

# Oracle 1Z0-1147

**Oracle AI Database@Azure Architect Professional**

**For More Information – Visit link below:**

**<https://www.examsempire.com/>**

**Product Version**

- 1. Up to Date products, reliable and verified.**
- 2. Questions and Answers in PDF Format.**



**<https://examsempire.com/>**

**Visit us at: <https://www.examsempire.com/1z0-1147>**

# Latest Version: 6.0

## Question: 1

An application team is deploying a cloud-native microservices application on Azure that has highly variable, unpredictable workloads and requires database auto-scaling with zero database administration effort. Which Oracle AI Database@Azure service should the architect provision to satisfy these requirements?

- A. Oracle Exadata Database Service on Dedicated Infrastructure
- B. Oracle Autonomous AI Database
- C. Oracle Exadata Database Service on Exascale Infrastructure
- D. Oracle Base Database Service

**Answer: B**

Explanation:

Oracle Autonomous AI Database is specifically designed for workloads where elasticity and minimal administrative overhead are primary requirements. Oracle describes it as a fully autonomous, self-governing database service that handles provisioning, backups, patching, upgrading, performance management, and resource scaling without traditional database-administration effort.

For unpredictable workloads, compute auto scaling is particularly significant. When enabled, Autonomous AI Database can automatically consume additional CPU and I/O resources as workload demand rises without requiring an administrator to manually resize the database. Storage can likewise be scaled and can use separate auto-scaling functionality.

Exadata Database Service on Dedicated Infrastructure provides extremely high performance and consolidation capabilities but retains customer database-management responsibilities. Exascale Infrastructure improves Exadata's elastic infrastructure model but is still not equivalent to Autonomous Database's fully managed, self-governing operational model. Base Database Service similarly requires substantially more customer administration.

Therefore, a cloud-native application with volatile usage and a requirement for essentially zero database administration maps directly to Oracle Autonomous AI Database .

Study Guide reference: Provision and operate Oracle AI Database@Azure Resources — Autonomous AI Database, elastic scaling, automated management and serverless operations.

## Question: 2

What is the default CIDR block range used to create the backup subnet for an Exadata VM Cluster in Oracle AI Database@Azure if the field is left blank during provisioning?

- A. 10.0.252.0/22
- B. 192.168.252.0/24
- C. 192.168.252.0/22
- D. 172.16.252.0/22

**Answer: C**

Explanation:

During Exadata VM Cluster provisioning, the administrator selects the client delegated subnet and can explicitly provide a CIDR range for the Exadata backup network. If no backup-subnet CIDR is supplied, Oracle automatically uses the documented default of 192.168.252.0/22 .

The backup subnet serves Oracle database backup-network requirements and is architecturally distinct from the Azure delegated client subnet. Oracle's deployment-planning guidance emphasizes that the default backup CIDR is not part of the Azure VNet CIDR allocation. Nevertheless, address planning remains important. In disaster-recovery designs involving multiple Exadata deployments—especially cross-availability-zone or cross-region configurations—the administrator should explicitly choose non-overlapping backup CIDR blocks so that replication, routing, and backup traffic do not encounter conflicting address ranges.

Option B uses the correct network prefix but an incorrect /24 mask. Options A and D use different base networks and therefore do not match the Oracle provisioning default.

Accordingly, the correct response is C: 192.168.252.0/22 .

Study Guide reference: Configure Oracle AI Database@Azure Networking — Exadata VM Cluster client and backup networks, CIDR planning, default backup-subnet configuration and non-overlapping network design.

### Question: 3

You want to temporarily minimize compute license billing costs for an idle guest Exadata VM Cluster on Oracle AI Database@Azure without destroying the virtual machines or losing database storage. What is the minimum number of OCPUs per database node you can scale the cluster down to in order to achieve this state?

- A. One OCPU
- B. Zero OCPUs
- C. Two OCPUs
- D. Four OCPUs

**Answer: B**

Explanation:

Oracle Exadata VM Cluster compute resources can be scaled down to zero OCPUs when the objective is to suspend compute consumption without deleting the VM Cluster definition or its persistent database storage. Oracle explicitly notes that simply stopping VMs does not stop the relevant compute billing; to stop VM Cluster compute-related billing, the enabled OCPU—or ECPU on applicable generations—allocation must be reduced to zero.

At zero enabled compute, the guest VMs are shut down while the infrastructure configuration and persistent database storage remain available. The administrator can later restore service by scaling the enabled compute allocation back above zero. This makes the feature useful for development, testing, seasonal, or otherwise intermittent environments where deleting and recreating the database environment would be operationally undesirable.

Values of one, two, or four OCPUs still represent allocated active compute and therefore do not achieve the requested zero-compute billing state.

On Exascale Infrastructure, Oracle documents the equivalent ECPU behavior explicitly: enabled ECPUs can be scaled to zero, shutting down VM Cluster VMs and stopping enabled-ECPU usage billing.

Study Guide reference: Provision and operate Oracle AI Database@Azure Resources — VM Cluster scaling, OCPU/ECPU allocation and cost control.

### Question: 4

Which two network routing capabilities or controls are supported when configuring User Defined Routes (UDR) within the delegated subnets of Oracle AI Database@Azure? (Choose two.)

- A. Custom network routing paths between on-premises systems and Exadata databases
- B. Automatic dynamic path selection across dual Azure ExpressRoute circuits
- C. Granular routing control of traffic flowing across hybrid cloud infrastructure
- D. Bypassing OCI control plane scheduled patching routes

**Answer: A, C**

Explanation:

Azure User-Defined Routes provide explicit Layer-3 forwarding control for traffic entering or leaving an Oracle AI Database@Azure delegated subnet. Oracle's documented network topologies use UDRs to direct application, on-premises, and database traffic through prescribed next hops such as an Azure Firewall or Network Virtual Appliance. This enables architects to construct custom network routing paths between on-premises systems and Exadata databases and exercise granular routing control across hybrid infrastructure .

In a Hub-and-Spoke deployment, for example, UDRs on application spokes can force database-bound traffic through the hub firewall, while corresponding routes on the Oracle AI Database@Azure spoke provide a symmetric return path. This is an architectural routing control, not a dynamic-routing protocol. ExpressRoute path selection is governed by BGP and ExpressRoute configuration rather than by a UDR automatically choosing between circuits.

Likewise, UDRs do not override Oracle-managed control-plane connectivity or scheduled lifecycle-management routes. Those functions remain under the service architecture.

Therefore, A and C are the supported routing capabilities. Study Guide reference: Configure Oracle AI Database@Azure Networking — UDRs, delegated subnets, hybrid connectivity, Hub-and-Spoke routing.

### Question: 5

A multicloud database administrator is deploying an Oracle Exadata VM Cluster in Oracle AI Database@Azure and has configured a client delegated subnet with the CIDR range 10.0.0.0/24. The administrator needs to verify which specific IP addresses within this delegated subnet CIDR range are automatically reserved by the system for database networking services. Based on the official network design requirements, which IP address ranges are reserved for these platform networking services? Choose ONE.

- A. 10.0.0.0 to 10.0.0.7, 10.0.0.32 to 10.0.0.39, and 10.0.0.254

- B. 10.0.0.0 to 10.0.0.3, 10.0.0.8 to 10.0.0.15, and 10.0.0.255
- C. 10.0.0.0 to 10.0.0.4, 10.0.0.10 to 10.0.0.15, and 10.0.0.255
- D. 10.0.0.1 to 10.0.0.5, 10.0.0.12 to 10.0.0.19, and 10.0.0.250

**Answer: B**

Explanation:

Oracle AI Database@Azure has specific client-subnet IP consumption requirements that must be considered during capacity planning. For the delegated client subnet, 13 IP addresses are reserved for networking services regardless of the number of VM Clusters deployed in that subnet .

Microsoft's Oracle AI Database@Azure networking documentation defines these 13 addresses as the first four addresses, the ninth through sixteenth addresses, and the final address of the subnet. Applied to 10.0.0.0/24, the reserved ranges are therefore 10.0.0.0–10.0.0.3, 10.0.0.8–10.0.0.15, and 10.0.0.255. These platform reservations must be counted in addition to addresses needed for VM nodes and SCAN addresses. For example, every VM consumes multiple client-network addresses, and each VM Cluster requires SCAN addresses. Failing to account for the reserved block can therefore result in insufficient address capacity even when a subnet initially appears large enough.

Only option B precisely represents the documented reserved address positions.

Study Guide reference: Configure Oracle AI Database@Azure Networking — delegated subnet planning, client subnet sizing, reserved networking-service IP addresses and Exadata VM Cluster capacity.

## Question: 6

By default, how many delegated subnets can Oracle AI Database@Azure use before a service limit increase must be requested?

- A. 5
- B. 1
- C. 3
- D. 10

**Answer: A**

Explanation:

Oracle AI Database@Azure has a default regional service limit of five delegated subnets . Once a deployment requires more than five delegated subnet links in a region, the customer must request a service-limit increase.

This aggregate regional limit must not be confused with the separate VNet design rule. An individual Azure VNet can contain a maximum of one delegated subnet for Oracle AI Database@Azure, but an organization can have several VNets and therefore several delegated subnets across the same region up to the default service limit.

Oracle identifies the relevant service-limit metadata as:

Service: Multicloud

Resource: Delegated Subnet Multicloud Links

Limit name: azure-delegated-subnet-count

The limit operates at regional scope. Therefore, creating additional peerings or Route Servers does not bypass it; the correct approach is to request additional service capacity through Oracle's service-limit process.

Values of 1, 3, and 10 do not represent the default regional limit.

Therefore, A. 5 is correct.

Study Guide reference: Configure Oracle AI Database@Azure Networking — delegated subnet quotas and service-limit increases.

## Question: 7

You have configured SAML-based identity federation between Microsoft Entra ID and your OCI IAM identity domain for Oracle AI Database@Azure. A database administrator now needs to perform container database and pluggable database management tasks. Where should the administrator sign in?

- A. The Azure Cloud Shell
- B. The Azure Marketplace private offer page
- C. Oracle Enterprise Manager standalone console
- D. The OCI Console

**Answer: D**

Explanation:

The administrator must use the OCI Console for database-management operations that remain exposed through Oracle Cloud Infrastructure. Oracle AI Database@Azure deliberately integrates many infrastructure lifecycle operations into Azure, but the service still maintains an associated OCI tenancy and OCI control plane. Consequently, certain Oracle-native operations—particularly detailed Container Database (CDB) and Pluggable Database (PDB) administration—require access to OCI.

The configured Microsoft Entra ID federation eliminates the need for a separate set of OCI credentials. Oracle documents federation as an optional capability that enables users to authenticate to the OCI tenancy by using their existing Microsoft Entra ID credentials. Oracle specifically notes that although most daily operations can be performed within Azure, some database-management tasks require OCI Console access.

Azure Cloud Shell operates Azure CLI/PowerShell workflows and is not the required CDB/PDB management interface. The Marketplace page controls purchasing and subscriptions, while a standalone Enterprise Manager installation is not the primary management portal for this service.

Therefore, D is correct.

Study Guide reference: Provision and operate Oracle AI Database@Azure Resources — OCI Console access and database lifecycle management.

## Question: 8

If a customer does not configure a customer-managed key, where does Oracle AI Database@Azure store the default TDE master encryption key for an Exadata database?

- A. Azure Key Vault
- B. OCI Vault

- C. Oracle Wallet
- D. Oracle Key Vault

**Answer: C**

Explanation:

The default encryption-key model for an Exadata database on Oracle AI Database@Azure is Oracle-Managed Keys (OMK) . Under this model, Oracle automatically generates the Transparent Data Encryption master encryption key and stores it in an Oracle Wallet secured within the database environment.

Oracle manages the complete lifecycle of this default key, including generation, secure storage, rotation, backup, and restoration. No additional customer-side key-management infrastructure is required.

OCI Vault, Oracle Key Vault, and Azure Key Vault represent customer-managed key alternatives. When an organization selects one of these mechanisms, the TDE master key is managed outside the local Oracle Wallet, enabling centralized customer control, auditing, rotation policy, and separation of duties. However, none is the default unless the customer explicitly configures a customer-managed key architecture.

Oracle's current security documentation explicitly identifies Oracle Wallet as the default encryption-key setting for Exadata Database Service on Oracle AI Database@Azure.

Therefore, C. Oracle Wallet is correct.

Study Guide reference: Secure Oracle AI Database@Azure — TDE, Oracle-managed keys and Oracle Wallet.

### Question: 9

Where is a long-term retention backup in the Oracle Database Autonomous Recovery Service automatically moved after its first 24 hours of storage?

- A. OCI Archive Storage tier
- B. Azure Archive Blob Storage tier
- C. OCI Legacy Backup tier
- D. OCI Infrequent Access Storage tier

**Answer: D**

Explanation:

Oracle Database Autonomous Recovery Service supports Long-Term Retention (LTR) backups for organizations that need to preserve full database backups for compliance, regulatory, legal, or business requirements. The correct destination is the OCI Object Storage Infrequent Access tier .

The question source specifies that the LTR backup initially remains in the standard storage tier and transitions after the first 24 hours to Infrequent Access storage for the remainder of its retention period. Oracle's Recovery Service documentation independently confirms that long-term retention backups are stored in the Object Storage Infrequent Access tier and can be retained for periods extending up to ten years.

The Infrequent Access tier is appropriate because LTR backups are retained for extended periods but are generally restored less frequently than operational backups. Archive Storage is not the managed

destination used by Recovery Service for this workflow. Azure Archive Blob Storage is likewise not the Oracle-managed OCI storage tier identified for these backups, and there is no "OCI Legacy Backup" storage tier.

Therefore, D is correct.

Study Guide reference: Provision and operate Oracle AI Database@Azure Resources — Autonomous Recovery Service, LTR backups, retention and storage lifecycle.

## Question: 10

Your team is preparing to migrate an on-premises database to Autonomous Database Serverless on Oracle Database@Azure using ZDM and wants to ensure that the Data Pump dump files are staged entirely within the Microsoft Azure subscription. Which storage option should you configure in the ZDM response file as the data transfer medium to meet this requirement?

- A. OCI Object Storage
- B. OCI File Storage Service
- C. Azure Files NFS share
- D. Azure Blob Storage API

**Answer: C**

Explanation:

The correct staging medium is an Azure Files NFS share . Oracle's ZDM documentation for Autonomous AI Database Serverless on Oracle Database@Azure explicitly supports Azure Files as an NFS file-share source for Data Pump migration staging.

In the logical migration workflow, the NFS share is made accessible to the source database and attached to the Autonomous AI Database target. Data Pump export files can therefore be written to and read from the shared location without being staged in OCI storage. Oracle specifically documents mounting Azure Files NFSv4.x and exposing the path through a database directory for Data Pump operations. OCI Object Storage and OCI File Storage would place the staging data within OCI, violating the explicit requirement that dump files remain in the Microsoft Azure subscription. Azure Blob Storage can be accessed by Autonomous Database through DBMS\_CLOUD for supported object-storage workflows, but the standard ZDM NFS-based Data Pump staging mechanism described in this question uses a mounted filesystem rather than a direct Blob REST API.

Therefore, C is correct.

Study Guide reference: Migrate Oracle Databases to Oracle AI Database@Azure — ZDM, Autonomous AI Database Serverless, Data Pump and Azure Files NFS.

**Thank You for Trying Our Product**

**Special 16 USD Discount Coupon: NSZUBG3X**

**Email:** [support@examsempire.com](mailto:support@examsempire.com)

**Check our Customer Testimonials and ratings  
available on every product page.**

**Visit our website.**

**<https://examsempire.com/>**