




AZ-120 Domain 1: Migrate SAP Workloads to Azure

AZ-120 skills outline (migrate SAP workloads), SAP on Azure planning guide, as of 2026 | Assessment and sizing (SAPS, HANA memory), migration methods (homogeneous, heterogeneous, DMO, HSR), data transfer, downtime planning, tooling and licensing

AZ-120 Domain 1 SAP migration DMO Sizing

Plan and assess



SAP to Azure migration project with go/no-go gates

Input	Source	Used for
Peak CPU, memory, IOPS per server	OS monitoring, SAP EarlyWatch Alert (EWA), ST06	reference-based SAPS and storage sizing
SAPS rating of current hardware	SAP benchmark site or vendor sheet	$\text{SAPS needed} = \text{rating} \times \text{peak utilisation}$
HANA memory need	/SDF/HDB_SIZING (from any DB), HANA_Sizing for HANA sources	pick M / Mv2 / E-series memory size
Greenfield S/4HANA	SAP Quick Sizer (user- or throughput-based)	SAPS, memory, disk for a new system
Supported VM sizes and SAPS	SAP Note 1928533 (NetWeaver), SAP HANA hardware directory (HANA)	choose only certified sizes
Interfaces and dependencies	SM59, interface lists, Azure Migrate dependency analysis	move-group design, latency-sensitive pairs

Reference-based sizing of the application tier (illustrative SAPS per VM)

PROBLEM

Three on-premises app servers are rated 30,000 SAPS each and peak at 60 % CPU. Keep 30 % headroom. The candidate Azure VM is rated 25,000 SAPS (illustrative: read the real value in SAP Note 1928533). How many VMs?

GIVEN

- $R = 30,000$ SAPS per server, $n = 3$ servers
- $u = 0.60$ peak utilisation
- $h = 1.30$ headroom factor
- $V = 25,000$ SAPS per Azure VM (illustrative)

1 SAPS consumed at peak

$$n \times R \times u = 3 \times 30,000 \times 0.60 = 54,000 \text{ SAPS}$$

Why: size on what is used, not on installed capacity

2 Add headroom

$$54,000 \times 1.30 = 70,200 \text{ SAPS}$$

Why: growth, month-end, failover of one server

3 VM count

$$\text{ceil}(70,200 / 25,000) = \text{ceil}(2.81) = 3 \text{ VMs}$$

Why: round up to whole VMs

4 Check N+1

$$2 \times 25,000 = 50,000 < 54,000$$

Why: losing one VM would overload the others at peak

ANSWER

3 VMs meet the sizing; 4 VMs if the tier must survive one VM loss at peak

CHECK

$$4 \times 25,000 = 100,000 \text{ SAPS; peak load } 54,000 \rightarrow 54 \% \text{ utilisation}$$

→ Choose the migration method

Conditions	HSR	B/R	DMO	Conv	Classic
Source database is already HANA	Y	Y	N	N	N
Database or OS changes	N	N	Y	Y	Y
Target is S/4HANA conversion	N	N	N	Y	N
Need minimal downtime	Y	N	--	--	N
Release upgrade wanted at the same time	N	N	Y	Y	N
Actions					
HANA system replication to Azure, then takeover	X				
HANA backup / restore (files via AzCopy or Data Box)		X			
SUM DMO with System Move			X		
SUM DMO with S/4HANA conversion				X	
Classical migration (SWPM export/import, R3load)					X

Method by source and target (Y = true, - = does not matter)

PORTFOLIO: 8 applications

retire 1	retain 1	rehost 2	relocate 1	replatform 1	repurchase 1	refactor 1
		EFFORT BENEFIT	APPLICATIONS (REASON)			
Retire (1) <i>decommission: switch it off: nobody needs it any more</i>			Old CRM 7.0 archive --- data archived, system shut down			
Retain (1) <i>revisit: keep it where it is for now (not ready, recent upgrade, compliance)</i>			SAP Solution Manager 7.2 --- kept until replaced by SAP Cloud ALM			
Rehost (2) <i>lift and shift: same servers moved to cloud VMs, no code change</i>			BW/4HANA on HANA --- HSR to Azure, takeover at cutover Content Server --- VM + Azure Files / blob			
Relocate (1) <i>hypervisor-level lift and shift: move the virtualised estate (e.g. VMware) as is</i>			Web Dispatcher on VMware --- Azure VMware Solution during transition			
Replatform (1) <i>lift, tinker and shift: a few cloud optimisations: managed database, containers</i>			ECC 6.0 on Oracle/AIX --- DMO with System Move to HANA on Mv2			
Repurchase (1) <i>drop and shop: replace with a SaaS or another product</i>			SAP PO interfaces --- SAP Integration Suite (BTP)			
Refactor (1) <i>re-architect: rebuild cloud-native: serverless, microservices, managed services</i>			S/4HANA target (2027 plan) --- brownfield conversion after the move			

Dots: effort (top row) and cloud benefit (bottom row), 0 to 4, relative to the other Rs. Rehost and relocate move fastest; refactor costs most and gains most.

One SAP landscape placed on the cloud migration Rs

Large data transfer option	Throughput driver	When
ExpressRoute (existing)	circuit bandwidth (1-10 Gbps)	DMO with System Move, HSR initial sync
Site-to-site VPN	about 1 Gbps per tunnel, internet quality	smaller systems, initial copies
AzCopy with parallel jobs	many streams to blob storage	export dumps, backups
Azure Data Box	offline, about 10 days turnaround	tens of TB, weak links, non-production copies
Azure Migrate / ASR replication	block replication of whole VMs	rehost of app servers and non-DB VMs (not HANA)

🕒 Downtime and cutover

Quantity	Formula
----------	---------

Downtime
$$D = t_{pre} + \frac{S}{r} + t_{post} + t_{val}$$

where D = technical downtime (system unavailable) (hours); t_{pre} = shutdown, final backups, interface freeze (hours); S = database size migrated (TB); r = export/import rate measured in the mock migration (TB per hour); t_{post} = post-processing (SGEN, licence key, profile changes) (hours); t_{val} = business validation before go/no-go (hours)

Technical downtime of an export/import migration

📄 Worked example: fits a weekend?

$S = 8$ TB, mock rate $r = 1.2$ TB/h, $t_{pre} = 2$ h, $t_{post} = 3$ h, $t_{val} = 4$ h. $D = 2 + 8 / 1.2 + 3 + 4 = 2 + 6.7 + 3 + 4 = 15.7$ h. A 48-hour weekend leaves 32 h for a roll-back (restart the source system), so the plan passes the gate. If the window were 12 h: use downtime-optimized DMO, archive data before migration, or HSR-based migration when the source is HANA.

Cutover checklist

✓ baseline / mandatory □ recommended – verify in your context

- ✓ Request a new SAP licence key for the new hardware key (message server host)
- ✓ Raise Azure vCPU quotas per VM family before the mock migrations
- ✓ Update DNS (or alias names) for SAP GUI, RFC destinations and printers
- **MUST** Check SAP kernel, OS and DB versions against SAP Notes for Azure (1928533, 2015553)
- **MUST** Rerun the interface tests (RFC, IDoc, file shares) from Azure
- **SHOULD** Keep the source system read-only for rollback until sign-off

🔧 Tooling and licensing

- **Azure Center for SAP solutions** deploy new S/4HANA systems (infrastructure + software install) or register existing ones as a Virtual Instance for SAP solutions; quality checks against best practice

- **SAP Deployment Automation Framework** Terraform + Ansible from GitHub; repeatable landscapes for mock migrations
- **Azure Migrate** discovers and rehosts non-HANA VMs (app servers, Web Dispatcher); not for HANA databases (use HSR or backup/restore)
- **Licensing** SAP licences are BYOL on Azure; Windows Server and SQL Server can use Azure Hybrid Benefit; RHEL and SLES for SAP images are pay-as-you-go or BYOS
- **RISE with SAP on Azure** SAP runs the system in its own subscription; you connect via VNet peering or VPN and keep identity, network and integrations

YOU MAY HEAR <i>"Azure Site Recovery can migrate a running HANA database consistently"</i> BUSTED ✓ ACTUALLY Use HANA system replication or backup/restore for HANA ASR / Azure Migrate are for app servers and other VMs; HANA needs database-level consistency.	YOU MAY HEAR <i>"Size Azure VMs to match the old servers core for core"</i> BUSTED ✓ ACTUALLY Size on SAPS consumed at peak plus headroom Old hardware is often oversized; SAPS normalise across CPU generations.
YOU MAY HEAR <i>"A heterogeneous production migration can be done by anyone"</i> BUSTED ✓ ACTUALLY SAP requires a certified OS/DB migration consultant for production heterogeneous copies SAP Note 82478 describes the rule (check the current note).	YOU MAY HEAR <i>"The first mock migration is a formality"</i> BUSTED ✓ ACTUALLY It is the only reliable source of runtimes Downtime plans and parallelism come from mock measurements, not estimates.

Key takeaways

- 01 Size on consumption**
SAPS × peak + headroom; HANA memory from sizing reports
- 02 Certified only**
SAP Note 1928533 and the HANA directory
- 03 Method follows source and target**
HSR or B/R for HANA, DMO for DB change
- 04 Mock, measure, then commit**
downtime = pre + size/rate + post + validation
- 05 Automate the landscape**
ACSS or SDAF for repeatable builds

Q1 Move ECC on Oracle to Azure and to HANA in one step?
A SUM DMO with System Move (export on-prem, transfer, import on Azure HANA).

Q2 Migrate an existing HANA database to Azure with minimal downtime?
A HANA system replication (async) to an Azure HANA VM, then takeover at cutover.

Q3 Which report estimates HANA memory for a non-HANA source?
A /SDF/HDB_SIZING (SAP Note 1872170).

Q4 Greenfield S/4HANA sizing tool?
A SAP Quick Sizer.

Q5 What must you request after moving the message server to new hardware?
A A new SAP licence key (hardware key changes).

Q6 Which Azure tool rehosts SAP application server VMs?
A Azure Migrate (server migration) or Azure Site Recovery.

Q7 Shrink downtime of a DMO migration of a huge database?
A Downtime-optimized DMO (and archiving before migration).

Q8 Repeatable SAP infrastructure and installation as code from Microsoft?
A SAP Deployment Automation Framework (Terraform + Ansible), or ACSS for a managed experience.