confirmed that all Microsoft SQL Server database functionalities were executed seamlessly without any issues. A range of database operations was validated under various workloads using the HammerDB performance benchmarking tool, with all workload types achieving optimal performance results.

Additional tests were conducted by running multiple workloads in parallel over extended periods, demonstrating strong performance for both OLTP and OLAP transactions. These results highlight that Hitachi Vantara hardware, including storage, processors, and networking components, provides a robust and high-performance infrastructure, ensuring reliable database operations and consistent performance even under heavy workloads.

The Hitachi Integrated Systems platform, using the combination of the Hitachi Virtual Storage Platform One Block 26 and the Hitachi Advanced Server HA820 G3 offers a powerful and stable environment for handling intensive transactions and workloads. The system maintains consistent memory and CPU resource utilization in a non-virtualized environment, ensuring high availability and reliability throughout the testing phase.

Note: Contact our Hitachi sales and engineering team for more details about performance results and best configuration practices. https://support.hitachivantara.com/en_us/contact-us.html.

Product descriptions

These products are used in this solution.

Hitachi Integrated Systems Platform

The Hitachi Integrated Systems platform is a high-performance, low-latency, integrated, converged solution using Hitachi Virtual Storage Platform One Block storage, Hitachi Advanced Server HA820 G3, as well as HA810 G3 with Sapphire Rapids Scalable Processors.

Hitachi Virtual Storage Platform One Block

The Hitachi Virtual Storage Platform One Block series simplifies system setup and management through the new VSP 360 management offering. Dynamic Drive Protection reduces RAID complexity, and always-on compression and deduplication enhance simplicity.

Dynamic Carbon Reduction optimizes energy usage by switching CPUs to ECO mode during low activity. Adaptive Data Reduction (ADR) is always on, enhancing efficiency and reducing the overall CO2 footprint.

Thin Image Advanced (TIA) integrates with major snapshot ecosystems, prioritizing security by defending against threats and ensuring data confidentiality. CyberArk Privileged Access Manager plugins enhance block storage system security by prioritizing data confidentiality, ensuring compliance, and actively defending against security threats.

Hitachi Virtual Storage Platform One Block includes the following 3 dedicated models:

- VSP One Block 24 – 256 GB Cache + SW Advanced Data Reduction (ADR) + 24 cores
- VSP One Block 26 – 768 GB Cache + 2x Compression Accelerator Module (CAM) + 24 cores
- VSP One Block 28 – 1 TB Cache + 4x CAM + 64 cores

All have the same drive count (72 NVMe flash drives, the appliance, and 2 × media trays) and they support Fibre Channel, iSCSI, and NVMe TCP connectivity. The new capabilities remove complexity such as data reduction always being on, Dynamic Drive Protection removes complicated RAID setup, and Dynamic Carbon Reduction delivers real-world reduction in power consumption. In addition, the models are FIPS compliant.

In short, the Hitachi Virtual Storage Platform One Block series combines simplicity, sustainability, and robust security features to optimize system management, energy efficiency, and data protection.

See https://www.hitachivantara.com/en-us/products/storage-platforms/block-storage/midrange/vsp-one-block for more information.

Hitachi Advanced Server HA820 G3

Hitachi Advanced Server HA820 G3 is a high-performance two-socket rackmount server designed for optimal performance and power efficiency. This allows owners to upgrade computing performance without overextending power consumption and offers non-latency support to virtualization environments that require maximum memory capacity. Hitachi Advanced Server HA820 G3 provides flexible I/O scalability for today's diverse data center application requirements.

Windows Server 2022 Datacenter Edition operating system

Windows Server 2022 is the latest server operating system from Microsoft, designed to meet the needs of modern data centers, hybrid cloud environments, and business-critical workloads. The Datacenter edition provides the highest level of performance, scalability, and security, making it ideal for large-scale virtualized and cloud-based environments. It builds upon the features of Windows Server 2019 and offers a wide array of new features that enhance security, networking, and management.

Microsoft SQL Server 2022 Enterprise Edition – Database Server

SQL Server 2022 Enterprise Edition builds on its legacy of high-end data management capabilities by delivering cutting-edge performance, security, and scalability for mission-critical workloads. With enhanced hybrid cloud integration and advancements in data security and business intelligence, SQL Server 2022 enables organizations to unlock new insights from their data, whether on-premises or in the cloud.

Choice of Language and Platform — You can run SQL Server 2022 from anywhere using Windows, Linux, and Kubernetes

Unparalleled high availability — Gain mission-critical uptime, fast failover, and improved disaster recovery by Availability groups.

Unlimited Virtualization — With SQL Server 2022, users gain access to unlimited virtualization rights, enabling them to run an unlimited number of virtual instances of SQL Server, making it ideal for large-scale virtualized environments.

Brocade switches from Broadcom

Brocade and Hitachi Vantara have partnered to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

Brocade Fibre Channel switches deliver industry-leading performance with seventh generation 64Gb/sec Fibre Channel interfaces, simplifying scale-out network architectures. Get the high-performance, availability, ease of management, and support for the next generation of Hitachi Virtual Storage Platform storage systems on a solid storage network foundation that can grow as your need grows.

See https://www.broadcom.com/products/fibre-channel-networking/switches for more information.

Cisco Nexus switches

The Cisco Nexus switch product line offers a range of solutions that simplify the connection and management of disparate data center resources through software-defined networking (SDN). Leveraging the Cisco Unified Fabric, which unifies storage, data, and networking (Ethernet/IP) services, the Nexus switches create an open, programmable network foundation built to support a virtualized data center environment.

HammerDB

HammerDB is the most popular, advanced, open-source, easy to use performance benchmarking tool available in the market. Source code for benchmarking workloads is based on the TPC specification. HammerDB can be used to run OLTP and OLAP workloads. In HammerDB, an OLAP workload is called TPROC-H which is derived from the TPC-H specification and an OLTP workload is called TPROC-C which is derived from the TPC-C specification.

TPROC-H Benchmarking is a decision support benchmark. It consists of a business-oriented standard set of 22 complex queries running against a database. The queries in this benchmark are designed in a way that it scans a large volume data with a high degree of complexity such as table joins, data sorting, and use of arithmetic functions. To increase the load on a system for capacity testing we run workloads using concurrent users or run workloads with additional parallel threads.

TPROC-C Benchmarking is an online transaction processing benchmark. It consists of queries or transactions running to the database. The transactions running in this workload have a combination of select, insert, update or delete records on tables. We can execute transactions from multiple users and the result should be minimal response time. We get two values— one is new orders per minute (NOPM) and the other is transactions per minute (TPM).

For more information, see https://www.hammerdb.com/.

Note: HammerDB benchmark workloads are derived from the TPC specification, but results should not be compared with official TPC benchmark results.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact