

Migrating SAP Applications to Azure



Table of contents

Migrating SAP Applications to Azure

- 01** [Introduction](#)
- 02** [Journey to the Cloud: Roadmap for SAP Customers](#)
- 03** [Azure Offerings for SAP](#)
 - SAP HANA on Azure Virtual Machines (VMs)
 - SAP HANA on Azure Large Instances
- 04** [SAP Certifications for Azure](#)
- 05** [Migration Tools and Methodologies](#)
 - Backup and Restore
 - SAP HANA Classical Migration
 - SAP HANA HSR
 - DMO
- 06** [Migration Scenarios](#)
 - Migrating from Any DB to Any DB on Azure
 - Migrating from Any DB to HANA on Azure
 - Migrating from on-premises HANA to Azure
- 07** [Strategies for Minimizing Downtime](#)
- 08** [Prerequisites for Migration](#)
 - Specialized Knowledge
 - SAP on Azure Sizing
 - Instance Sizing for SAP
 - Network Design
 - Scale-up and Scale-out Sizing
- 09** [Post-Migration Steps](#)
- 10** [High Level Technical Health Checks](#)
- 11** [Conclusion](#)

01 / Introduction

One thing is certain – all SAP NetWeaver landscapes will migrate to SAP HANA and the cloud. And all SAP app and industry modules will also become available in S/4HANA. It's also clear that this migration will take a long time.

For many SAP managers, it's the journey to SAP HANA and the cloud that concerns them, because many things need to work side by side until that journey is complete. SAP managers need help plotting this journey to the cloud where all versions of their applications and systems – whether SAP NetWeaver or S/4HANA, Windows or Linux, AnyDB or SAP HANA, non-production or production, temporary or permanent, smallest or largest, can exist side by side until that journey is complete.

With the broadest global footprint, largest compliance portfolio, enterprise-grade SLAs, and world-class support, Azure provides a robust, resilient, and reliable environment for your SAP applications.

This guide is for IT professionals who manage enterprise SAP environments and want to learn more about migrating SAP landscapes to Azure. This guide provides the tools and processes required for migrating mission-critical SAP landscapes to the cloud.

Why migrating to Azure makes sense

Most SAP customers are looking for an agile deployment architecture. SAP managers are feeling pressured to start moving business functionality to S/4HANA. But they don't know how their SAP landscape will look on S/4HANA. This is because SAP roadmaps for S/4HANA doesn't yet address many of the customer's businesses processes that are currently in use. As a result, SAP customers face the following situation:

- They don't know how many S/4HANA systems they need to deploy.
- Many short-lived S/4HANA systems must be created for project purposes.
- It's unclear when they can decommission their NetWeaver-based systems.
- With so many unknowns, it's very hard to get an accurate picture of the on-premise hardware investment.

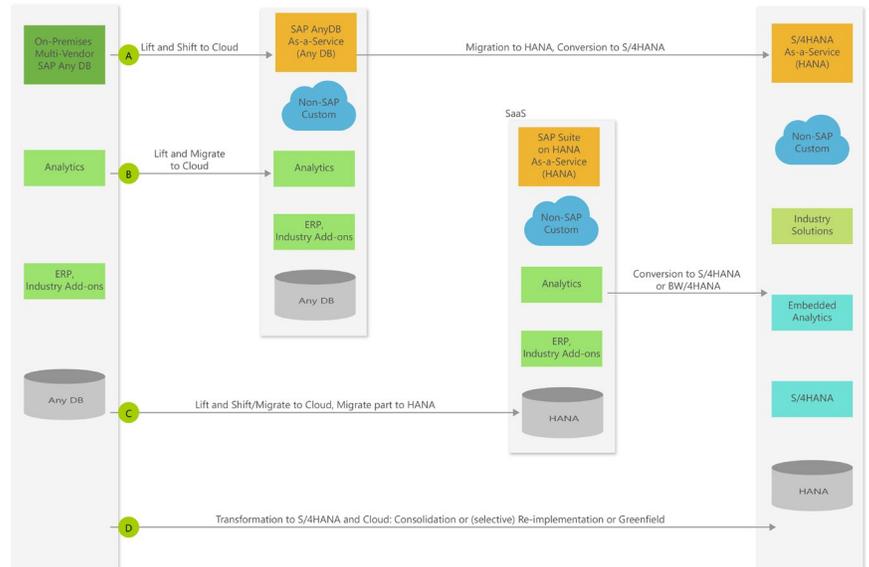
Azure provides the answer to this infrastructure investment question. Azure allows SAP customers to deploy S/4HANA systems on demand and pay only for the active infrastructure usage. Azure also allows SAP customers to move their existing SAP systems running on all SAP supported databases. Furthermore, customers can right size the Azure infrastructure of older NetWeaver systems as business processes move out of those systems into S/4HANA.

Customers can accommodate infrastructure demand on S/4HANA systems by changing to larger and larger VM types for the SAP HANA database and scale-out the application layer. Amongst all the public cloud providers, Azure uniquely offers a purpose-built SAP HANA Large Instance option supporting up to 20TB on a single node.

Finally, Azure is the only hyperscale provider that offers 99.99% SLA for SAP HANA infrastructure. Azure is the enterprise-grade scalable, reliable, and agile infrastructure solution that allows you to optimize your infrastructure as you transition from NetWeaver-based application to the SAP S/4HANA platform.

Journey to the Cloud

Roadmap for SAP Customers



Currently, many customers are running their SAP Business Suite applications on-premise with Any DB. Any DB refers to any non-HANA SAP-certified database, such as SQL Server, Oracle, DB2, Sybase ASE, or MaxDB.

The diagram above displays a roadmap for different types of SAP customers moving to Azure who have a goal of being on the S/4HANA platform at the end of 2025 or before:

A: This is a lift and shift move to Azure. The SAP customer runs their SAP systems on an OS/DBMS combination that they want to keep in Azure as well. This type usually includes customers already running their SAP systems on Windows and Linux on-premises. The main driver to move their existing SAP systems into Azure is the customer wanting to be able to apply right sizing to those systems over time and run them on the same platform as their S/4HANA platform. Typical examples of these kinds of customers include [Microsoft](#) or Zespri.

B: Some customers need to abandon their OS and even their DBMS platform in the move to Azure. Drivers for such a move include not being able to run their proprietary Unix, AS400, or mainframe OS in Azure. Or the customer may simply want to move to a more cost-efficient DBMS system. A classic case for such a move is [Malaysia Airlines](#). These customers usually do not see any benefit to move their NetWeaver-based

business suite to SAP HANA, but look at redesigning their SAP landscape for S/4HANA.

C: This type includes customers who already run SAP HANA on-premise and the on-premise appliance hardware is running out of service. As Azure accommodates even the largest SAP HANA deployments, SAP customers can choose to move to Azure with their SAP HANA-based NetWeaver application performing a heterogeneous platform migration.

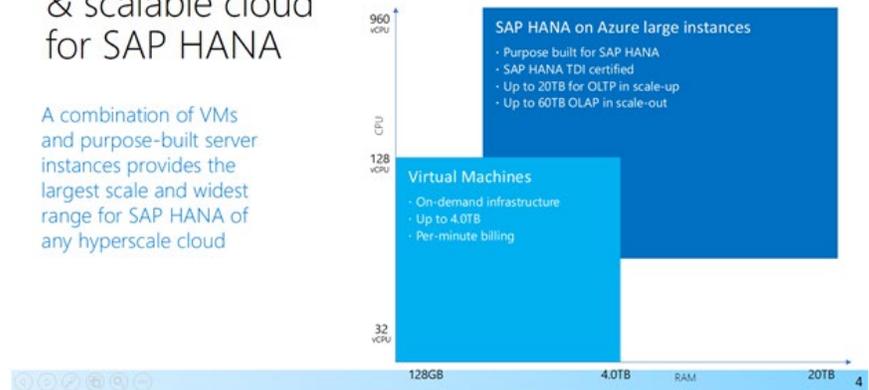
D: This type includes customers that deploy their new SAP S/4HANA systems on Azure. This includes those customers that have already deployed their SAP NetWeaver-based systems in Azure and have now set up their S/4HANA systems to transform their NetWeaver-based business processes to S/4HANA.

Even though every customer's SAP landscape is different, Microsoft Azure covers all types: Any DB, Suite on HANA, and S/4HANA.

HANA on Azure Virtual Machines (VMs)

Most powerful & scalable cloud for SAP HANA

A combination of VMs and purpose-built server instances provides the largest scale and widest range for SAP HANA of any hyperscale cloud



When it comes to SAP, we understand that you want to pick infrastructure that grows with your business needs. For example, you might have a 3 TB database today, but your transaction volume might quickly outgrow it.

We provide two different infrastructure technologies, both which are fully aligned with SAP sizing guidance and certification criteria.

SAP HANA on Azure Large Instances is SAP TDI certified HANA hardware which is purpose-built bare-metal hosted in Azure data centers and specifically targets large and extra-large SAP HANA workloads. HANA on Azure Large Instances offer the largest OLTP nodes up to 20TB which in High Availability (HA) configurations can provide 99.99% SLAs, and the largest OLAP nodes with 4TB usable. Azure also supports N+M scale out clusters and provides single node SLAs of 99.9% to deliver production class SAP HANA footprints.

Azure Virtual Machines are on-demand VMs for all sizes of SAP, Any DB, and small and medium SAP HANA, offering up to 4TB of RAM. HANA certified VMs give you all the benefits of on-demand infrastructure, including agility, the ability to shut down when not in use, and lower cost systems for dev/test and training use.

Azure M-Series VMs are optimized for large in-memory workloads such as SAP HANA and SQL Hekaton, utilizing high memory (RAM) and virtual CPUs (vCPU) powered by Intel® Xeon® 2.5 GHz E7-8890 v3 processors. M-Series VMs support up to 128 vCPUs and between 1 TiB and 3.8 TiB (4 TB) RAM. With Azure Premium Storage and Azure Write Accelerator for low latency workloads M-Series VMs are designed for the largest enterprise database workloads.

The Ev3-series features the Intel® Xeon® E5-2673 v4 2.3 GHz processor in a hyper-threaded configuration, providing a better value proposition for most general-purpose workloads, and bringing the Ev3 into alignment with the general purpose VMs of most other clouds. Memory has been expanded (from 7 GiB/vCPU to 8 GiB/vCPU) while disk and network limits have been adjusted on a per core basis to align with the move to hyperthreading. The Ev3 is ideal for SAP application layers and medium sized SAP databases.

The Dv3-series provides a better value proposition for most general-purpose workloads, and bring the Dv3 into alignment with the general purpose VMs of most other clouds. Memory has been expanded (from ~3.5 GiB/vCPU to 4 GiB/vCPU) while disk and network limits have been adjusted on a per core basis to align with the move to hyperthreading. The Dv3 is ideal for SAP application layers and small to medium SAP databases.

The documents listed below provide additional guidance for HANA on Azure VMs:

- [Getting started with SAP HANA on Azure](#)
- [High Availability of SAP HANA on Azure Virtual Machines](#)
- [Backup guide for SAP HANA on Azure Virtual Machines](#)

Microsoft works with SAP to run benchmark tests and certify Azure Infrastructure (VMs and SAP HANA Large Instances) for production SAP workloads, both AnyDB and SAP HANA.

SAP HANA on Azure VM certification

Category	Scenario	SAP Product	Guest OS	Database	VM/Server Type
SAP HANA	OLAP	SAP Business Warehouse (BW); BW/4HANA, HANA Enterprise, HANA Side Car	SUSE, Red Hat	SAP HANA	GS5, HANA on Azure (Large Instances), M-Series
	OLTP	S/4HANA	SUSE, Red Hat	SAP HANA	GS5 (** Controlled Availability), HANA on Azure (Large Instances), M-Series
	OLTP	SAP Business Suite on HANA, SAP NetWeaver	SUSE, Red Hat	SAP HANA	HANA on Azure (Large Instances), M-Series
	OLTP	SAP Business One on HANA	SUSE, Red Hat	SAP HANA	DS14_v2, M-Series
	-	SAP HANA One	SUSE, Red Hat	SAP HANA	DS14_v2

SAP (AnyDB) on Azure VM certification

Category	Scenario	SAP Product	Guest OS	Database	VM/Server Type
SAP Any DB	-	SAP Business Suite 6.x (SAP NetWeaver Application Server 7.x)	Windows, SU SE, Red Hat, Oracle Linux *1	SQL Server, ASE, Oracle, MaxDB, DB2/UDB	D/DS11_v2 – D/DS15_v2 GS1 – GS5 D_v3, M-Series
		SAP Business All-in-One			
		SAP NetWeaver Application Server 7.x (ABAP and Java)			
		SAP BPC NetWeaver			
		SAP BPC Microsoft 10.01	Windows	SQL Server	
		TREX 7.10	Windows, SUSE, Red Hat *	N/A	
		SAP LiveCache, SAP Content Server 6.50	Windows, SUSE, Red Hat *	MaxDB	
		SAP BusinessObjects BI/Dashboards/WebIntelligence/Explorer 4.1, Crystal Reports 2013 4.1, Design Studio BI Platform, Data Services 4.2	Windows Server, SUSE, Red Hat	N/A	
		SAP Business One	Windows Server	SQL Server	

SAP HANA on Azure Large Instances certification

SKU	SAP Products	Workload	CPUs	CPU Threads	Memory	Storage	Available Regions
S72	SAP BW, SAP BW/4HANA or SAP HANA, SAP Business Suite on SAP HANA, S/4HANA	OLAP/OLTP	36 2 x Intel® Xeon® Processor E7-8890 v3	72	768GB	3TB	US West US East Australia East Australia Southeast Japan East Japan West West Europe North Europe
S192*			96 4 x Intel® Xeon® Processor E7-8890 v4	192	2.0TB	8TB	
S384			192 8 x Intel® Processor E7-8890 v4	384	4.0TB	16TB	
S72m	SAP Business Suite on SAP HANA or SAP S/4HANA	OLTP	36 2 x Intel® Xeon® Processor E7-8890 v3	72	1.5TB	6TB	
S192m			96 4 x Intel® Xeon® Processor E7-8890 v4	192	4.0TB	16TB	
S384m*			192 8 x Intel® Xeon® Processor E7-8890 v4	384	6.0TB	18TB	
S384xm			192 8 x Intel® Xeon® Processor E7-8890 v4	384	8.0TB	22TB	
S576m*			288 12 x Intel® Xeon® Processor E7-8890 v4	576	12.0TB	28TB	
S768m			384 16 x Intel® Xeon® Processor E7-8890 v4	768	16.0TB	36TB	

*Certifications coming soon

SKU	SAP Products	Workload	CPUs	CPU Threads	Memory	Storage	Available Regions
S960m			480 20 x Intel® Xeon® Processor E7-8890 v4	960	20.0TB	46TB	

*Certifications coming soon

The tables above contain SAP HANA and Any DB products that are certified to run in Microsoft Azure. All SAP products with Any DB are fully certified, including SAP Business Suite, NetWeaver, Business Objects, and Business One.

Supported platforms include Windows, Linux, SQL Server, ASE, Oracle, MaxDB, and DB2/UDB. You can use D v2, D v3, E v3, G and M Series VMs.

All SAP products with HANA are also fully certified. OLAP solutions such as Business Warehouse, HANA Enterprise, and sidecar are certified with GS5 VM, M Series VMs, and HANA Large Instances. OLTP solutions such as S/4HANA and Business Suite on HANA are certified with Large Instances, and E v3 Series VMs will be certified this calendar year.

[Find out more](#) about SAP HANA architecture on Azure Large Instances.

This section outlines the various tools and methodologies that you can use for your SAP system migration to Azure.

Backup and restore

Backup and restore options are commonly used for saving data on a source system and restoring it on a different system. Azure has many storage options available to help with data transfer.

Backup and restore options are tried-and-true mechanisms for saving data on a source system and restoring it to another destination. Microsoft Azure has various storage options available to help facilitate data transfer to Azure. For more information, see the following:

- [Overview of the features in Azure Backup](#)
- [Use Azure Storage for SQL Server Backup and Restore](#)
- [Backup guide for SAP HANA on Azure Virtual Machines](#)

Database replication features

Database replication features include SQL Server AlwaysOn, HANA System Replication (HSR), and Oracle Dataguard.

SQL Server AlwaysOn

One of the advantages of the SQL Server AlwaysOn methodology is that it replicates the SAP HANA data directly from the source to another database. You can extend on-premises Availability Groups to Microsoft Azure by provisioning one or more Azure VMs with SQL Server and then adding them as replicas to your on-premises [Availability Groups](#).

SAP HANA System Replication (HSR)

One of the advantages of the HSR methodology is that it replicates SAP HANA data directly from the source to another database. You can use this separate database as the primary one in case of a takeover, as it contains the exact same information from the SAP HANA database.

For more information, see [High Availability of SAP HANA on Azure Virtual Machines](#) (VMs).

Oracle Dataguard

Oracle Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive disasters and data corruptions.

SAP System Copy

SAP's system copy and migration services of Software Provisioning Manager 1.0 comes with the latest SAPinst version which includes the latest features and fixes. System copy lets you create consistent copies of your SAP systems which are adapted to your needs, whether you have an SAP system distributed to several hosts, a high-availability system, or a system with an [operating system/database combination](#).

SAP HANA Classical Migration

SAP's HANA classical migration option lets you migrate your database to SAP HANA using SAP heterogeneous copy process and tools.

For more information about classical migration, see SAP's [Classical Migration to SAP HANA](#).

SAP Software Update Manager (SUM) Database Migration Option (DMO)

SUM DMO option is a database migration process step that is integrated into SAP Update Manager (SUM). The integration into SUM allows performing SAP release upgrades in combination with SAP database platform migrations and a move into Azure. One of the advantages of SAP's SUM DMO approach is that it provides more export/import optimizations out of the box than the classical migration toolset.

For more information about the SUM DMO approach, see [SAP's Database Migration Option \(DMO\) of SUM - Introduction](#).

For the migration scenarios below, AnyDB refers to all SAP-certified relational databases except for SAP HANA. This includes SQL Server, Oracle, DB2/UDB, SAP ASE, and MaxDB.

Migrating from AnyDB to AnyDB on Azure VM

This scenario involves migrating an SAP instance from on-premises to an Azure VM while you continue to use AnyDB. This scenario has two different use cases. One case is that the SAP customer stays with the same DBMS. A classical example is a customer moving the SQL Server-based applications to Azure where the systems will still run on SQL Server. In exceptional cases, customers perform such system copies by exporting the data out of the SAP databases to adjust and align the number of database files. The usual way is to perform a backup/restore process that runs while the source SAP systems are still up and running. The second case is that the customer wants to perform a heterogeneous database migration (e.g., from Oracle to SQL Server) while moving to Azure. In such cases, an export of the databases is necessary. These export files need to be moved to Azure and imported into the target DBMS. [Find out more here.](#)

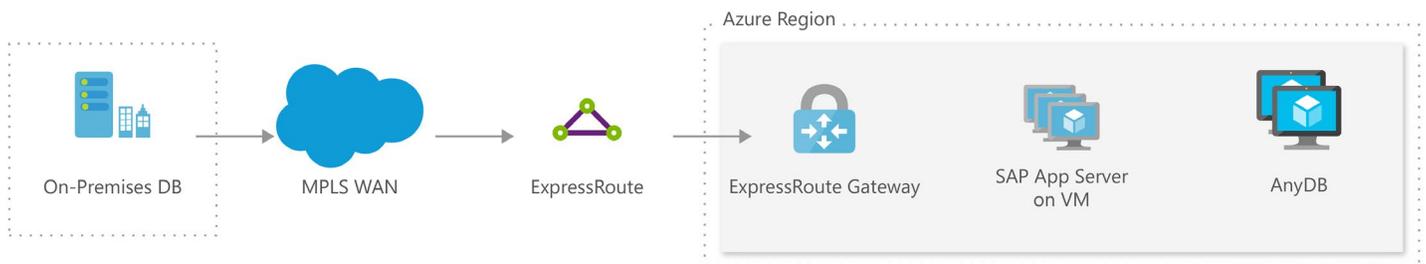
The steps for each use case is as follows:

Backup/Restore method

1. While the SAP system is up and running, perform a backup of source system.
2. Copy backup into Azure.
3. Restore in Azure VM that runs your DBMS instance.
4. Set up log-shipping so that the changes applied in the meantime can be moved into the Azure hosted instances and applied there.
5. Stop the on-premise SAP instance. Take a last transaction log backup. Make sure the last transaction log backup gets copied to the Azure hosted instance and applied.
6. Start SAP system that you installed in Azure against the Azure hosted DBMS instance.
7. Perform post-system copy tasks.

Export/Import method

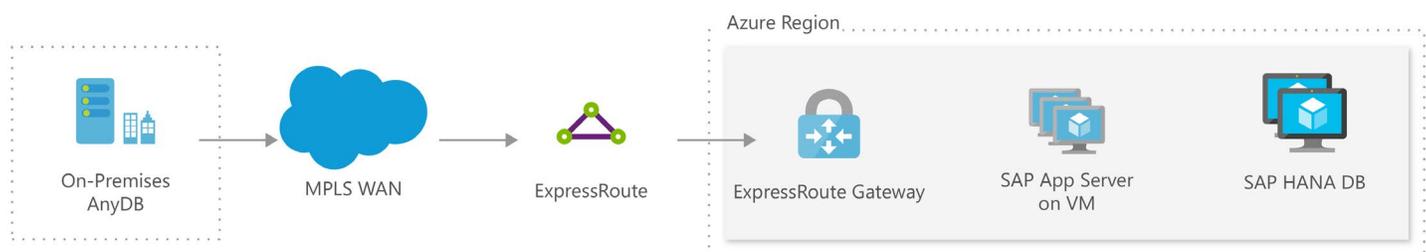
1. Set up migration tools on-premise for exporting the SAP data and set up the components needed for importing on Azure.
2. Shut down the SAP system hosted on-premise.
3. Start export processes on-premise and start copying finished export packages immediately into Azure.
4. Start the import of copied packages into the Azure hosted DBMS instance while the export processes still create new export packages on-premise.
5. Continue until all packages are exported and imported.
6. Perform post-migration tasks.



Migrating from AnyDB to HANA on Azure

This scenario involves migrating any SAP Application from on-premises to Azure while you switch your database from AnyDB to HANA

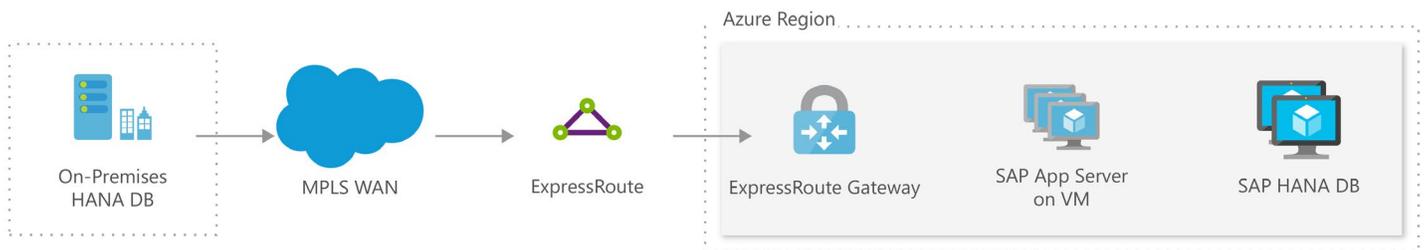
1. Install a new SAP application server and HANA Database on Azure. This can be done through predefined [Azure ARM templates](#).
2. Run export using the SAP Data Migration Options (DMO) system move tool for the on-premise SAP server
3. There are 2 ways of uploading the finished export files to Azure through ExpressRoute from the on premise SAPserver . One is sequential mode and the other is Parallel mode. Parallel mode is used to minimize downtime.
4. Run import using SAP DMO System Move option in Azure.
5. Run all post-migration steps for the Azure SAP App server.



Migrating On-Premises HANA to Azure

This scenario involves migrating an SAP Application from on-premises to Azure while you continue to use HANA (scale-up).

1. Install a new SAP application server and HANA DB server on Azure. This can be done through predefined [Azure ARM templates](#).
2. Take a full database backup for the on-premises SAP HANA Database server.
3. Upload backup files to Azure through ExpressRoute.
4. Restore HANA database in Azure.
5. Run post-migration steps for the Azure SAP App server.
6. (Optional) Set up HANA System Replication to replicate production database to Azure to minimize migration downtime.



There are two main types of migrations to Azure: homogeneous and heterogeneous. In homogeneous migrations, the database and OS remain the same, although software versions may change. With heterogeneous migrations, either the OS or database, or both, change.

Homogeneous migrations

Homogeneous system migrations to the cloud are fairly straight forward. For homogeneous migrations to Azure, use database-specific methods to migrate the database.

Build target systems and set up data synchronization ahead of time. For example, when migrating an SAP system running on Microsoft SQL server 2016/Windows server 2012 R2, you can set up target database synchronization using log shipping. When setting up log shipping, you need to restore the latest full backup in Azure. For HANA databases, you can use HSR. You can also use your current DR techniques to migrate systems to the cloud.

For cases where you cannot use methods similar to log shipping, you can use backup and restore. If you are using the backup and restore method, it is imperative to take a full backup, move the backup to the cloud, and restore the database at the target to extend the downtime. You can recover a full database backup ahead of time, then use a differential or incremental backup to proceed with final cutover. While this method is not optimal, you can use it for non-production systems or for production systems if other methods do not work.

Heterogeneous migrations

Heterogeneous migrations to Azure require more diligence and planning than homogeneous migrations. In general, with adequate time for preparation and testing, almost any customer system of any size can be moved to Azure.

Very large database migration (VLDB) migrations require do considerable skill and attention to detail. For example, you should measure and analyze the impact of table splitting.

Splitting a large table into more than 50 parallel exports may considerably decrease the time needed to export a table, but too many table splits may result in drastically increased Import times. Therefore, the net impact of table splitting must be calculated and tested. An expert licensed OS/DB migration consultant will be familiar with the concepts and tools.

Find out more about [very large database migration to Azure recommendations](#) and [SAP OS/DB migration to SQL server recommendations](#).

Important points to consider

There are some important points to consider when performing a homogeneous or heterogeneous migration to Azure.

ExpressRoute connectivity

It is highly recommended, and mandatory in some cases, to have ExpressRoute connectivity to Azure. ExpressRoute ensure that the connectivity is consistent and predictable. It is also vital to ensure that there are no unnecessary Network Virtual Appliances (NVA) between on-premise corporate networks and Azure.

Disk configuration

You should consider using premium disks for migration to take advantage of local SSD wherever possible. Premium disks and local SSDs provide higher disk throughput and IOPS. In addition, striping multiple disks into one volume may also improve performance. [Find out more](#) about scalability and performance targets.

Compute configuration

It is important to note that compute configuration is just not a function of memory and CPU. Each compute SKU has storage and network throughput metrics. You can find these VM metrics via the Azure portal. To fulfil the characteristics of the SAP workload, you may need a bigger VM SKU even though you may not need the memory and CPU.

Mock migration runs

It is important to conduct mock runs of migrations to avoid surprises. Before production migration, at least pre-production and few non-production systems would have been migrated. The prior runs will provide significant knowledge. With this knowledge, it is important to potentially run mock multiple migrations before the actual production migration. Mock migrations provide more accurate estimates of the actual production downtime.

HANA scale-out migrations using HSR

Hana System Replication (HSR) provides continuous synchronization of a HANA database to a secondary location. HSR can be used to migrate HANA databases. While using HSR to migrate HANA scale out deployment, it is important to note source and target worker nodes are the same. Any

changes to the number of nodes should be done post the cutover process. Refer to [SAP Note 1984882](#)
- Using HANA System Replication for Hardware Exchange with minimum/zero Downtime.

It is possible to migrate from multi node HANA database to a single node HANA database. [SAP Note 2093572](#) provides more details on the procedure.

Should there be a scenario to migrate from single node to multi node using HANA backup, refer to [SAP Note 2130603](#).

Specialized Knowledge

To migrate SAP systems, you should have moderate to high-level knowledge of the source and target IT technologies and environments, including the following:

- Azure Virtual Machines (VMs)
- Azure Storage
- Azure Virtual Network
- Azure ExpressRoute
- Azure Backup
- Azure Site Recovery
- HANA on Azure Large Instances

In addition, we recommend reviewing the following documentation, [Using Azure for hosting and running SAP workload scenarios](#)

Instance Sizing in Azure

Before picking the instance size in Azure, it's important to understand the workload and corresponding sizing. We can broadly classify SAP workloads in the cloud as HANA and non-HANA.

It is well known that HANA workloads are memory intensive. Memory predominantly dictates the size of the HANA compute.

HANA sizing for existing deployments

HANA sizing for existing deployments is relatively straightforward. There are many customers who have optimized their initial sizing and after a period of time, decide that they want to increase size. They take advantage of the cloud for any future growth. Some customers even require more memory when moving to the cloud.

Below is an equation you can use to estimate HANA sizing for existing deployments:

HANA sizing in cloud = current sizing +/- updated requirements

As an illustration of topology changes, for scale-out scenarios, there is a potential to consider options when migrating to the cloud. For example, if the current deployment is a 3+1 512 GB RAM scale-out scenario, one large VM (2TB) would be a good option.

HANA sizing for new deployments

For any new SAP deployments, SAP provides guidelines to size the new HANA system. For any greenfield implementations, SAP Quick Sizer is a great tool to understand what system size you need.

[Find out more](#) about details on scale-out vs scale-up sizing.

Non-HANA sizing

For any database sizing, it's important to consider the current hardware configuration and utilization, as well as EWA. Once you have a reasonable idea of the current configuration, you can select the appropriate SKUs in Azure.

One major aspect to consider when considering non-HANA sizing is storage throughput, IOPS (input output operations per second) and network throughput. VM SKUs have a cap on these parameters. Also, not all Azure VMs are certified to run SAP.

Network Design

When running SAP on Azure, network design has three critical components:

1. Connecting to Azure
2. VNET design
3. IP requirements
4. Network performance considerations

Connecting to Azure

You can connect to Azure two ways:

1. ExpressRoute
2. Site to Site (S2S) connectivity

SAP workloads are critical to your business. Communication between SAP systems running in the cloud and on-premise applications must have tight latency requirements. Your business needs reliable, consistent network performance.

You can use ExpressRoute to access Azure through a dedicated circuit. For HANA large instances, we highly recommend using ExpressRoute. ExpressRoute support bandwidth up to 10 Gbps. [Learn more about ExpressRoute](#)

The next best option is Site to Site connectivity. S2S connectivity is VPN-based and provides bandwidth up to 1.25 Gbps. You can start with S2S during your pilot phase, so you can accelerate pilot scenarios while setting up ExpressRoute connectivity. [Learn more about S2S connectivity.](#)

VNET design

In most of the customer cases, SAP is not the first and the only workload being hosted in Azure. In such cases, the network design for SAP needs to integrate into the existing VNet architecture. As a guiding principle for company-wide Azure network design, customers should follow the principles and architectures included in the [Azure Virtual Datacenter documentation](#). Virtual Network (VNET) is a virtual representation of the Azure network. Azure's VNET enables many types of Azure resources, such as Azure Virtual Machines (VMs), to securely communicate with each other, the Internet, and on-premise networks. [Learn more about Azure Virtual Network key capabilities.](#)

A VNET can be segmented into multiple subnets. Azure VMs are deployed in the subnets. By default, all VMs within subnets of the same VNET can communicate with each other, while also being restricted through NSGs. VMs within subnets of different VNETs cannot communicate with each other unless they are peered or connected.

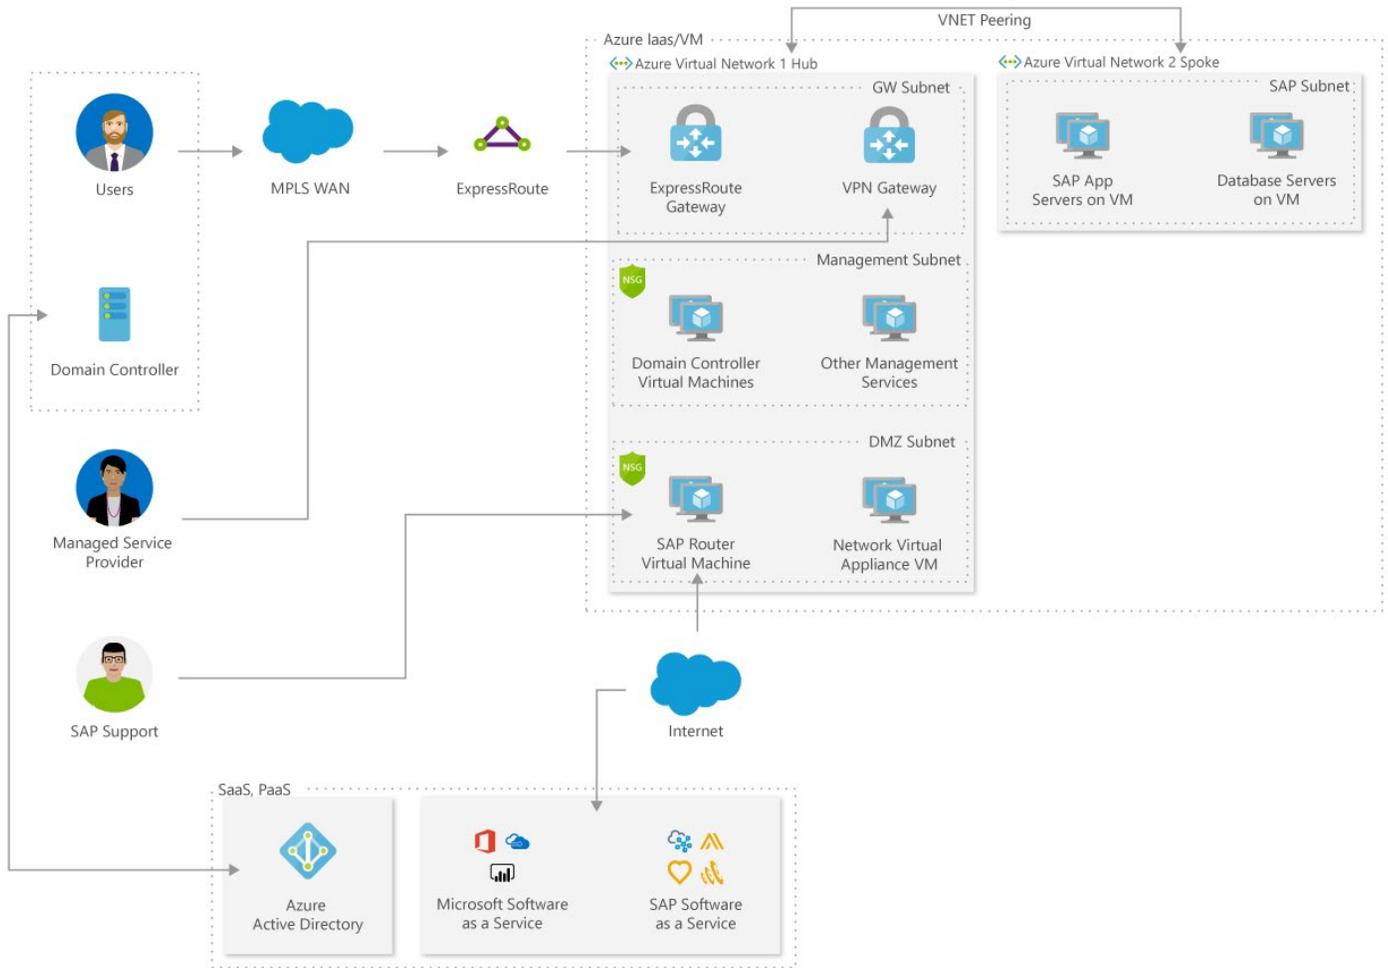
It is assumed that the customer implemented a hub/spoke VNet architecture in Azure where all external facing endpoints like ExpressRoute are going into one Azure VNet that only runs applications which are used for logging and routing of network requests, which is known as a hub VNet. All applications are running in separate VNets (spokes) which are peered with the hub VNet. There might be some deviating architecture items in the case of HANA Large Instances.

The two diagrams below display how to separate DMZ components from SAP components.

While there are many methods, the diagrams highlight two common scenarios.

1. The DMZ and SAP systems are in the same VNET but segmented in separate subnets
2. The DMZ and SAP systems are in different VNets

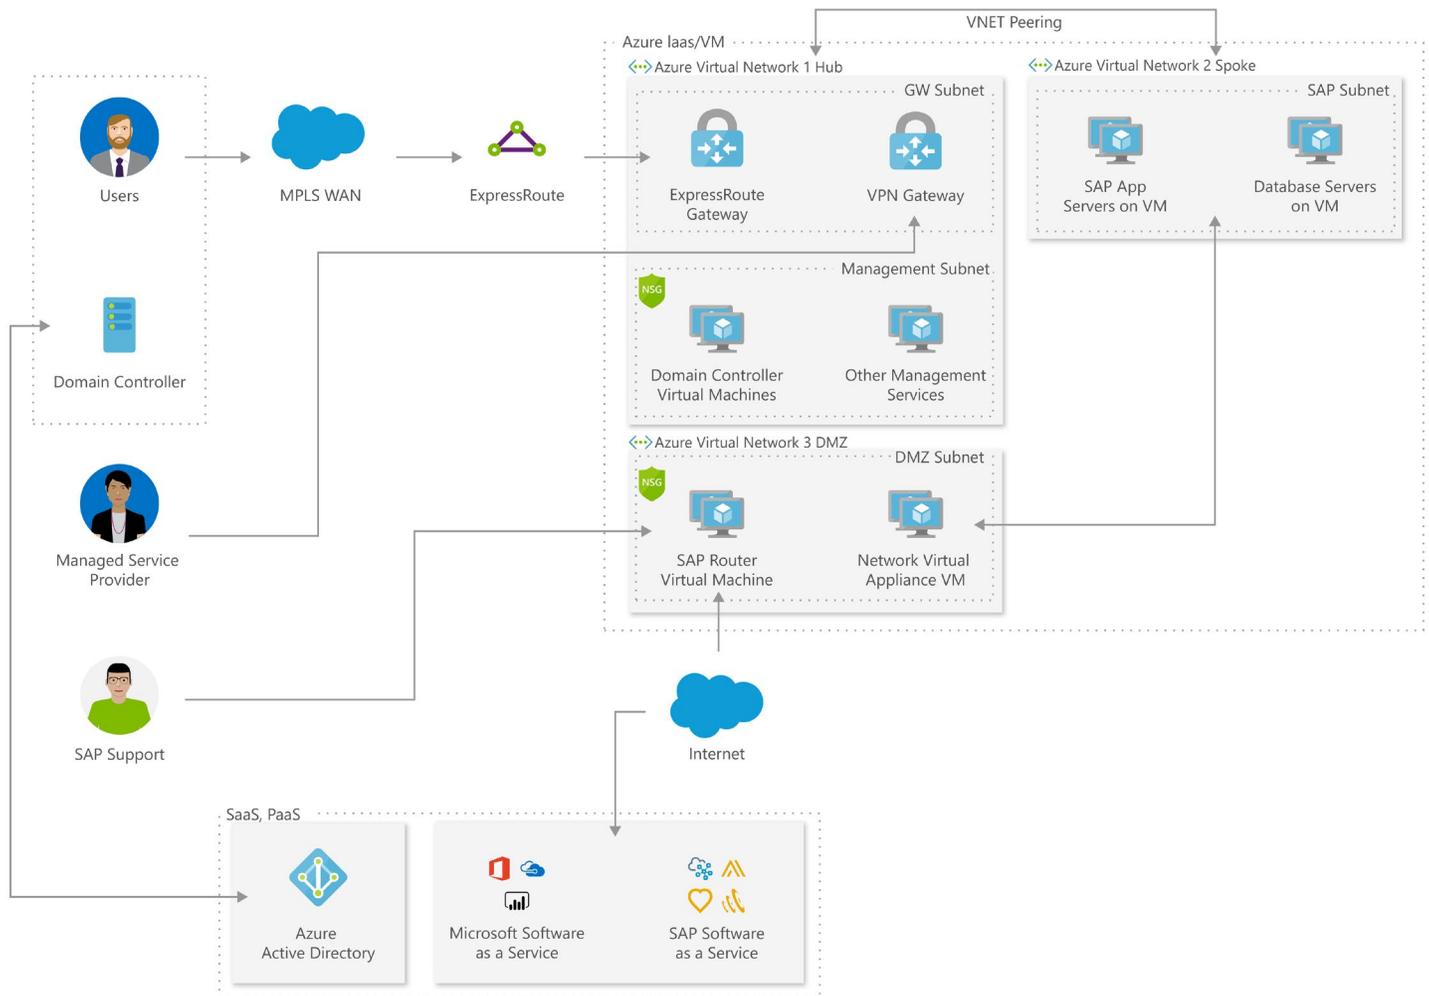
Scenario 1



In this scenario, the DMZ, SAP application, and active directory are in different subnets of the same VNET. Access between subnets is controlled via network security groups.

This diagram includes HANA larger instances (HLI). HLI is deployed in its own subnet and connectivity to HLI occurs via a seamless backend ExpressRoute.

Scenario 2



In this scenario, the DMZ VNET is separated from the SAP VNET. The VNETS can be connected via peering or other connectivity mechanisms.

The diagram in this scenario provides fundamental building blocks and possibilities so you can ask the right questions during design sessions. Please note that the final design may be much more detailed, depending on your requirements.

IP requirements

Since the SAP application performs reverse lookup, many customers prefer using a separate IP for the SAP alias name, which Azure supports. Azure VMs support multiple IP addresses, although most customers will not need more than two. [Learn about Virtual Network Interface Controllers \(vNIC\)](#)

You can add multiple IP addresses using one of the following methods:

- **Adding the IP address of the primary vNIC:** The default gateway is added to the primary vNIC and no additional routes are required. [Learn more about this method](#)
- **Adding additional vNIC(s) to the Azure VM:** By adding additional vNIC(s) to the VM, the IP address is added to the VM. [Learn more about this method](#)

The following steps should be performed post-migration.

Smoke testing

Before the migrated system is released for production operations, you should test all critical business processes. This includes logistics, financial processes, BW reports, and loads in the ECC system, as well as third party screening in GTS, etc. Typically, smoke testing should be read-only business processes.

Typically, SAP ECC systems have middleware connectivity along with a variety of interfaces. Smoke testing should also include connectivity tests to all critical interfaces.

Catch-up background jobs

All cloud migrations result in some downtime. During the SAP system migration downtime, background jobs do not run, and data from critical interfaces, IDOCs, etc., are not processed.

Post-migration, schedule any data processing and catch-up background jobs.

Backup system

A newly migrated system is a fresh state for the system. A full system backup should be triggered after the migration is complete.

Typically, a full system backup takes time so this should be done after opening the system for user transactions. One other way to preserve the fresh state is to take a snapshot of the backup, if possible.

High availability testing

We recommend configuring high availability at all Single Points of Failures (SPoF). Wherever high availability has been configured, perform failover testing to avoid any unintended downtime.

Monitor data synchronization to DR

After migration is complete, set up and monitor data synchronization to the DR setup.

Infrastructure and application monitoring

Migration involves changes at the hardware and software layers. Set up infrastructure and application monitoring and analyze data/alerts for any anomalies.

SAP Support requirements

Should you need any assistance from SAP Support, SAP service marketplace settings and SAP router (customer) information must be up to date.

Detailed migration checks depend on the type of migration, if it is a migration to HANA, or a migration to the cloud, or both. Depending on the type of migration, the checks would differ.

The list below is a high-level checklist to ensure that basic checks at the SAP application, database, and infrastructure levels are in place.

It is recommended to test failover post migration, wherever HANA is configured for SPOFs. This will ensure application downtime is minimized in case of any failures.

SAP application level health checks

You should perform the following health checks at the SAP application level:

- SM21 (logs from all instances)
- ST22
- ST04
- SMQ1/SMQ2
- Mail (SCOT)
- Check RZ03, SMLG, RZ12, SM61, SM14 for availability and correctness of all instances
- Check LC10, DB59
- Check all the important RFC destinations

Database level health checks

You should perform the following health checks at the database level:

- Ensure no errors are detected in the database logs
- Ensure that there is enough database space
- Make sure that no obvious database level blocking is occurring
- Monitor database memory usage

Infrastructure level health checks

You should perform the following health checks at the infrastructure level:

- Check for any network latencies
- Monitor CPU and memory usage

Microsoft Azure provides a robust, resilient, and reliable environment for your SAP applications. This guide helped you understand what choices you need to make, what to expect, how to prepare, and the tools and background information you need to migrate your mission-critical SAP landscapes to Azure. It also introduced the different migration approaches available to you.

If you need any further information, please see [Using Azure for hosting and running SAP workload scenarios](#).