




## DP-300 Domain 1: Plan and Deploy Azure SQL

Microsoft Learn DP-300 skills outline and Azure SQL documentation, as of 2026 | Azure SQL Database vs Managed Instance vs SQL Server on Azure VMs, DTU and vCore, service tiers, serverless, elastic pools, sizing, deployment with az CLI and Bicep, migration paths

DP-300 Domain 1 Azure SQL Plan and implement

### Choose the deployment option

|                                          | Recommended                                     |                                     |                                                   |
|------------------------------------------|-------------------------------------------------|-------------------------------------|---------------------------------------------------|
|                                          | Azure SQL Database                              | SQL Managed Instance                | SQL Server on Azure VM                            |
| Who patches OS and engine                | Microsoft                                       | Microsoft                           | you (automated patching helps)                    |
| Engine compatibility                     | database-scoped, latest engine                  | near 100 % of SQL Server            | 100 %, any version you install                    |
| Instance features (Agent, cross-DB, CLR) | no (elastic jobs instead of Agent)              | ✓                                   | yes, plus OS access, SSRS, SSIS, FCI              |
| Network placement                        | public endpoint + firewall, or private endpoint | inside your VNet (delegated subnet) | your VNet, your NSG                               |
| Max size (as of 2026)                    | 4 TB GP/BC, 128 TB Hyperscale                   | 16 TB GP/BC, 32 TB next-gen GP      | VM disk limits                                    |
| Typical pick                             | new cloud apps, SaaS per tenant                 | lift-and-shift with few changes     | OS access, unsupported features, 3rd-party agents |



Exam reasoning path for 'which Azure SQL option'

### EXAM TIP

Keywords that point to **Managed Instance**: 'minimal code changes', 'SQL Agent jobs', 'cross-database queries', 'native backup restore from .bak', 'VNet-only'. Keywords for **SQL VM**: 'operating system access', 'SQL Server Reporting Services on the same server', 'FILESTREAM', 'specific SQL Server version'. Keywords for **SQL Database Hyperscale**: 'more than 4 TB', 'fast restore regardless of size', 'rapid read scale-out'.

## ③ Purchasing models and service tiers

### Azure SQL Database vCore tiers

|                    | General Purpose<br>provisioned or serverless | Recommended<br>Business Critical<br>provisioned | Hyperscale<br>provisioned or serverless |
|--------------------|----------------------------------------------|-------------------------------------------------|-----------------------------------------|
| Storage            | remote premium storage                       | local SSD                                       | page servers + Azure Storage            |
| Replicas           | 1 compute + spare                            | 4-node availability group                       | 0-4 HA, up to 30 named                  |
| Readable secondary | —                                            | 1 included (read scale-out)                     | HA and named replicas                   |
| Zone redundancy    | yes (selected regions)                       | ✓                                               | ✓                                       |
| In-Memory OLTP     | —                                            | ✓                                               | —                                       |
| Latency            | 5-10 ms IO                                   | 1-2 ms IO                                       | cache-dependent                         |
| Max size           | 4 TB                                         | 4 TB                                            | 128 TB                                  |

| Model             | Unit you buy                                     | Tiers                                                    | Use when                                  | Gotcha                                                                           |
|-------------------|--------------------------------------------------|----------------------------------------------------------|-------------------------------------------|----------------------------------------------------------------------------------|
| DTU               | DTUs (bundled CPU, memory, IO)                   | Basic (5 DTU, 2 GB), Standard (S0-S12), Premium (P1-P15) | small, simple, predictable databases      | no Azure Hybrid Benefit, no reservations, cannot scale storage and compute apart |
| vCore provisioned | vCores + storage GB                              | General Purpose, Business Critical, Hyperscale           | most production workloads                 | pay every hour even idle                                                         |
| vCore serverless  | vCore-seconds used between min and max           | General Purpose, Hyperscale                              | intermittent, unpredictable use, dev/test | cold start after auto-pause; auto-pause only on General Purpose                  |
| Elastic pool      | eDTUs or pool vCores                             | Standard/Premium DTU or GP/BC/Hyperscale vCore           | SaaS with one DB per tenant               | a noisy DB can starve others: set per-DB max                                     |
| Free offer        | 100,000 vCore-seconds and 32 GB per DB per month | General Purpose serverless                               | learning, prototypes                      | up to 10 free DBs per subscription; check the current terms                      |

### 💡 Switching models

DTU to vCore is an online scale operation: as a rough rule, 100 Standard DTUs  $\approx$  1 General Purpose vCore and 125 Premium DTUs  $\approx$  1 Business Critical vCore (Microsoft migration guidance). Hardware: standard-series (Gen5) is the default; premium-series and memory-optimised premium-series give faster CPUs and more memory per vCore.

## 🔗 Size and deploy

| az CLI                                                                                                                                                                                                       | What it does                                                             |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|
| <pre>az sql server create -n sql-app-prod -g rg-data -l westeurope --enable-ad-only-auth --external-admin-principal-type Group --external-admin-name sql-admins --external-admin-sid &lt;objectId&gt;</pre>  | logical server with Microsoft Entra-only authentication (no SQL logins)  |
| <pre>az sql db create -s sql-app-prod -g rg-data -n app --edition GeneralPurpose --family Gen5 --capacity 2 --compute-model Serverless --min-capacity 0.5 --auto-pause-delay 60 --zone-redundant false</pre> | serverless GP database, pauses after 60 idle minutes                     |
| <pre>az sql elastic-pool create -s sql-app-prod -g rg-data -n pool1 --edition GeneralPurpose --family Gen5 --capacity 8 --db-max-capacity 4 --db-min-capacity 0</pre>                                        | 8-vCore pool, each database capped at 4 vCores                           |
| <pre>az sql db update ... --service-objective HS_Gen5_4</pre>                                                                                                                                                | scale up or change tier (online; a short reconnect at the end)           |
| <pre>az sql mi create -n sqlmi-prod -g rg-data --subnet &lt;subnetId&gt; --capacity 8 --edition GeneralPurpose --family Gen5 --storage 512GB --license-type BasePrice</pre>                                  | Managed Instance in a delegated subnet; BasePrice = Azure Hybrid Benefit |
| <pre>az sql vm create -n sqlvm01 -g rg-data --license-type AHUB --sql-mgmt-type Full</pre>                                                                                                                   | register a VM with the SQL IaaS Agent extension                          |

```
main.bicep: serverless database on an existing logical server
resource srv 'Microsoft.Sql/servers@2023-08-01' existing = {
  name: 'sql-app-prod'
}
resource db 'Microsoft.Sql/servers/databases@2023-08-01' = {
  parent: srv
  name: 'app'
  location: resourceGroup().location
  sku: { name: 'GP_S_Gen5_2', tier: 'GeneralPurpose' }
  properties: {
    autoPauseDelay: 60 // minutes, -1 disables
    minCapacity: json('0.5')
    requestedBackupStorageRedundancy: 'Zone'
    zoneRedundant: false
  }
}
```

## Worked example: elastic pool or single databases

### PROBLEM

A SaaS app has 20 tenant databases. Each peaks at 4 vCores for about one hour a day, at different times; telemetry shows the busiest hour needs 12 vCores in total across all tenants. Pool or singles, and how big?

#### 1. Singles sized for their own peak

⇒  $n \times p = 20 \times 4 = 80$  vCores provisioned  
why: each single database must hold its own peak

#### 2. Pool sized for the concurrent peak

⇒  $c \times (1 + h) = 12 \times 1.25 = 15 \rightarrow$  round up to the next pool size, 16 vCores  
why: peaks do not coincide, so they share capacity

#### 3. Compare

⇒  $16 / 80 = 20\%$  of the singles' vCores → 80 % fewer vCores  
why: same price per vCore, so cost follows vCores

#### 4. Protect neighbours

⇒ `--db-max-capacity 4 --db-min-capacity 0`  
why: a tenant cannot take more than its peak; nothing reserved when idle

### GIVEN

- $n = 20$  databases
- $p = 4$  vCores peak per database
- $c = 12$  vCores concurrent peak, measured
- $h = 25\%$  headroom target
- pool vCore unit price = single-database vCore unit price (vCore model)

### ANSWER

Elastic pool of 16 vCores with per-database max 4: about 80 % cheaper than 20 single 4-vCore databases.

Check: If  $c$  approaches  $n \times p$  (peaks coincide), the pool saves nothing: keep singles or split into several pools.

## SQL Server on Azure VMs: build it right

### SQL VM configuration baseline


✓ baseline / mandatory    □ recommended – verify in your context

- ✓ Memory-optimised series (E, M) or constrained-vCPU sizes to cut per-core licences constrained vCPU keeps memory and IO, halves or quarters the licensed cores
- ✓ Premium SSD v2 or Premium SSD (or Ultra Disk for log latency); data and log on separate disks
- ✓ Host caching: ReadOnly on data disks, None on log disks
- ✓ tempdb on the local ephemeral SSD (D: on Windows) where the VM size has one the IaaS Agent extension recreates the folder at boot
- **MUST** Storage Spaces stripe several disks for IOPS; 64 KB allocation unit for data and log volumes
- **MUST** Register with the SQL IaaS Agent extension: licence mode, automated patching window, automated backup
- **SHOULD** Enable instant file initialisation and lock pages in memory for the service account
- **SHOULD** Place AG replicas in different availability zones or an availability set

### ⚠ Licensing traps

Pay-as-you-go SQL VM images include the SQL licence per vCore-hour; switching to Azure Hybrid Benefit (`--license-type AHUB`) needs licences with Software Assurance. A free passive secondary for HA/DR is a licence benefit of Software Assurance; Developer and Express editions cost nothing but Developer is not for production.

## → Migrate to Azure SQL



| Target                 | Offline method                                                              | Online (minimal downtime) method                                            |
|------------------------|-----------------------------------------------------------------------------|-----------------------------------------------------------------------------|
| Azure SQL Database     | BACPAC import (SqlPackage), DMS offline                                     | transactional replication (source as publisher), DMS online where supported |
| SQL Managed Instance   | native RESTORE FROM URL of .bak in Blob Storage, DMS offline                | Managed Instance link (distributed AG), Log Replay Service, DMS online      |
| SQL Server on Azure VM | backup/restore, detach/attach, Azure Migrate lift-and-shift of the whole VM | log shipping, Always On AG with an Azure replica, DMS                       |

### YOU MAY HEAR

"Azure SQL Database restores native .bak files"

**BUSTED**

### ✓ ACTUALLY

Only Managed Instance and SQL VMs restore .bak  
SQL Database imports BACPAC or uses DMS / replication; a .bak needs MI (RESTORE FROM URL) or a VM.

### YOU MAY HEAR

"Serverless always costs less"

**BUSTED**

### ✓ ACTUALLY

Serverless wins only when the database is idle or low for long periods  
Its per-vCore-second price is higher than provisioned; a busy 24/7 database is cheaper provisioned.

### YOU MAY HEAR

"Business Critical is just faster storage"

**BUSTED**

### ✓ ACTUALLY

It is a four-replica availability group  
That brings fast failover, a free readable secondary (ApplicationIntent=ReadOnly) and In-Memory OLTP.

### YOU MAY HEAR

"Hyperscale cannot go back"

**BUSTED**

### ✓ ACTUALLY

Reverse migration from Hyperscale to General Purpose is supported within 45 days  
Only for databases migrated into Hyperscale from another tier; check current limits.

**BACPAC** — A zip of schema and data exported by SqlPackage or the portal; the portable format for Azure SQL Database import and export (not transactionally consistent unless copied first).

**Compatibility level** — Database setting (for example 150, 160) that selects query-optimiser and T-SQL behaviour; keep the source level at migration, raise it after a Query Store baseline.

## 🚩 Recap and self-test

*Tick a box only when you could explain the point to someone else*



### Database, instance or VM

PaaS DB first; MI for instance features; VM for OS access



### Business Critical = 4-replica AG

low latency, free readable secondary



### Serverless for intermittent load

auto-pause only on General Purpose



### vCore over DTU for production

Hybrid Benefit, reservations, independent storage



### Hyperscale for > 4 TB

up to 128 TB, fast snapshot backups, named replicas



### Pool when peaks do not coincide

size for concurrent peak, cap each DB

**Q1 Which Azure SQL option supports SQL Agent, cross-database queries and linked servers without a VM?**

**A** SQL Managed Instance; SQL Database uses elastic jobs instead of Agent and has no cross-database queries (only elastic query, limited).

**Q2 A database will grow to 30 TB. Which tier?**

**A** Hyperscale (up to 128 TB); General Purpose and Business Critical stop at 4 TB for a single database.

**Q3 Which service tier gives one readable secondary at no extra cost?**

**A** Business Critical (and Premium DTU), via read scale-out with `ApplicationIntent=Read Only`.

**Q4 Which purchasing model allows Azure Hybrid Benefit?**

**A** vCore (provisioned; not serverless for SQL Database, and not DTU).

**Q5 What does auto-pause do and where is it available?**

**A** Stops compute billing after the configured idle delay (storage still billed); General Purpose serverless only.

**Q6 Online migration from SQL Server 2019 to Managed Instance with the smallest cutover?**

**A** Managed Instance link (distributed availability group), or Log Replay Service / DMS online.

**Q7 How should host caching be set for a SQL VM log disk?**

**A** None (log writes are sequential and must be durable); `ReadOnly` on data disks.

**Q8 Where does tempdb go on a SQL VM?**

**A** On the local ephemeral SSD of the VM size (D: on Windows): fastest, and tempdb is recreated at restart anyway.

**Q9 What is a logical server in Azure SQL Database?**

**A** A management container for databases, firewall rules and admins; it is not an instance and has no shared buffer pool.