

Microsoft SQL Server 2022 Deployment on Cisco UCS X210c M6/M7 for Bare-Metal and Hybrid Cloud

June 2023

Contents

Overview	3
Scope of this document	4
Cisco UCS X-Series platform for SQL Server databases	4
Solution overview	5
Deploying hardware and software	15
Deploying highly available SQL Server databases using Always On AG	20
Disaster recovery with Azure SQL Managed Instance	21
Accelerating Microsoft SQL Server backup compression performance with Intel integrated Accelerator Engines	27
Summary	29
For more information	30

Overview

Microsoft SQL Server is one of the most widely adapted database platforms by different customer segments. SQL Server instances are simple to deploy and flexible, and they have evolved over years, offering a rich set of features for achieving higher transactional throughput, higher availability, and disaster recovery of the databases. Many organizations deploy SQL Server databases on bare-metal servers for performance-critical databases because this type of deployment offers cost-effective, high-performing, and simple solution. Also, bare-metal database deployments will benefit from dedicated compute, memory, and low-latency local storage performance. SQL Server 2022 is the latest offering from Microsoft, and it has many new features and enhancements, now delivering industry-leading performance, seamless analytics over on-premises operational data, and business continuity through Azure.

The Cisco Unified Computing System™ (Cisco UCS®) X-Series is a new modular computing system, configured and managed from the cloud. It is designed to meet the needs of modern applications and improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design. The Cisco UCS X-Series is equipped with the latest-generation Intel® Xeon™ scalable processors, and it supports Double Data Rate 4 and 5 (DDR4 and 5) memory modules only or a mixture with Intel Optane Persistent Memory. The chassis is designed to support up to 200-Gbps network bandwidth per compute node using 5th-generation fabric interconnects. The Cisco Intersight® platform is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support.

The Cisco UCS X210c Compute Node is the first computing device to integrate into the Cisco UCS X-Series Modular System. Up to eight compute nodes can reside in the 7-Rack-Unit (7RU) Cisco UCS X9508 Chassis, offering one of the highest densities of compute, Input/Output (I/O), and storage per rack unit in the industry. Bringing a high-performing storage tier closer to the compute node opens up new deployment types, thereby improving data center efficiencies. The combination of Intel Xeon scalable processors, support for up to 8-TB DDR4/DDR5 memory only or 12-TB memory with Intel Optane, and up to six Nonvolatile Memory Express (NVMe) or Serial Attached SCSI (SAS) solid-state drives or a mix of both makes the Cisco UCS X210c the perfect choice for running performance-critical enterprise databases such as Microsoft SQL Server.

This document discusses a reference architecture that illustrates the benefits of using a Cisco UCS X210c M6 and M7 Compute Node for bare-metal deployments, taking advantage of the local storage of the blade. It also provides Cisco UCS configuration best practices for deploying high-performing, highly available, and Disaster Recovery (DR)-capable Microsoft SQL Server transactional databases to use both on-premises and cloud SQL instances. The document also discusses the benefits of using UCS X210c M7 powered by Intel 4th generation scalable processors and integrated accelerators for achieving better ROI by offloading Microsoft SQL Server backup compression operations.

Scope of this document

This document discusses Cisco UCS deployment and specific configuration best practices for hosting performance-critical operational Microsoft SQL Server databases using Cisco UCS X210c M6/M7 blades with local attached NVMe and SAS Solid-State Drives (SSDs). It presents recommendations for deploying SQL Server Always On Availability Groups for achieving highly available databases with automatic failover capabilities with in an on-premises datacenter. It covers DR scenarios of the on-premises databases, where an Azure SQL Managed Instance (SQL MI) link feature is used to replicate an on-premises database to the Azure Cloud. It also discusses the performance benefits of offloading SQL Server backup compression to Intel Integrated accelerators.

Note that SQL Server 2022 Highly Available (HA) and Disaster Recovery (DR) options are applicable to both UCS X210c M6 and M7 compute nodes featured with the latest Intel Xeon Scalable processors and have similar local storage configurations. However, Intel Integrated accelerators are supported only with Intel 4th generation Scalable processors. Therefore, the UCS X210c M7 compute node is used in order to leverage the Intel Quick Assist Technology (QAT) accelerator for offloading SQL Server backup compression.

Cisco UCS X-Series platform for SQL Server databases

This section highlights a few innovative capabilities that the Cisco UCS X-Series platform and UCS X210c M6/M7 blade offer for performance-critical databases.

- You can simplify and standardize SQL Server deployment with a stateless and programmable computing platform: Cisco UCS provides a stateless and programmable computing platform that envisions servers as resources whose identity, configuration, and connectivity you can manage through software rather than the tedious, time-consuming, error-prone manual processes of the day. You can dynamically create Cisco UCS server profiles that consist of critical server information such as network, storage, boot order, virtual LANs (VLANs), and so on, and associate them to any physical server within a few minutes rather than hours. It facilitates rapid bare-metal provisioning and replacement of failed servers by simply migrating service profiles among servers, thereby greatly reducing the application downtimes hosted on these servers. You can define once and use multiple times the nature of the Cisco UCS Server templates and policies, enabling your customers to standardize SQL Server database deployments across their organization with the same server configuration and firmware versions.
- You can manage the entire infrastructure using a centralized cloud-based platform, called Cisco Intersight. The Cisco Intersight cloud-operations platform allows you to handle full lifecycle management of on-premises infrastructure; remote, branch-office, and edge locations; and the public cloud. From a single cloud-based interface, you can consistently manage the entire infrastructure, no matter where it resides.

- You can achieve blazing IO performance using ultra-latency nonvolatile memory express (NVMe) disks of UCS X210c M6/M7 Compute Node. UCS X210c M6/M7 supports up to 6 high performing NVMe drives or a mixture of up to six SATA/SAS or NVMe drives. The NVMe drives are directly connected to CPUs over 16 high speed PCIe generation 4 buses supporting millions of IO operations per second (IOPS) at sub milli second latency. Using six 15.3TB NVMe disks, a single X210 M6/M7 Compute Node can support up to 91.8TB extreme performing NVMe storage. Using six 7.6TB SAS SSD disks, a single X210 M6/M7 Compute Node can support up to 45.6 TB SAS SSD storage. Traditional Mega RAID (Redundant Array of Independent Disks) controller can be used to create RAID on the SSD drives. Optionally, you can create RAID on the NVMe disks as well using technologies like [Intel vROC](#) (virtual Raid On Chip). The NVMe storage is used for storing data files while traditional SAS SSDs are used for storing transaction log files.

Solution overview

This solution for the SQL Server 2022 databases is validated with the Cisco UCS X-Series, and the Cisco Intersight platform delivers cloud-managed infrastructure on the latest supported Cisco UCS hardware. The Cisco Intersight cloud-management platform configures and manages the infrastructure. This section covers the solution requirements and design details.

Reference architecture

The solution is validated using a Cisco UCS X9508 chassis equipped with Cisco UCS 9108-25G Intelligent Fabric Modules (IFMs) and UCS X210c M6/M7 Compute Nodes in the on-premises datacenter. Cisco UCS X210c M6/M7 nodes are configured with latest Intel scalable processors and 512-GB DDR4/DDR5 memory. Two M.2 SATA SSDs are used in a RAID 1 configuration for the OS boot drive. Four high-performing NVMe drives store SQL Server data files, and two SAS SSDs store database transaction log files. RAID 1 is configured on two SAS SSD drives using Mega RAID controller. In this example deployment, two X210c M6/M7 Compute Nodes host stand-alone SQL Server instances. Cisco recommends that you install the compute nodes on two different Cisco UCS X9508 chassis to achieve chassis-level redundancy.

A Windows Failover Cluster is deployed and configured on the two Cisco UCS X210c blades. Installed standalone SQL Server 2022 instances are on each blade, and they are configured with Always On availability Groups (AG) for providing high availability for a set of user databases. The two SQL Server instances (replicas) are configured with synchronous-commit mode replication with automatic failover such that the databases are always synchronized and both replicas have the same copy of the database. One instance is configured as a primary replica serving read-write transactions, and you can optionally configure the other replica to serve read-only transactions. If the primary replica is not available for any reason (due to hardware or software failures), AG will initiate the failover to the secondary replica and the secondary replica will become the primary replica and start serving read-write transactions. The former primary replica then becomes the secondary and tries to catch up with the new primary replica when it comes back online. You can offload read-only and analytical workloads to the secondary replica within the on-premises datacenter.

SQL Server 2022 introduced the “Azure Managed Instance Link” feature, which enables you to replicate an on-premises database Azure SQL Managed Instance using Distributed Availability Group (DAG) technology. The SQL MI replica running in the cloud receives the data changes from the on-premises database and supports only manual failover from the on-premises database to the Azure cloud. You can take advantage of the SQL MI running in the cloud to offload the read-only and analytical workloads from the primary replica. If a complete disaster occurs in the on-premises datacenter, the administrator can manually move the DAG from the on-premises database to the cloud and start using the SQL MI replica for read-write workloads.

Figure 1 shows the overall architecture of the solution validated using two X210c M6/M7 compute nodes for highly available and DR-capable SQL Server 2022 deployment.

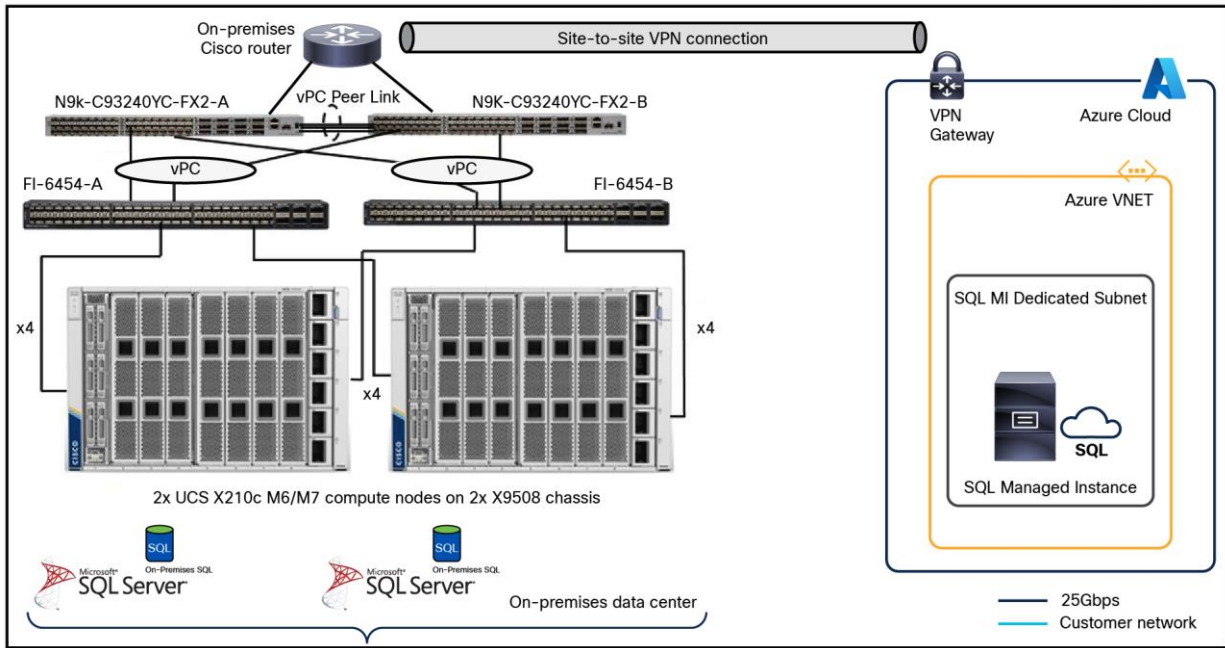


Figure 1. Reference architecture for highly available SQL Server 2022 deployment using Cisco UCS X210c M6/M7 compute node

Figure 2 shows the logical diagram depicting the implementation of the SQL Server Always On AG for highly available databases within the on-premises datacenter and DAG for disaster recovery between the on-premises and Azure cloud using the SQL MI Link feature.

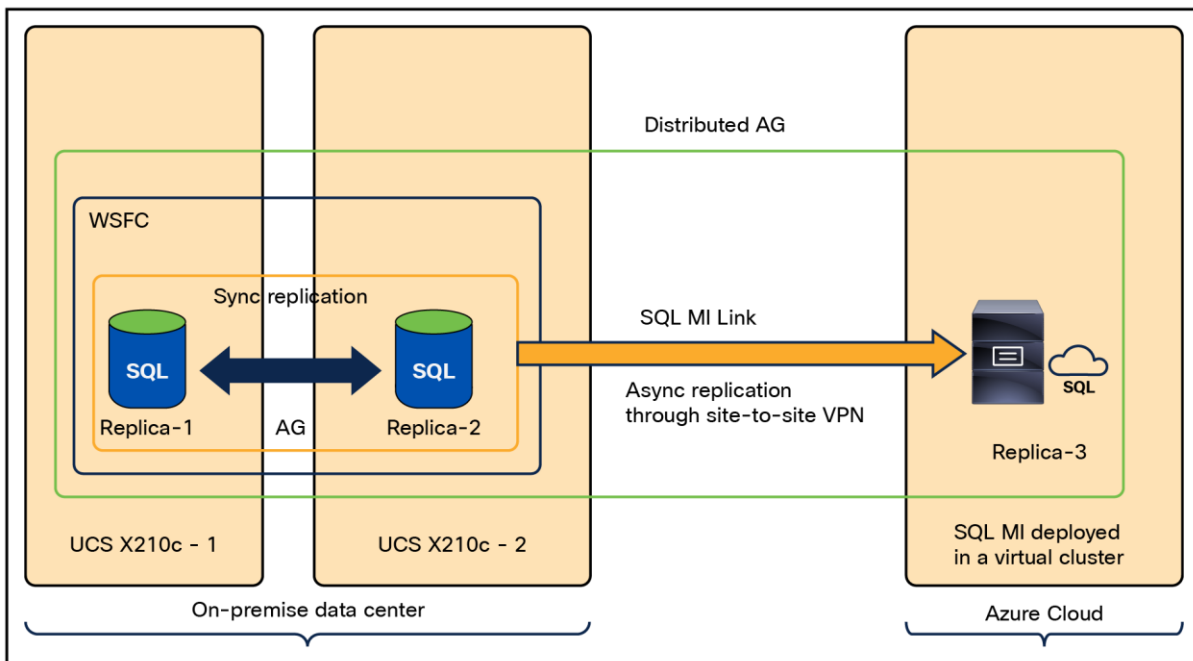


Figure 2. SQL Server Always On AG and Distributed AG implementation for HA and DR

Table 1 lists the software and hardware components used along with their versions.

Table 1. Hardware and software components used for HA/DR validation

Component	Model	Software version
Management	2x Cisco UCS 6454 Fabric Interconnects	4.2(1h)
Compute	2x Cisco UCS X210c M6 or M7 Blade Servers	5.(4a) / 5.1(1b)
Processor	Intel Xeon 3 rd generation scalable processors	
Memory	512GB (16x 32GB) memory	
Adapters	Cisco UCS Virtual Interface Card (VIC) 14425 or VIC 15420	Driver: 7.716.2.0 / 5.12.9.1
Storage	4 x 3.2-TB 2.5-in. U.2 Intel P5600 NVMe high-performance, medium-endurance drives for database data files 2 x 480-GB SATA SSD drives for database T-Log files 2 x 240-GB M.2 Serial Advanced Technology Attachment (SATA) SSDs and associated RAID controllers for OS boot drive	
Cisco software	Cisco Intersight software	
Software	Windows Server 2022 (Datacenter Edition) Intel vROC drivers: (required only when vROC is used) SQL Server 2022 Evaluation Edition	

The following sections provides more information about the hardware and software components used in the solution.

Cisco UCS X-Series Modular System

The Cisco UCS X-Series Modular System begins with the Cisco UCS X9508 Chassis (Figure 1), engineered to be adaptable and future-ready. The X-Series is a standards-based open system designed to be deployed and automated quickly in a hybrid cloud environment.

With a midplane-free design, I/O connectivity for the X9508 chassis is accomplished with front-loading vertically oriented computing nodes that intersect with horizontally oriented I/O connectivity modules in the rear of the chassis. A unified Ethernet fabric is supplied with the Cisco UCS 9108 IFMs. In the future, Cisco UCS X-Fabric Technology interconnects will supply other industry-standard protocols as standards emerge. You can easily update interconnections with new modules.

The Cisco UCS X-Series is powered by [Cisco Intersight software](#), so it is simple to deploy and manage at scale.

The Cisco UCS X9508 Chassis (Figure 3) provides the following features and benefits:

- The 7RU chassis has 8 front-facing flexible slots. These slots can house a combination of computing nodes and a pool of future I/O resources, which may include Graphics Processing Unit (GPU) accelerators, disk storage, and nonvolatile memory.
- Two Cisco UCS 9108 IFMs at the top of the chassis connect the chassis to upstream Cisco UCS 6400 Series Fabric Interconnects (FIs). Each IFM offers these features:
 - The module provides up to 100 Gbps of unified fabric connectivity per computing node.
 - The module provides eight 25-Gbps Small Form-Factor Pluggable 28 (SFP28) uplink ports.
 - The unified fabric carries management traffic to the Cisco Intersight cloud-operations platform, Fibre Channel over Ethernet (FCoE) traffic, and production Ethernet traffic to the fabric interconnects.
- At the bottom of the chassis are slots used to house UCS X9416 X-Fabric Modules which enables GPU connectivity to the UCS X210c Compute Nodes.
- Six 2800-watt (W) Power Supply Units (PSUs) provide 54 Volts (V) of power to the chassis with N, N+1, and N+N redundancy. A higher voltage allows efficient power delivery with less copper wiring needed and reduced power loss.
- Efficient, 4 x 100-mm, dual counter-rotating fans deliver industry-leading airflow and power efficiency. Optimized thermal algorithms enable different cooling modes to best support the network environment. Cooling is modular, so future enhancements can potentially handle open- or closed-loop liquid cooling to support even higher-power processors.



Figure 3.
Cisco UCS 9508 X-Series Chassis, front (left) and back (right)

Since Cisco first delivered the Cisco Unified Computing System in 2009, our goal has been to simplify the datacenter. We pulled management out of servers and into the network. We simplified multiple networks into a single unified fabric. And we eliminated network layers in favor of a flat topology wrapped into a single unified system. With the Cisco UCS X-Series Modular System, the simplicity is extended even further:

- Simplify with cloud-operated infrastructure. We move management from the network into the cloud so that you can respond at the speed and scale of your business and manage all your infrastructure.

You can shape Cisco UCS X-Series Modular System resources to workload requirements with the Cisco Intersight cloud-operations platform. You can integrate third-party devices, including storage from NetApp, Pure Storage, and Hitachi. In addition, you gain intelligent visualization, optimization, and orchestration for all your applications and infrastructure.

- Simplify with an adaptable system designed for modern applications. Today's cloud-native, hybrid applications are inherently unpredictable. They are deployed and redeployed as part of an iterative DevOps practice. Requirements change often, and you need a system that doesn't lock you into one set of resources when you find that you need a different set. For hybrid applications, and for a range of traditional datacenter applications, you can consolidate your resources on a single platform that combines the density and efficiency of blade servers with the expandability of rack servers. The result is better performance, automation, and efficiency.
- Simplify with a system engineered for the future. Embrace emerging technology and reduce risk with a modular system designed to support future generations of processors, storage, nonvolatile memory, accelerators, and interconnects. Gone is the need to purchase, configure, maintain, power, and cool discrete management modules and servers. Cloud-based management is kept up-to-date automatically with a constant stream of new capabilities delivered by the Cisco Intersight Software-as-a-Service (SaaS) model.
- Support a broader range of workloads. A single server type supporting a broader range of workloads means fewer different products to support, reduced training costs, and increased flexibility.

Cisco UCS X210c Series Servers

The Cisco UCS X-Series Modular System simplifies your datacenter, adapting to the unpredictable needs of modern applications while also accommodating traditional scale-out and enterprise workloads. It reduces the number of server types that you need to maintain, helping to improve operational efficiency and agility by reducing complexity. Powered by the Cisco Intersight cloud-operations platform, it shifts your thinking from administrative details to business outcomes with hybrid cloud infrastructure that is assembled from the cloud, shaped to your workloads, and continuously optimized.

The Cisco UCS X210c M6 Compute Node is the first computing device integrated into the Cisco UCS X-Series Modular System (Figure 4). Up to 8 computing nodes can reside in the 7RU Cisco UCS X9508 Chassis, offering one of the highest densities of computing, I/O, and storage resources per rack unit in the industry. The Cisco UCS X210c M6/M7 harnesses the power of the latest Intel Xeon scalable processors. It includes the following features:

- CPU: Install up to two 3rd or 4th generation Intel Xeon scalable processors with up to 40/60 cores per processor and 1.5/2.6 MB of Level 3 cache per core.
- Memory: Install up to thirty-two 256-GB DDR4 3200/4800-MHz Dual Inline Memory Modules (DIMMs) for up to 8 TB of main memory. Configuring up to sixteen 512-GB Intel Optane persistent-memory DIMMs can yield up to 12 TB of memory.
- Storage: Install up to six hot-pluggable SSDs or NVMe 2.5-inch drives with a choice of enterprise-class RAID or pass-through controllers with four lanes each of PCIe generation 4 connectivity and up to two M.2 SATA drives for flexible boot and local storage capabilities.

- Modular LAN-on-motherboard (mLOM) Virtual Interface Card (VIC): The Cisco UCS VIC 14425 occupies the mLOM slot on the server, enabling up to 50-Gbps unified fabric connectivity to each of the chassis IFMs for 100-Gbps connectivity per UCS X210c M6 server.

Similarly, Cisco UCS 5th generation VIC 15420 occupies the mLOM slot on the server, enabling up to 50-Gbps unified fabric connectivity to each of the chassis IFMs for 100-Gbps connectivity per UCS X210c M7 server. Concurrently, UCS VIC 15231 occupies the mLOM slot on the server, enabling up to 100-Gbps unified fabric connectivity to each of the chassis IFMs for 200-Gbps connectivity per UCS X210c M7 server.

- Optional mezzanine VIC: The Cisco UCS VIC 14825 can occupy the mezzanine slot on the server at the bottom rear of the chassis. The I/O connectors of this card link to Cisco UCS X-Fabric Technology that is planned for future I/O expansion. An included bridge card extends the two 50-Gbps network connections of this VIC through IFM connectors, bringing the total bandwidth to 100 Gbps per fabric (for a total of 200 Gbps per UCS X210c M6 server). Similarly, Cisco UCS 5th generation VIC 15422 can occupy the mezzanine slot on the server at the bottom rear of the chassis. The I/O connectors of this card link to Cisco UCS X-Fabric Technology that is planned for future I/O expansion. An included bridge card extends the two 50-Gbps network connections of this VIC through IFM connectors, bringing the total bandwidth to 100 Gbps per fabric (for a total of 200 Gbps per UCS X210c M7 server).
- Security: The server supports an optional Trusted Platform Module (TPM). Additional features include a secure boot Field-Programmable Gateway (FPGA) and Anti-Counterfeit Technology 2 (ACT2) provisions.




Figure 4.
Front view of X210c M6/M7 Compute Node

A specifications sheet for the X210c M6 compute node is available at:
<https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/x210c-specsheet.pdf>.

A specifications sheet for the X210c M7 compute node is available at:
<https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/x210cm7-specsheet.pdf>

Cisco UCS VIC 15231

Cisco UCS X210c M6 and M7 compute nodes support 5th-generation Cisco UCS VIC 15231 along with 4th-generation VICs (14425 and 14825). UCS VIC 15231 fits in to the mLOM slot in the Cisco X210c Compute Node and enables up to 100 Gbps of unified fabric connectivity to each of the chassis IFMs, for a total of 200 Gbps of connectivity per server. Cisco UCS VIC 15231 connectivity to the IFM 9108 and to the fabric interconnects is delivered through 2x 100-Gbps connections. The Cisco UCS VIC 15231 supports 256 virtual interfaces (both Fibre Channel and Ethernet) along with the latest networking innovations such as Non-Volatile Memory Express over RDMA (ROCEv2), Virtual Extensible LAN (VXLAN), network virtualization generic routing encapsulation (NVGRE) offload, and so on. Refer to figure 5.

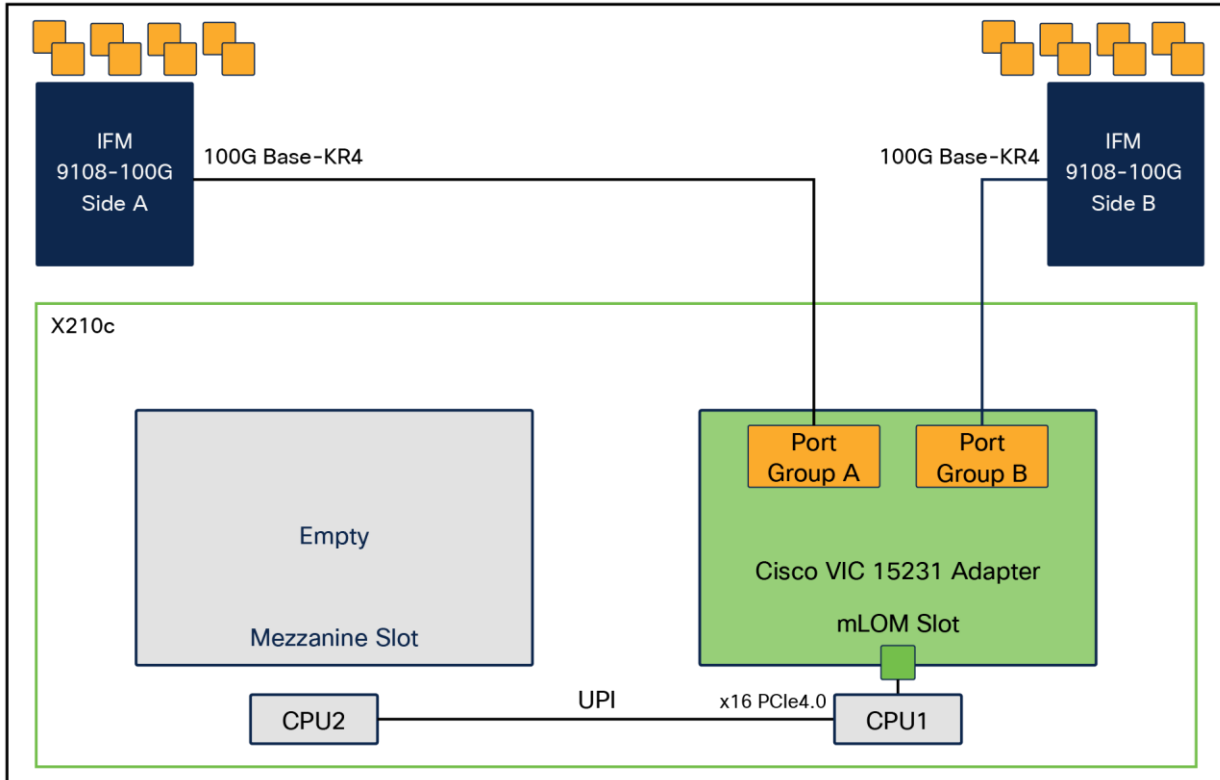


Figure 5.
Single Cisco VIC 15231 in Cisco UCS X210c

Cisco UCS VIC 14425

For this validation, Cisco UCS VIC 14425 is used; it fits in to the mLOM slot in the Cisco X210c Compute Node and enables up to 50 Gbps of unified fabric connectivity to each of the chassis IFMs, for a total of 100 Gbps of connectivity per server. Cisco UCS VIC 14425 connectivity to the IFM and to the fabric interconnects is delivered through four 25-Gbps connections, which are configured automatically as two 50-Gbps port channels. The Cisco UCS VIC 14425 supports 256 virtual interfaces (both Fibre Channel and Ethernet) along with the latest networking innovations such as Non-Volatile Memory Express over Fabrics (NVMeoF) over RDMA (ROCEv2), Virtual Extensible LAN (VXLAN), Network Virtualization Generic Routing Encapsulation (NVGRE) offload, and so on. Refer to Figure 6.

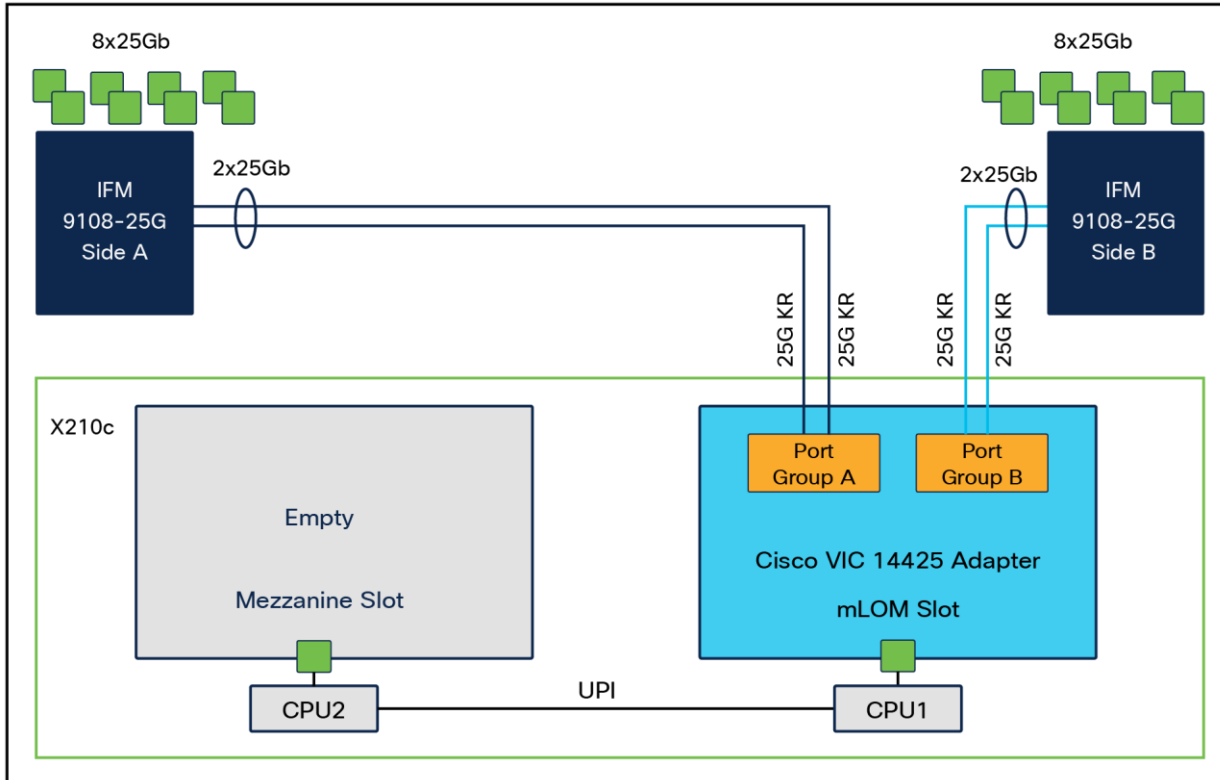


Figure 6.
Single Cisco VIC 14425 in Cisco UCS X210c

For additional network bandwidth, you can also use the Cisco VIC 14825; it fits the mezzanine slot on the server. A bridge card (UCSX-V4-BRIDGE) extends the 2x 50 Gbps of network connections of this VIC up to the mLOM slot and out through the IFM connectors of the mLOM, bringing the total bandwidth to 100 Gbps per fabric for a total bandwidth of 200 Gbps per server (Figure 7).

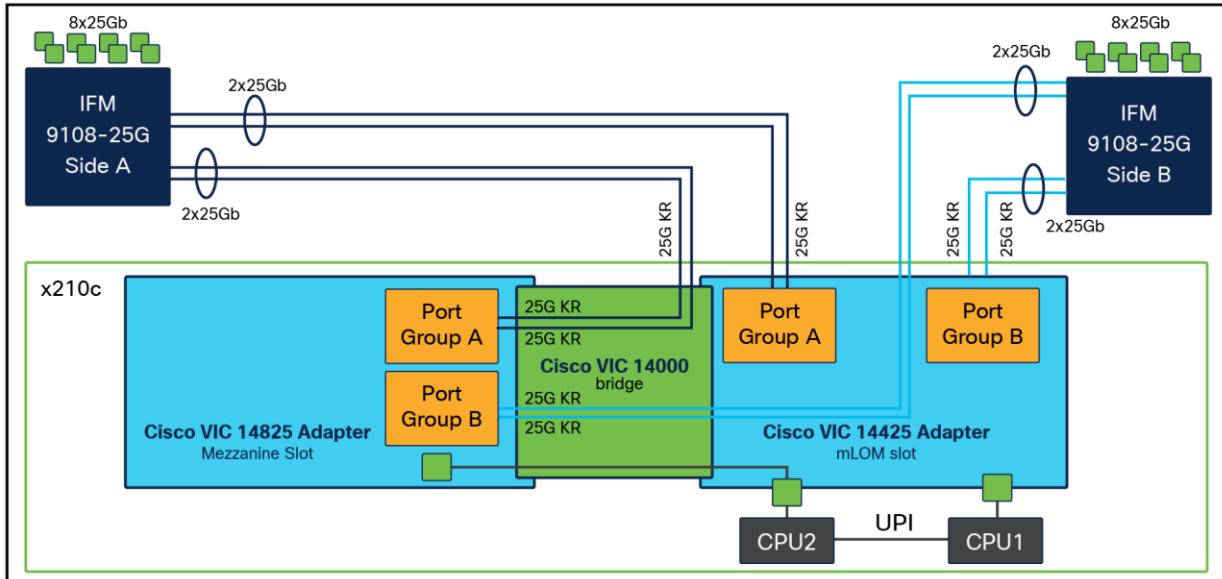


Figure 7.
Cisco VIC 14425 and 14825 in Cisco UCS X210c M6

Cisco UCS 6400 Series Fabric Interconnect

The Cisco UCS fabric interconnects provide a single point for connectivity and management for the entire Cisco UCS deployment (Figure 7). Typically deployed as an active-active pair, fabric interconnects of the system integrate all components into a single, highly available management domain that the Cisco UCS Manager or Cisco Intersight platform controls. Cisco UCS fabric interconnects provide a single unified fabric for the system, with low-latency, lossless, cut-through switching that supports LAN, storage-area network (SAN), and management traffic using a single set of cables.

Front view



Rear view



Figure 8.
Cisco UCS 6400 Series Fabric Interconnect

The Cisco UCS 6454 Fabric Interconnect used in the current design is a 54-port fabric interconnect. This 1RU device includes 28 x 10-/25-Gbps Ethernet ports, 4 x 1-/10-/25-Gbps Ethernet ports, 6 x 40-/100-Gbps Ethernet uplink ports, and 16 unified ports that can support 10-/25-Gigabit Ethernet or 8-/16-/32-Gbps Fibre Channel, depending on the SFP design.

Note: To support the Cisco UCS X-Series, the fabric interconnects must be configured in Cisco Intersight managed mode. This option replaces the local management with Cisco Intersight cloud- or appliance-based management.

Cisco Intersight platform

The Cisco Intersight platform is a SaaS infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support (Figure 9). This platform is designed to be modular, so you can adopt services based on your individual requirements. It significantly simplifies IT operations by bridging applications with infrastructure, providing visibility and management from bare-metal servers and hypervisors to serverless applications, thereby reducing costs and mitigating risk. This unified SaaS platform uses a unified Open Application Programming Interface (API) design that natively integrates with third-party platforms and tools.

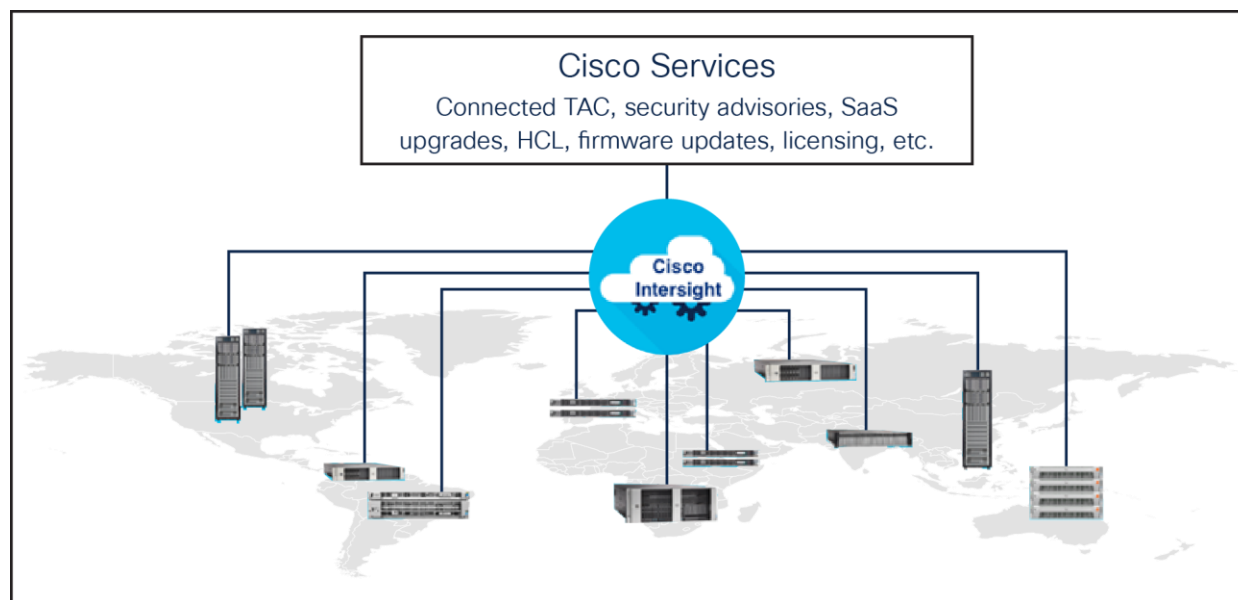


Figure 9.
Cisco Intersight overview

The main benefits of Cisco Intersight infrastructure services follow:

- Simplify daily operations by automating many daily manual tasks.
- Combine the convenience of a SaaS platform with the capability to connect from anywhere and manage infrastructure through a browser or mobile app.
- Stay ahead of problems and accelerate trouble resolution through advanced support capabilities.
- Gain global visibility of infrastructure health and status along with advanced management and support capabilities.
- Upgrade to add workload optimization and Kubernetes services when needed.

Microsoft SQL Server 2022

Microsoft SQL Server 2022 is the latest relational database engine from Microsoft. It offers many new features and enhancements to the relational and analytical engines and is available in both Linux and Windows versions. This latest version comes with many features in different areas such as analytics, security, availability, performance, management, and so on.

Azure SQL Managed Instance

Azure SQL Managed Instance is the intelligent, scalable cloud database service that combines the broadest SQL Server database engine compatibility with all the benefits of a fully managed and evergreen platform as a service. SQL Managed Instance has near 100-percent compatibility with the latest SQL Server (Enterprise Edition) database engine, providing a native [virtual network \(VNet\)](#) implementation that addresses common security concerns, and a [business model](#) favorable for existing SQL Server customers. For more details about architecture and network connectivity of SQL MI, please visit: <https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/managed-instance-link-feature-overview?view=azuresql>.

Link to Azure SQL Managed Instance

Link to Azure SQL MI is new capability that enables near real-time data replication from on-premises SQL Server to Azure SQL Managed Instance. The link provides hybrid flexibility and database mobility as it unlocks several scenarios, such as scaling read-only workloads, offloading analytics and reporting to Azure, and migrating to the cloud. And, with SQL Server 2022, the link feature enables disaster recovery. In this solution, this feature has been validated to replicate SQL Server databases installed on X210c M6/M7 blade servers to Azure SQL MI server.

Deploying hardware and software

- This section discusses some of the important Cisco UCS policies used for creating Cisco UCS Server profile templates to be consumed by Cisco UCS X210c M6/M7 blades for hosting SQL Server databases for a bare-metal deployment use case. Note that this document does not cover all the standard policies or pools used for creating the server profile templates.
- With a service profile template, you can create several server profiles quickly with the same configurations, such as the number of vNICs, local storage configuration, boot order, and so on, and with identity information drawn from the same pools.

Virtual Network Interface Cards (vNICs) and LAN connectivity policy

The following vNICs are created for different types of traffic and used in the LAN connectivity policy that is later used in the server profile template.

- Client-Mgmt: This vNIC is used for database client connectivity traffic.
- AG-Replication: This vNIC carries SQL Server Availability Group (AG) replication data traffic between the Cisco UCS X210c compute nodes located in an on-premises datacenter. SQL Server does not commit the transactions on the primary replica until the data changes are replicated to the secondary synchronous replicas.

Table 2 provides more details about the vNICs used for this validation.

Table 2. vNIC configuration details

Details	Client-Mgmt	AG-Replication
Placement - slot ID	MLOM	MLOM
Placement - switch ID	A	A
Placement - PCI order	0	1
Failover	Enabled	Enabled
CDN - source	vNIC name	vNIC name
Network group policy - VLANs	314	500
Network control policy	CDP and LLDP enabled	CDP and LLDP enabled
Ethernet Quality of Service (QoS)	Priority: Best-effort, MTU: 1500 bytes	Priority: Best-effort, MTU: 9000 bytes
Ethernet adapter policy	Pre-defined default Windows adapter policy	Pre-defined default Windows adapter policy

BIOS policy

BIOS policy plays an important role to achieve maximum performance from the server. Table 3 lists a few important settings used for this validation; these settings are important to achieve maximum performance from the system. The rest of the settings are set to platform-default.

Table 3. BIOS policy

BIOS token	BIOS token value
Boot options: VMD enablement (required for vROC)	Enabled
Intel directed I/O: Intel VT for directed I/O	Enabled
PCI: NVMe SSD hot-plug support	Enabled
Power and performance: L1 stream hardware prefetcher	Enabled
Power and performance: L2 stream hardware prefetcher	Enabled
Power and performance: Virtual NUMA	Disabled
Power and performance: XPT remote prefetcher	Enabled
Processor: Adjacent cache line prefetcher	Enabled
Processor: Boot performance mode	Maximum performance
Processor: Power technology	Performance
Processor: Intel HyperThreading mode	Enabled
Processor: Intel Turbo Boot Tech	Enabled
Processor: LLC prefetch	Enabled
Processor: Package C state limit	C0 C1 state
Processor: C1E, C3, and C6 report and CPU C state	Disabled
Processor: P-STATE coordination	HW ALL
Processor: Power performance tuning	OS
Processor: UPI link frequency select	Auto
Processor: Sub NUMA clustering	Disabled
Processor: XPT prefetch	Enabled

Note: Cisco requires that you thoroughly test, validate, and analyze any changes in the BIOS settings in the test or development environments before you implement them in the production environments.

Storage policy for database T-LOG files

For this validation, local attached high-performing SSD drives are used for storing SQL Server Transaction log (T-LOG) files. A storage policy is used for implementing RAID 1 configuration on the SAS SSD drives installed in the disk slots 5 and 6 that are controlled by a Mega RAID controller. This RAID 1 volume will be used for storing the T-LOG files. Table 4 provides more details about the storage policy used for creating RAID 10 volume.

Table 4. Storage policy

Parameter	Values
Policy name	RAID1_Slots_5to6
Disk group name	SQLLOG-DG
Number of spans	1
Disk array spans	Span 0: 5-6
Virtual drive name	SQLLOG
Expand to available	Yes
Stripe size	64KiB
Write policy	Write through

Note that two M.2 SSD drives are used in the RAID 1 configuration for the OS boot drive.

The following screenshot shows the local disk policy used for this validation.

The screenshot shows the 'Create' page for a storage policy. The 'Policy Details' tab is selected. Under 'M.2 RAID Configuration', the 'Slot of the M.2 RAID controller for virtual drive creation' is set to 'MSTOR-RAID-1 (MSTOR-RAID)'. Under 'MRAID/RAID Controller Configuration', 'Global Hot Spares' is set to 0. There are two tables:

Drive Group Name	RAID Level	Number of Spans	Dedicated Hot Spares	Drive Array Spans
SQLLOG-DG	RAID1			{5-6}

Virtual Drive Name	Drive Group	Size (MiB)	Expand to Available	Set as Boot Drive
SQLLOG	SQLLOG-DG	-	Yes	No

Note that four NVMe drives installed in slots 1 to 4 are directly connected to CPUs and can be controlled from a BIOS or vROC tool installed in the operating system.

After you create the Cisco UCS Server Profile template with required policies and pools, you can clone as many server profiles as you need and associate them to the Cisco UCS X210c M6/M7 servers. After you associate the server profile to a Cisco UCS X210c M6/M7 blade, you can power on the servers.

Updating Cisco VIC and storage drivers

Windows 2022 Server does not have Cisco and storage drivers natively. After installing the Windows Operating System, download the latest Windows drivers for the Cisco UCS X210c M6/M7 server from the Cisco website and update the Cisco VIC and storage drivers. Refer to the following link for detailed steps for downloading and updating drivers:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vic_drivers/install/Windows/b_Cisco_VIC_Drivers_for_Windows_Installation_Guide.pdf.

After updating the Windows drivers, the network interfaces and storage devices will be discovered and shown in the operating system. Assign the IP addresses to the network interfaces and initialize, partition and create the volumes for storing database data and log files.

Updating Windows power setting

For maximum performance from the system, ensure that the power policy of the Windows server is set to “High performance”:

```
PS C:\Users\Administrator.UCSTME>
PS C:\Users\Administrator.UCSTME> powercfg.exe /setactive 8c5e7fda-e8bf-4a96-9a85-a6e23a8c635c
PS C:\Users\Administrator.UCSTME> powercfg /list

Existing Power Schemes (* Active)
-----
Power Scheme GUID: 381b4222-f694-41f0-9685-ff5bb260df2e (Balanced)
Power Scheme GUID: 8c5e7fda-e8bf-4a96-9a85-a6e23a8c635c (High performance) *
Power Scheme GUID: a1841308-3541-4fab-bc81-f71556f20b4a (Power saver)
PS C:\Users\Administrator.UCSTME>
```

Setting MTU for AG network at Windows OS level

As explained earlier in the vNIC policy, each Cisco UCS X210c blade is configured with two networks. The client-Mgmt interface is by clients to connect to the database, and the AG-Replication interface is for Availability Group replication traffic. For better network throughput and latencies for replication traffic between the two Cisco UCS X210c blades, set the Maximum Transmission Unit (MTU) or Jumbo Frames to 9000 on both the hosts on AG-Replication interfaces and verify that the large packets can be transmitted without breaking them. The following screenshot shows how to verify Jumbo Frames:

```
PS C:\Users\Administrator.UCSTME> Get-NetAdapter

Name                InterfaceDescription      ifIndex Status      MacAddress      LinkSpeed
----                -
AG-Replication      Cisco VIC Ethernet Interface 6 Up         00-25-B5-AA-01-0F 50 Gbps
Client-Mgmt        Cisco VIC Ethernet Interface #2 4 Up         00-25-B5-AA-01-10 50 Gbps

PS C:\Users\Administrator.UCSTME> Set-NetAdapterAdvancedProperty -Name AG-Replication -RegistryKeyword *JumboPacket -RegistryValue 9014
PS C:\Users\Administrator.UCSTME>
PS C:\Users\Administrator.UCSTME> ping 192.168.100.31 -S 192.168.100.32 -f -l 8972

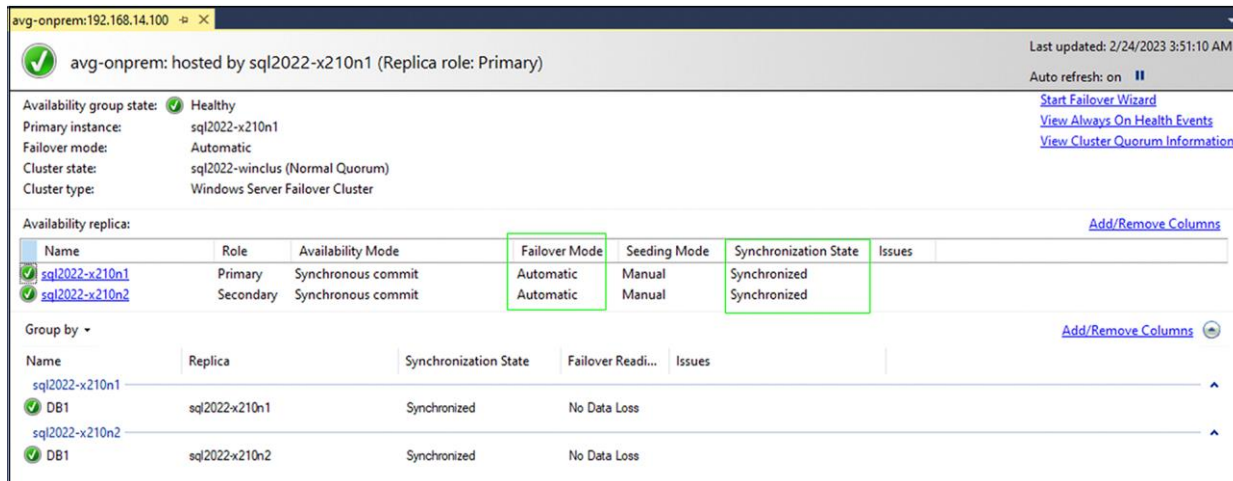
Pinging 192.168.100.31 from 192.168.100.32 with 8972 bytes of data:
Reply from 192.168.100.31: bytes=8972 time<1ms TTL=128
Reply from 192.168.100.31: bytes=8972 time<1ms TTL=128
```

Deploying highly available SQL Server databases using Always On AG

Microsoft SQL Server Always On Availability Group (AG) is the most commonly adapted technology by customers; it enables them to achieve both High Availability (HA) and Disaster Recovery (DR) of the SQL Server database. This feature takes HA and DR to a new level by enabling multiple copies of the database to be highly available. As a result, read-only workloads are enabled, in addition to the ability to offload management tasks such as backups, database consistency checks, and so on to the secondary replicas.

For this validation, use two Cisco UCS X210c blades to deploy AG with synchronous-commit replication and configure it with automatic failover. If the current primary replica is not reachable for any reason, automatic failover of Availability Group will be initiated from the current primary to the secondary synchronous replica located within the same datacenter. Since the secondary replica already has the latest copy of the data, it can quickly recover the database and start servicing the workload.

The following screenshot shows the SQL Server Management Studio view of validated deployment of AG using two Cisco UCS X210c blades. Each blade is configured with Intel Xeon 3rd-generation scalable processors, 256-GB memory, and four NVMe SSDs configured with RAID 10 for storing database data files and two SAS SSDs configured with RAID 1 for storing database transaction log files.



The screenshot displays the SQL Server Enterprise Manager interface for an Availability Group named 'avg-onprem'. The group is hosted by 'sql2022-x210n1' (Primary role) and is in a 'Healthy' state. The configuration includes:

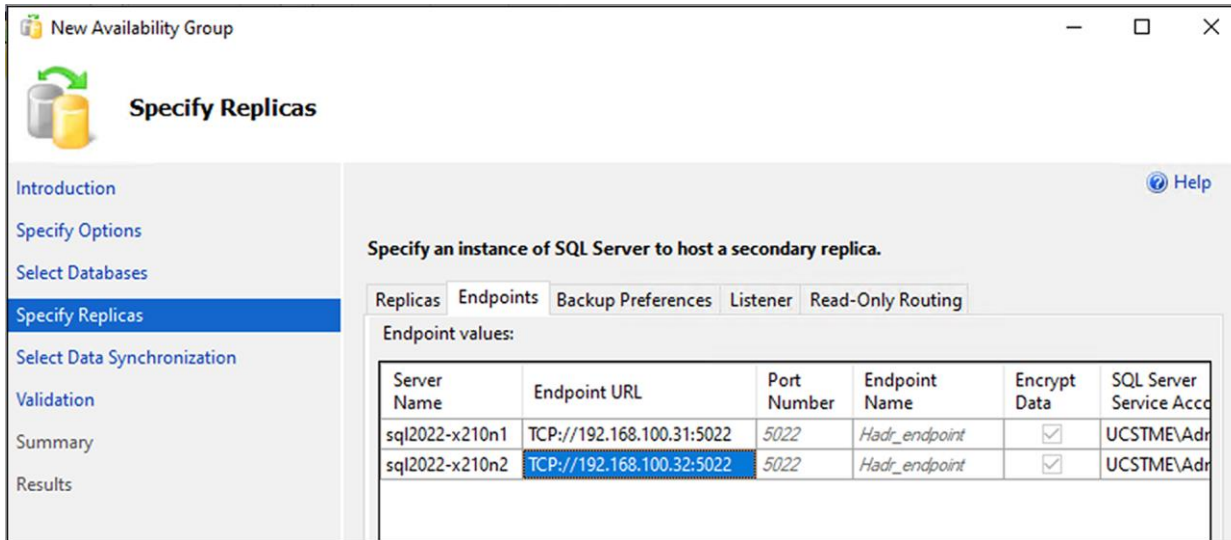
- Availability group state: Healthy
- Primary instance: sql2022-x210n1
- Failover mode: Automatic
- Cluster state: sql2022-winclus (Normal Quorum)
- Cluster type: Windows Server Failover Cluster

The 'Availability replica' table shows two replicas:

Name	Role	Availability Mode	Failover Mode	Seeding Mode	Synchronization State	Issues
sql2022-x210n1	Primary	Synchronous commit	Automatic	Manual	Synchronized	
sql2022-x210n2	Secondary	Synchronous commit	Automatic	Manual	Synchronized	

The 'Group by' view shows the database 'DB1' on both replicas, with a 'Synchronization State' of 'Synchronized' and 'No Data Loss'.

For this validation, a dedicated network (AG-Replication) is used for AG replication traffic between the two Cisco UCS X210c nodes configured with synchronous replication. This network is configured with a 9000-MTU for higher network bandwidth. The following screenshot shows the endpoint configuration using the dedicated network:



Details of deploying Availability Group on a Windows-based Failover cluster is available here:

<https://learn.microsoft.com/en-us/sql/database-engine/availability-groups/windows/use-the-availability-group-wizard-sql-server-management-studio?view=sql-server-ver16>

<https://techcommunity.microsoft.com/t5/itops-talk-blog/step-by-step-creating-a-sql-server-always-on-availability-group/ba-p/648772>

Disaster recovery with Azure SQL Managed Instance

Traditionally, disaster recovery of Availability Group databases hosted in the primary site is achieved by replicating data to a secondary replica deployed in a remote datacenter located a few miles away from the primary site. This secondary replica is configured to receive updates from the primary replica in asynchronous-commit mode. Asynchronous replication mode means that only manual failover from primary to this secondary site is supported.

Microsoft SQL Server 2022 introduced the Managed Instance (MI) Link feature, which enables near real-time data replication from an SQL Server instance deployed anywhere in an on-premises datacenter to Azure SQL Managed Instance. Replicating data to Azure cloud enables disaster recovery and also unlocks several other scenarios such as scaling read-only workloads, offloading analytics to the cloud, and so on. You can configure replication to SQL MI from a standalone SQL Server instance or Availability Group listener.

Note that disaster recovery with an Azure SQL MI Link is possible only when you are using SQL Server 2022 in the on-premises datacenter; in addition, this feature is in public preview at the time of writing this document, and it is limited to the users or customers who are signed up for the feature using this link:

<https://forms.office.com/pages/responsepage.aspx?id=v4j5cvGGr0GRqy180BHbR5o5Tvh3duNFm5f7UUZcVpZUNFZDVkrJTKhfFSEw0MksZUVZOOEUzRERFM4u>.

For this validation, site-to-site secure Virtual Private Network (VPN) is configured between the on-premises datacenter and Azure cloud. After the VPN is established between the two sites, ensure your private on-premises network can reach Azure Virtual Network and subnets defined in it. For detailed steps to configure site-to-site VPN between on-premises datacenter and Azure cloud, please visit:

<https://learn.microsoft.com/en-us/azure/vpn-gateway/tutorial-site-to-site-portal>.

For deploying SQL Managed Instance in Azure cloud, follow the steps detailed in this link:
<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/instance-create-quickstart?view=azuresql>.

SQL Managed Instance is a set of service components that are hosted on a dedicated set of isolated virtual machines that are joined to a virtual cluster. A virtual cluster can host multiple managed instances, and the cluster automatically expands or contracts as needed to accommodate new and removed instances. SQL MI deployment requires a dedicated subnet in the Azure Virtual Network, and it cannot be used by any other services. After you deploy SQL MI with the required configuration, in order to access it from the on-premises datacenter through the secure VPN, you must update the Network Security Group (NSG) of Azure SQL MI to allow network traffic on the following ports:

- Inbound traffic on port 5022 and port range 11000 to 11999 from the on-premises network

The following screenshot shows the SQL MI NSG configured to allow traffic from on-premises SQL Server on the required ports:

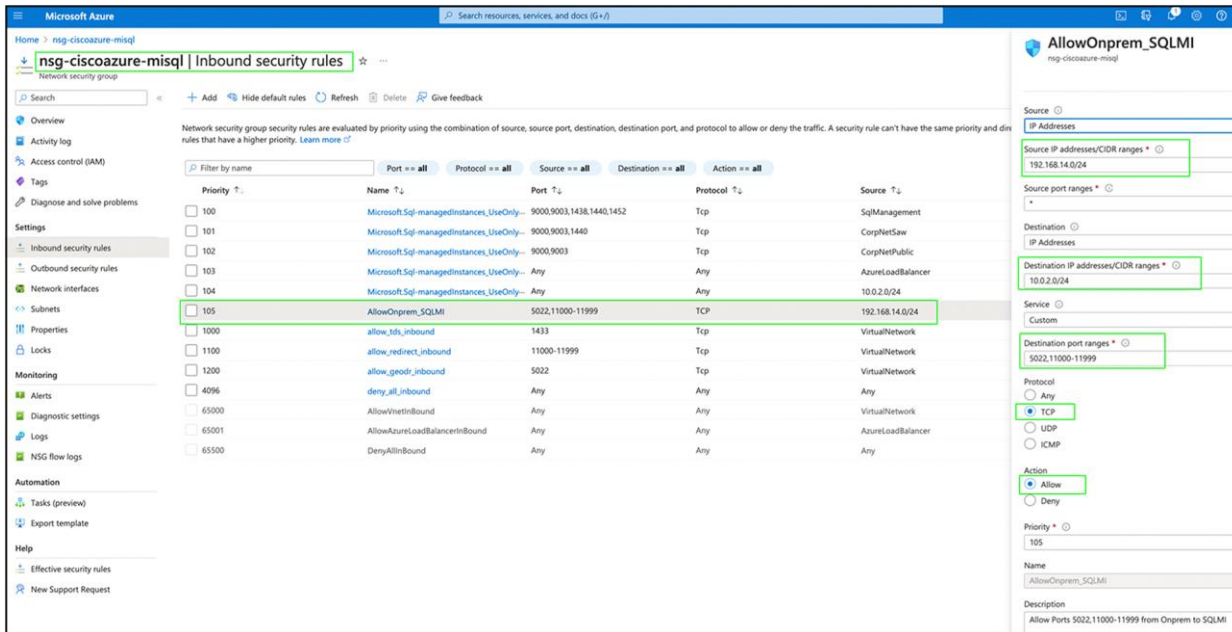


Table 5 provides more details about the NSG inbound rule configuration shown in the previous screenshot.

Table 5. SQL MI Network Security Group Inbound rule to allow communication from on-premises SQL Server Instance

Parameter	Values
On-premises DC network/subnet	192.168.14.0/24
SQL MI name	ciscoazure-misql
SQL MI dedicated subnet	sql-subnet (10.0.2.0/24)
Source IP range	On-premises DC network (192.168.14.0/24)
Source port range	5022 and 11000 to 11999
Destination IP range	SQL MI dedicated subnet (10.0.2.0/24)
Destination port range	Any (*)
Protocol	TCP
Action	Allow

Similarly, in the on-premises datacenter, the following firewall ports must be allowed on all Windows Servers hosting on-premises SQL Server Instances:

- Inbound traffic on port 5022 must be allowed from the entire subnet range hosting SQL Managed Instance
- Ensure that you have added your organization’s Domain Name System (DNS) IP address or 8.8.8.8 so that the Managed Instance name is resolvable from the Windows Server hosts running SQL Server Instances.

Figure 10 shows a high-level network configuration between on-premises and Azure SQL MI network subnets.

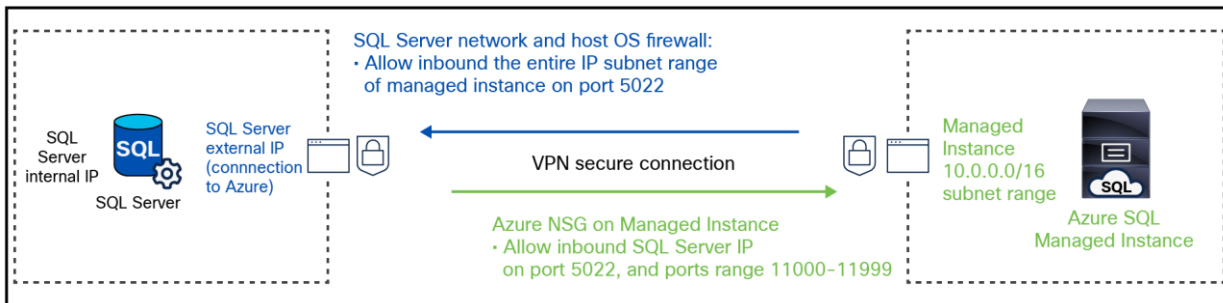
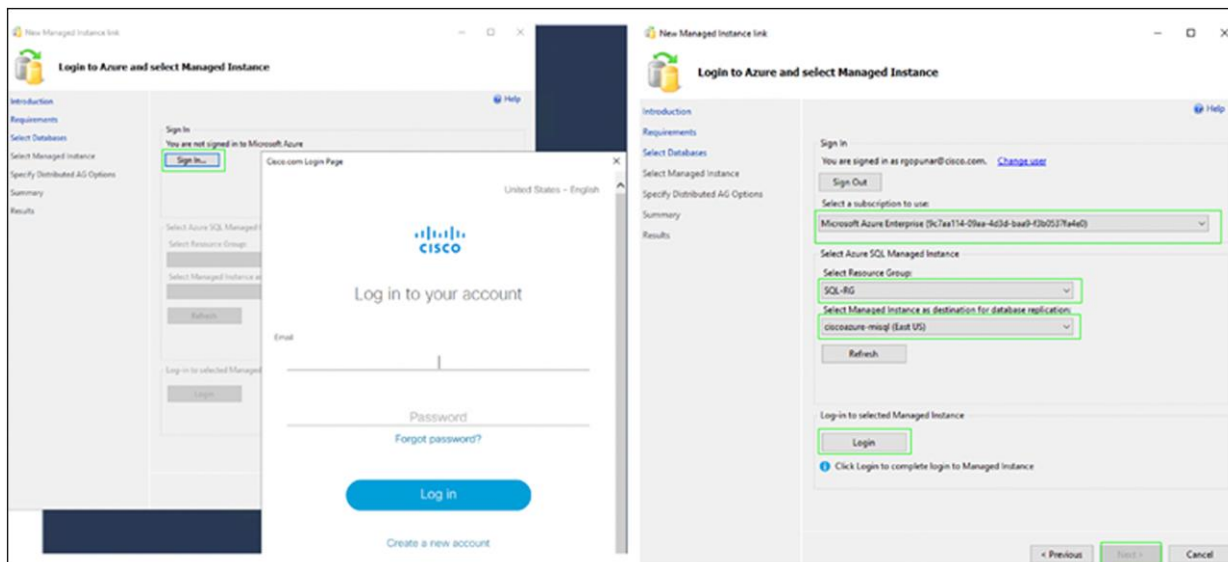


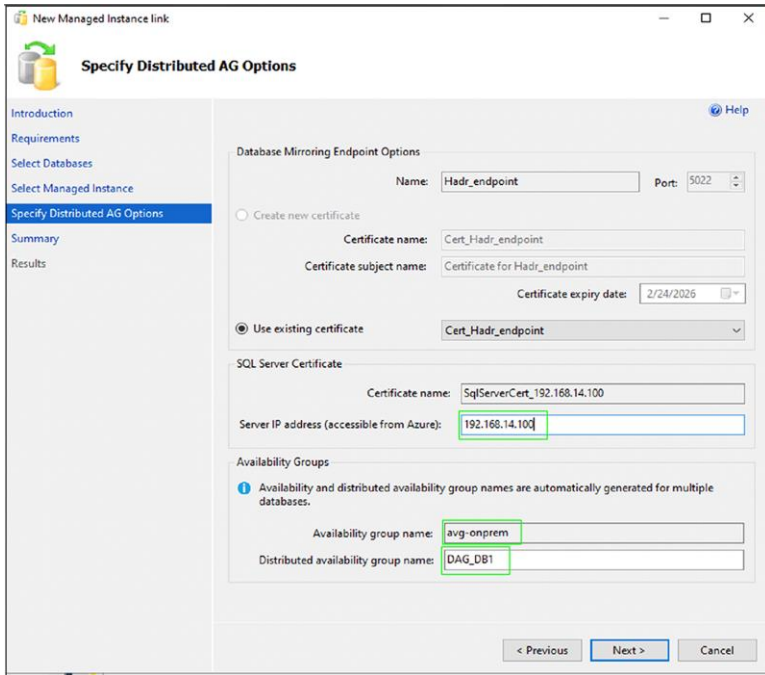
Figure 10. Network connectivity between on-premises DC and Azure SQL MI network subnets

After you have created and configured the Azure SQL MI to be accessible from the on-premises datacenter, you can configure replication to the Azure SQL MI from the on-premises datacenter SQL Server Instance by doing the following:

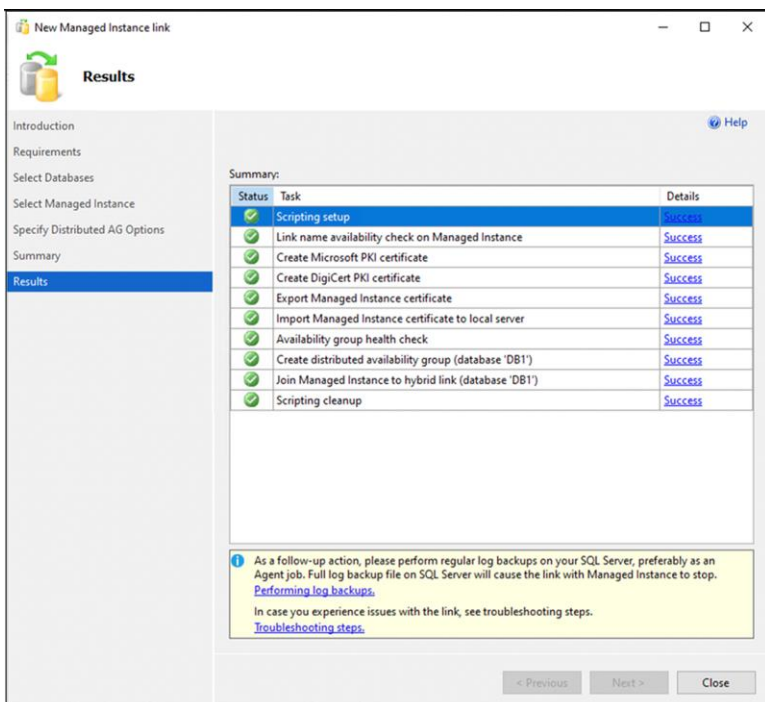
1. Download and install SQL Server Management Studio Version 19.0 or later.
2. Select the database from the on-premises SQL Server Instance. Right click on the database and select the Azure SQL Managed Instance link -> Replicate Database. Click Next on the introduction screen.
3. Ensure that all the requirements are met as called out on the Requirements page. Add trace flags 1800 and 9567 to the SQL Server.
4. Enable the Always On AG feature on all the on-premises instances and restart them. Encrypt the AG endpoints that AG uses with database master keys. Click Next.
5. Select the database you want to replicate to the Azure SQL MI instance. The database must be backed up in order to be eligible for replication. Click Next.
6. Click the **Sign In** button to direct the authentication page because you need to be authenticated by both your organization and Microsoft using Multi Factor Authentication (MFA). When you are successfully authenticated, select your Azure subscription, Resource Group and Azure SQL MI. Click the Login button to connect to your Azure SQL MI Instance. When you are connected, click Next (refer to the screenshot that follows).



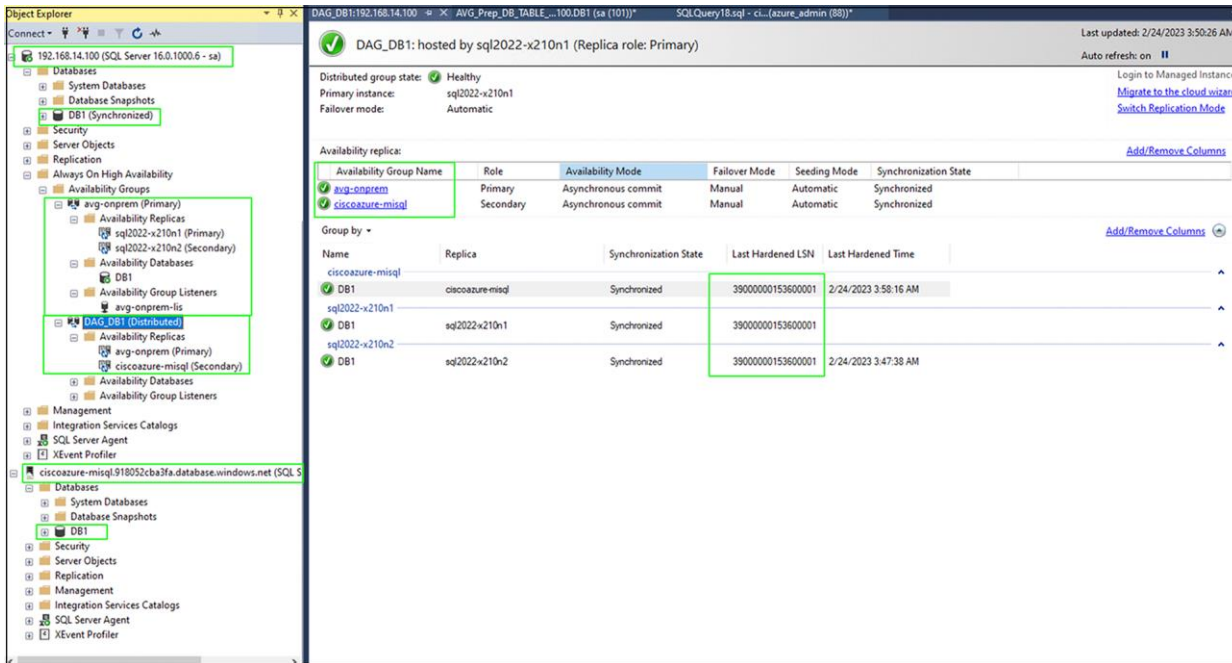
- On the Distributed AG Options page, most of the values are auto-populated. Provide your on-premises AG listener IP address if you wish to replicate the database from the existing on-premises AG with an AG listener. If not, just provide the IP address of a standalone SQL Server Instance. Provide a name for the Distributed Availability Group. In the following screenshot, 192.168.14.100 is the on-premises AG listener IP address configured for on-premises AG configured with two synchronous replicas with an automatic failover option.



- On the Summary page, review the summary of actions and click finish. The following should appear if you have successfully executed all the steps:



9. You can verify the on-premises AG and link to Azure SQL MI from the AG dashboard, as shown in the following screenshot:



10. Start running your transactional workload on the on-premises databases and verify that your data changes are replicated to on-premises synchronous replicas as well as to an Azure SQL MI Instance. You can look at the “Last Hardened LSN” number to verify that all the replicas are synchronized and have the latest copy of the database.

For more detailed steps for replicating on-premises databases to an Azure SQL Managed Instance using the link feature in the SQL Server Management Studio (SSMS) refer to: <https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/managed-instance-link-use-ssms-to-replicate-database?view=azuresql>.

Failover and failback validation

Within the on-premises datacenter, redundant hardware and software components are used to achieve higher availability to the deployed SQL Server databases. Use two separate Cisco UCS X210c blades (installing them in two different Cisco UCS X9508 Chassis is recommended) with RAID-configured local attached drives for hosting two independent SQL Server Instances. For replicating database data changes synchronously between two SQL Server instances, configure the Availability Group (AG) feature with the automatic failover option. If one instance is not available for any reason, failover is initiated automatically to the secondary SQL Server Instance. The secondary instance then becomes primary and starts servicing the workload. When the original primary is back online, it becomes the secondary and tries to catch up and synchronize with the current primary database.

As discussed in the previous sections, use the Azure SQL Managed Instance link (using Distributed AG) for replicating data changes asynchronously from the on-premises datacenter to the Azure cloud for disaster-recovery purposes. If a complete disaster occurs in the on-premises datacenter, a manual failover must be triggered from the SQL Managed Instance to the Azure cloud. After failover to Azure cloud occurs, the Azure SQL MI starts servicing the read-write workload until the on-premises datacenter setup is brought online. Optionally, you can delete the existing link between the on-premises and Azure cloud. When the on-premises setup becomes available, you can execute a full backup of the database from the Azure SQL MI and restore it on the on-premises datacenter and then reconfigure the Azure link again for replicating the data changes to the cloud. The detailed steps for failover and failback from the on-premises datacenter to the Azure cloud and vice versa are captured here: <https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/managed-instance-link-use-ssms-to-failover-database?view=azuresql>

<https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/restore-database-to-sql-server?view=azuresql&tabs=managed-identity>.

Accelerating Microsoft SQL Server backup compression performance with Intel integrated Accelerator Engines

Intel Xeon Scalable processors feature the broadest and widest set of built-in accelerator engines for today's most demanding workloads. Whether on-prem, in the cloud, or at the edge, Intel Accelerator Engines can help take your business to new heights, increasing application performance, reducing costs, and improving power efficiency.

With 4th-generation Intel Xeon Scalable processors (called Sapphire Rapids), these accelerators are built into the processors without having to purchase additional specialized hardware for offloading most common tasks. Intel offers various accelerator engines delivering enhanced performance across various emerging workloads. For more details on the various built-in accelerator engines supported by Intel, refer to this link: <https://www.intel.in/content/www/in/en/now/xeon-accelerated/accelerators-eguide.html>

This section provides validation and the test results conducted in Cisco labs on the Intel (QAT) accelerator using Microsoft SQL Server databases for offloading backup compression.

Note that these integrated accelerators are shipped only with Intel 4th-generation Xeon Scalable processors. Thus for QAT testing, the UCS X210c M7 compute node is used as it is featured by Intel 4th-generation Scalable processors.

Table 6. Test configuration used for Intel QAT validation

Component	Details
Compute	One Cisco UCS X210c M7 compute node with 2x Intel 4 th -generation Xeon 8468H Scalable Processors with total of 96 physical cores
Memory	512GB (32G x 16) - at 4800 MTs
Storage	2x 1.6TB NVMe SSDs for storing database files 2x 1.6TB NVMe SSDs for storing backup files 2x M.2 drives with HW RAID1 for OS and SQL binaries

Component	Details
Operating system	Windows Server 2022 Datacenter Standard Edition
Database	Microsoft SQL Server 2022 Evaluation Edition
DB sizes tested	500GB/1000GB
Testing tool	HammerDB
Tests conducted	Full Database backup with SQL Server native compression Full Database backup with Intel QAT compression Transactional log backups with SQL native and Intel QAT with workload running on the database with 50 HammerDB users

After installing Windows Server 2022 on the Cisco UCS X210c M7 compute node, Intel QAT drivers are installed in order to detect the Intel QAT accelerator cores. The Microsoft SQL Server 2022 instance is installed, and the 500GB operational database is created and loaded with the HammerDB tool.

Refer to this link for Intel QAT accelerator drivers:

<https://www.intel.in/content/www/in/en/download/765502/intel-quickassist-technology-driver-for-windows-hw-version-2-0.html?>

Microsoft and Intel worked together to enable the support for offloading backup compression using the Intel QAT hardware accelerator. The Microsoft SQL Server 2022 supports offloading and backup acceleration using the Intel QAT accelerator. The SQL Server 2022 is configured for offloading the backup compression to Intel QAT as described here: <https://learn.microsoft.com/en-us/sql/relational-databases/integrated-acceleration/use-integrated-acceleration-and-offloading?source=recommendations&view=sql-server-ver16>

Figure 11 illustrates backup compression performance improvements with Intel QAT as compared to SQL Server native compression for a 500G full database backup use case.




Figure 11.
SQL Server backup compression performance improvements with Intel QAT

As shown in the above screenshot, backup duration with Intel QAT compression is nearly 2.5 times reduced compared to the SQL server native backup compression algorithm. The system CPU utilization during compression with Intel QAT is around 3.45%, whereas CPU utilization with SQL Server native compression is 8.5%. This is because the compression is being offloaded from actual CPU cores to Intel integrated QAT accelerator cores. The freed CPU cycles can be used for executing actual customer workload. Also, Intel QAT can process the compression at nearly 3900 MB/s speed, while SQL native compression speed by comparison is only around 1600 MB/s.

Note that with increased database size of 1000GB (1TB), similar compression performance improvements are observed with Intel QAT.

Here are the advantages of using Cisco UCS X210c M7 compute nodes powered by Intel 4th-generation Xeon Scalable processors with the Intel Quick Assist Technology accelerator for backup compression.

1. Achieve better Return on Investment (ROI) by offloading CPU-intensive compression/decompression operations to Intel QAT accelerators.
2. Achieve nearly 2.5 times reduction in the duration of Microsoft SQL Server backup by offloading compression to Intel QAT.
3. Faster backup/restore times enable customers to reduce the maintenance window times for critical production environments, thereby increasing availability of the databases.
4. Freed CPU cycles of processors can be used for executing actual customer workloads.
5. To analyze the performance improvements with Intel QAT, similar tests were conducted with transactional log backup with the workload running on the database. In addition, nearly similar performance improvements were noticed even with the transactional log backup.

Summary

The Cisco UCS X-Series with Cisco Intersight software is a modular system managed from the cloud. It is designed to be shaped to meet the needs of modern applications and improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design. The Cisco UCS X-Series Modular System simplifies your datacenter, adapting to the unpredictable needs of modern applications while also providing for traditional scale-out and enterprise workloads.

This document has discussed the deployment best practices for configuring the Cisco UCS X210c M6/M7 blade for hosting SQL Server databases on local NVMe and SAS SSD drives. It addressed AG configuration best practices for achieving high availability of databases within an on-premises datacenter and provided steps to replicate an SQL 2022 databases to Azure SQL MI for disaster-recovery purposes.

This document has also discussed the benefits of using the Intel QAT accelerator for offloading compression/decompression operations, which reduces backup duration with reduced CPU usage and thus achieves better ROI with Cisco UCS X210c M7 compute node-based MSSQL deployments.

For more information

For additional information about UCS X- Series modular system, refer to:

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-x-series-modular-system/index.html>
or

<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/solution-overview-c22-2432175.html?ccid=cc002456&oid=sowcsm025665>.

The Cisco Intersight configuration guide is available here: <https://www.cisco.com/c/en/us/products/cloud-systems-management/intersight/index.html?ccid=cc001268>.

For information about BIOS tunings for different workloads, refer to the BIOS tuning guide at:

<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/performance-tuning-guide-ucs-m6-servers.pdf>

For more information about the Managed Instance link feature, refer to: <https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/managed-instance-link-feature-overview?view=azuresql>

For more information about Azure SQL Managed Instance, refer to: <https://learn.microsoft.com/en-us/azure/azure-sql/managed-instance/sql-managed-instance-paas-overview?view=azuresql>.

For more details about SQL Server Availability Groups, refer to: <https://learn.microsoft.com/en-us/sql/database-engine/availability-groups/windows/overview-of-always-on-availability-groups-sql-server?view=sql-server-ver16>.

To create a site-to-site secured VPN, refer to: <https://learn.microsoft.com/en-us/azure/vpn-gateway/tutorial-site-to-site-portal>.

For Intel Integrated Accelerators, refer to : <https://www.intel.in/content/www/in/en/now/xeon-accelerated/accelerators-eguide.html>

For configuring backup compression offloading with Intel QAT on Microsoft SQL Server 2022, refer to: <https://learn.microsoft.com/en-us/sql/relational-databases/integrated-acceleration/use-integrated-acceleration-and-offloading?source=recommendations&view=sql-server-ver16>

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)