Oracle.1z0-1147.v2026-10-02.q45

Exam Code:	1z0-1147
Exam Name:	Oracle AI Database@Azure Architect Professional
Certification Provider:	Oracle
Free Question Number:	45
Version:	v2026-10-02
# of views:	104
# of Questions views:	450
https://www.freecram.net/torrent/Oracle.1z0-1147.v2026-10-02.q45.html	

NEW QUESTION: 1
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: (SHOW ANSWER)
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.

NEW QUESTION: 2

An enterprise is implementing a security architecture for Oracle AI Database@Azure where database administrators must manage resources natively through the Azure Portal while authenticating seamlessly via Single Sign-On (SSO) with Microsoft Entra ID. To facilitate this, the organization plans to configure optional Identity Federation and Role-Based Access Control (RBAC) mapping between Azure and Oracle Cloud Infrastructure (OCI). The deployment team needs to coordinate the exact user personas and permissions required on both cloud platforms to configure this identity federation. To complete this task, which combination of administrative personas and cloud-specific permissions must be assigned? Choose ONE.

A. Azure Billing Account Owner with OCI root compartment tenancy administrator privileges

B. Azure Entra ID Administrator with OCI Identity Domain Administrator privileges

C. Azure Subscription Owner with OCI Support Owner privileges

D. Azure User Access Administrator with OCI Security Zone Administrator privileges

Answer: (SHOW ANSWER)

Oracle's Oracle AI Database@Azure prerequisites explicitly identify the personas required for optional identity federation. On the Azure side, the required persona is an Entra ID Administrator with appropriate Microsoft Entra administrative permissions. On the OCI side, the corresponding persona is an Identity Domain Administrator .

This combination is necessary because federation spans both identity systems. Microsoft Entra ID must be configured with the OCI enterprise application, SAML authentication details, groups/users, and provisioning configuration. OCI IAM Identity Domains must then trust the Entra identity provider, establish the corresponding IdP configuration and policies, and support any required user provisioning.

Oracle's prerequisite matrix specifically lists Azure permissions including External Identity Provider Administrator, Groups Administrator, and User Administrator capabilities, with OCI Identity Domain Administrator permission on the Oracle side.

Billing Account Owner, Subscription Owner, Support Owner, and Security Zone Administrator roles serve different administrative purposes and do not jointly provide the identity-domain configuration authority required for federation.

Therefore, B exactly matches Oracle's documented administrative personas.

Study Guide reference: Secure Oracle AI Database@Azure - identity federation, Microsoft Entra ID, OCI IAM Identity Domains, SSO and RBAC.

NEW QUESTION: 3

You are configuring logging for an Exadata Infrastructure resource in Oracle AI Database@Azure and attempt to save a second diagnostic setting that routes the same log categories to your primary Log Analytics workspace. What will occur when you attempt to save this configuration?

A. The configuration will fail because you cannot create two diagnostic settings that send duplicate log categories to the same destination.

B. The configuration will succeed but will trigger an automatic alert in Microsoft Sentinel regarding duplicate logging.

C. The Azure Portal will automatically consolidate both configurations into a single active setting without warning.
D. The configuration will succeed but will automatically override the existing log retention policy on your workspace.
Answer: (SHOW ANSWER)
Microsoft's Oracle Exadata Database on Dedicated Infrastructure observability guidance explicitly defines a no-duplicate-destination rule for these resources: two diagnostic settings cannot send the same log categories to the same destination. Each diagnostic-setting configuration must therefore remain unique.
An Exadata resource can have multiple diagnostic settings-Microsoft documents a maximum of five per resource-but those settings are intended to provide different category/destination combinations. If an administrator attempts to create another setting with the same log categories targeting the same Log Analytics workspace, the duplicate configuration is rejected rather than automatically consolidated.
Microsoft Sentinel does not serve as a conflict-management mechanism for the diagnostic-settings configuration, and the platform does not resolve the collision by modifying Log Analytics retention. Retention is a separate workspace/data-management concern.
This rule prevents unintended duplicate ingestion and avoids unnecessary telemetry duplication and associated ingestion cost.
Therefore, A accurately describes the outcome.
Study Guide reference: Provision and operate Oracle AI Database@Azure Resources - Exadata logging, Azure Monitor diagnostic settings, Log Analytics and observability constraints.

NEW QUESTION: 4
Which three are valid destination types that can be configured when setting up Azure Monitor diagnostic settings for Exadata VM Clusters in Oracle AI Database@Azure? (Choose three.)
A. Azure Storage account
B. Oracle Enterprise Manager Repository
C. Azure Event Hubs
D. Azure Log Analytics workspace
E. OCI Object Storage bucket
F. Azure SQL Database
Answer: (SHOW ANSWER)
Oracle AI Database@Azure integrates its operational telemetry with Azure Monitor diagnostic settings so organizations can apply existing Azure observability, compliance, and SIEM architectures to Oracle database infrastructure.
Three standard destinations are relevant here. An Azure Log Analytics workspace receives telemetry for centralized querying, dashboarding, alerting, and Kusto Query Language analysis. An Azure Storage account can archive logs or metrics for retention, compliance, or later processing. Azure Event Hubs provides a streaming destination for forwarding telemetry to SIEM products, custom consumers, event-processing pipelines, or other downstream platforms.

Oracle's Oracle AI Database@Azure monitoring documentation specifically lists Log Analytics Workspace, Storage Account, and Event Hub among the available Azure Monitor diagnostic-routing destinations. Oracle Enterprise Manager can monitor Oracle environments but is not an Azure Monitor diagnostic-settings destination. OCI Object Storage is likewise not selected directly as a destination from the Azure diagnostic- settings workflow. Azure SQL Database is a database service rather than a supported diagnostic sink.
Therefore, the three correct selections are A, C, and D .
Study Guide reference: Provision and operate Oracle AI Database@Azure Resources - monitoring, Azure diagnostic settings, Log Analytics, Storage accounts and Event Hubs.

NEW QUESTION: 5
You have provisioned two VNets for Oracle AI Database@Azure in the same Oracle Availability Zone. One VNet uses advanced network features and the other uses default networking, and you now want to peer them.
What should you expect?
A. VNet peering between them isn't supported and may cause asymmetric routing and traffic loss.
B. Peering will succeed with no additional configuration required.
C. Peering will automatically upgrade the default-networking VNet.
D. Peering will disable advanced network features on both VNets.
Answer: (SHOW ANSWER)
This configuration is explicitly identified as unsupported in Microsoft's Oracle AI Database@Azure network- planning guidance. If two Oracle AI Database@Azure VNets are deployed in the same Oracle Availability Zone and one uses Advanced Networking while the other uses the default networking implementation, VNet peering between them can introduce asymmetric routing . Asymmetric routing occurs when outbound and return traffic traverse different network paths. This is particularly problematic with stateful network elements and the differing routing models associated with the two Oracle AI Database@Azure network implementations. The result can be data-plane traffic loss toward the VNet using Advanced Networking.
VNet peering itself does not upgrade a previously provisioned VNet to Advanced Networking, nor does it downgrade the Advanced Networking environment. Advanced Networking capability is determined by the feature configuration and the state of the delegated subnet before Oracle resources/VNICs are provisioned.
Therefore, the two environments should not be directly peered in this mismatched state. The correct answer is A .
Study Guide reference: Configure Oracle AI Database@Azure Networking - Advanced Networking compatibility and VNet peering.

NEW QUESTION: 6
Which of the following does NOT describe a benefit of Oracle's multicloud strategy for Oracle AI Database@Azure?
A. Eliminates the requirement to maintain any Oracle Cloud Infrastructure tenancy.

B. A managed network connection between the Azure data center and the nearest OCI region.
C. Colocation of OCI infrastructure within Microsoft's Azure data centers.
D. Direct, low-latency connectivity between Oracle Exadata and Azure-based applications.
Answer: (SHOW ANSWER)
Oracle AI Database@Azure does not eliminate the need for an OCI tenancy. Oracle explicitly requires an Oracle Cloud account/OCI tenancy during onboarding. A Public Offer creates a new tenancy, while a Private Offer can link an eligible existing tenancy or create one.
The other choices describe genuine architectural advantages. Oracle deploys its database infrastructure inside isolated areas of selected Microsoft Azure data centers, which places Oracle database resources physically close to Azure applications. This enables low-latency application-to-database communication through Azure networking rather than forcing production data traffic across the public internet to a distant OCI region.
Oracle also configures a managed connection between the Azure Child Site and geographically paired OCI Parent Site. The Parent Site hosts the Oracle control-plane components used to operate and manage the Oracle infrastructure located inside Azure.
The result is an architecture combining Azure-native application proximity and management integration with Oracle's database infrastructure and operational control.
Therefore, the claim that OCI tenancy requirements disappear is false, making A the correct "NOT a benefit" answer.
Study Guide reference: Introduction to Multicloud - Oracle AI Database@Azure architecture and multicloud benefits.

NEW QUESTION: 7
You are provisioning a new guest Exadata VM Cluster inside the Azure Portal and must specify a unique identifier for the host name prefix. What is the maximum number of characters allowed for this host name prefix under Oracle AI Database@Azure network design rules?
A. 30 characters
B. 15 characters
C. 63 characters
D. 23 characters
Answer: D (LEAVE A REPLY)
Oracle AI Database@Azure applies stricter hostname-prefix validation than the generic 63-character DNS- label maximum. For an Exadata VM Cluster, Oracle's DNS and network-design documentation specifies a maximum hostname-prefix length of 23 characters .
The value also has syntactic constraints. It must begin with a letter, end with an alphanumeric character, and can contain hyphens. Certain reserved values are prohibited. These restrictions exist because the supplied prefix becomes part of system-generated host and database-network names; Oracle must reserve additional character capacity for generated suffixes and service-specific naming constructs.

Therefore, 63 characters represents a general DNS-label ceiling but is not valid for this Exadata field. Thirty characters likewise exceeds the Oracle-specific limit, while 15 characters is unnecessarily restrictive.
The correct maximum is consequently 23 characters .
This is a configuration detail worth memorizing because hostname validation occurs during VM Cluster provisioning and an invalid prefix can block resource creation before the underlying Exadata compute nodes are deployed.
Study Guide reference: Configure Oracle AI Database@Azure Networking - Exadata VM Cluster DNS, hostname prefix rules and naming constraints.

NEW QUESTION: 8
Which service is Oracle's native, recommended tool to orchestrate and automate database-level role transitions (switchover and failover) across regions for Oracle Database@Azure?
A. Azure Site Recovery
B. OCI Full Stack Disaster Recovery
C. OCI GoldenGate Service Console
D. Azure Traffic Manager
Answer: (SHOW ANSWER)
Oracle recommends OCI Full Stack Disaster Recovery (Full Stack DR) for orchestrating disaster-recovery operations around Oracle Database@Azure deployments.
Oracle's reference architecture for local and regional Oracle Database@Azure standbys explicitly recommends Full Stack DR to orchestrate database switchover and failover operations . The service can coordinate ordered disaster-recovery plans instead of requiring administrators to manually perform every database and infrastructure transition during an incident or DR drill.
Oracle Data Guard remains the underlying database replication and role-management technology, while Full Stack DR provides the higher-level orchestration framework for executing coordinated transitions.
Azure Site Recovery is primarily designed for Azure VM and workload replication/recovery and is not Oracle's recommended database-role orchestration service. GoldenGate provides logical replication and data integration rather than end-to-end DR orchestration. Azure Traffic Manager performs DNS-based traffic distribution and can participate in application routing, but it does not execute Oracle database switchovers or failovers.
Therefore, B is correct.
Study Guide reference: HA and DR with Oracle AI Database@Azure - Full Stack DR orchestration and Data Guard role transitions.

NEW QUESTION: 9
Which OCI Identity and Access Management (IAM) policy must be added to your OCI tenancy to prevent
404 NotAuthorizedOrNotFound errors when default OCI tags are automatically generated during Exadata VM cluster creation in Oracle AI Database@Azure?

A. Allow group administrators to manage tagging-family in tenancy
B. Allow any-user to inspect tag-namespaces in tenancy where request.user.id = 'multicloud'
C. Allow any-user to use tag-namespaces in tenancy where request.principal.type = 'multicloudlink'
D. Allow group odbaa-network-administrators to use tag-namespaces in compartment
Answer: (SHOW ANSWER)
Oracle documents a specific 404 NotAuthorizedOrNotFound condition that can occur during Exadata VM Cluster creation when the selected private DNS zone does not already contain tags. Oracle then attempts to generate the default OCI oracle-tags, but the multicloud service principal must have permission to use the associated tag namespace.
The required policy statement is:
Allow any user to use tag-namespaces in tenancy where request.principal.type = 'multicloudlink'
Oracle's current troubleshooting guidance also requires the companion permission:
Allow any user to manage tag-defaults in tenancy where request.principal.type = 'multicloudlink'
The crucial detail is that the provisioning operation executes under the multicloudlink principal type , not an ordinary administrator group or a principal with request.user.id='multicloud'. Merely granting inspect permission is insufficient because the service must actually apply the tag namespace/defaults.
Thus, among the supplied options, C is the required policy statement.
Study Guide reference: Secure Oracle AI Database@Azure - OCI IAM, MulticloudLink service principal and tagging permissions.

NEW QUESTION: 10
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?
Freecram.net
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: (SHOW ANSWER)
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.

NEW QUESTION: 11
Which capability is provided by a Network Anchor in Oracle AI Database@Azure?
A. Encrypt data in transit between Azure and OCI using IPsec exclusively.
B. Provide DNS forwarding and listening endpoints between Azure and OCI.
C. Serve as the customer-managed key store for TDE.
D. Automatically create Availability Domains inside Azure regions.
Answer: (SHOW ANSWER)
A Network Anchor provides the cross-cloud network integration context used by supported Oracle AI Database@Azure services, including important DNS-resolution capabilities between Azure and OCI .
For inbound DNS resolution toward Azure, Oracle's Network Control Plane can create an OCI forwarding DNS endpoint and associated forwarding rules. For Azure-to-OCI resolution, the service can create an OCI listening DNS endpoint and expose its address through the Network Anchor so Azure DNS infrastructure or on-premises DNS servers can forward queries for Oracle database namespaces.
When both directions are configured, the architecture provides bidirectional private DNS resolution between Azure and OCI environments. Network Anchors can also participate in private-zone synchronization depending on the database service and DNS configuration.
The Network Anchor is not a cryptographic key store; TDE keys belong in Oracle Wallet, OCI Vault, Oracle Key Vault, or Azure Key Vault. Nor does the Network Anchor create Azure Availability Zones/OCI Availability Domains or mandate an IPsec-only intercloud transport mechanism.
Therefore, B is correct.
Study Guide reference: Configure Oracle AI Database@Azure Networking - Network Anchors, DNS listening endpoints and forwarding endpoints.

NEW QUESTION: 12
You have enabled automatic backups on an Autonomous AI Database with a 30-day retention period. You subsequently terminate the Autonomous Container Database. What happens to the previously retained automatic backups?
A. They are deleted immediately upon termination.

B. They are retained for up to 72 hours or until the end of their retention period, whichever retention option applies.
C. They are automatically converted into long-term backups.
D. They are exported automatically to Azure Blob Storage.
Answer: (SHOW ANSWER)
Oracle Autonomous Container Database backup configuration includes an explicit backup-retention policy on termination . Administrators can select one of two supported behaviors: retain backups according to the normal configured backup-retention period, or retain them for 72 hours after termination and then delete them .
Oracle documentation identifies these options as RETAIN_PER_RETENTION_WINDOW and RETAIN_FOR_72_HOURS. The latter is the default in relevant API/Terraform configuration.
Therefore, terminating an Autonomous Container Database does not inherently result in the immediate deletion of all automatic backups. Their post-termination lifetime is governed by the configured termination-retention policy.
Standard automatic backups are also not automatically converted into Long-Term Retention backups when the database is terminated. LTR backups have their own lifecycle and are explicitly excluded from the ordinary ACD termination-retention setting. No automatic export to Azure Blob Storage occurs as part of ACD termination.
Thus, B correctly reflects Oracle's documented post-termination backup behavior.
Study Guide reference: Secure Oracle AI Database@Azure - Autonomous Container Database automatic backups, termination and retention policy.

NEW QUESTION: 13
Which two are officially documented multicloud use cases according to the Oracle Cloud Infrastructure (OCI) Cloud Adoption Framework? (Choose two.)
A. Automatic Cross-Cloud Identity Domain Replication
B. Running Applications in One Cloud and Databases in Another
C. Cross-Cloud Disaster Recovery
D. Federated Active Directory Domain Controller Synchronization
Answer: (SHOW ANSWER)
Oracle's multicloud guidance identifies business and architecture patterns rather than narrowly defined infrastructure features as core multicloud use cases. Two explicitly documented patterns are Cross-Cloud Disaster Recovery and Running Applications in One Cloud and Databases in Another .
Cross-cloud disaster recovery places a production environment in one cloud while maintaining a recovery environment in another. This design reduces dependence on a single cloud platform and enables business continuity if the primary environment experiences a major outage. It is particularly relevant where resilience requirements justify provider-level diversification.
The application/database split pattern allows an organization to operate its application tier in one cloud while placing its database in another cloud that provides preferable database capabilities,

performance, economics, or management characteristics. Oracle AI Database@Azure is a practical implementation of this model:
Azure applications can use Oracle database technology deployed in Azure data centers and managed through Oracle's multicloud architecture.
Automatic identity-domain replication and Active Directory domain-controller synchronization may support a multicloud environment, but they are implementation capabilities rather than the high-level multicloud use cases documented by Oracle.
Study Guide reference: Introduction to Multicloud - multicloud drivers, application/database separation, cross-cloud DR, resiliency and workload placement.

NEW QUESTION: 14
Your team is migrating a processing-intensive workload to Oracle Autonomous AI Database Serverless on Oracle AI Database@Azure and needs to scale up the compute allocation during provisioning. What is the maximum number of ECPUs that can be configured for a single database instance using the slider user interface?
A. 512 ECPUs
B. 1024 ECPUs
C. 256 ECPUs
D. 128 ECPUs
Answer: (SHOW ANSWER)
The standard Oracle Autonomous AI Database Serverless resource-allocation interface permits the administrator to configure the base compute allocation up to 512 ECPUs using the ECPU slider. This value represents the configured base ECPU count for a single database instance rather than the potential temporary compute consumption obtainable through compute auto scaling.
Oracle distinguishes configured ECPUs from auto-scaled usage. For eligible Autonomous AI Database configurations, compute auto scaling can allow the database to consume additional processing capacity above the configured base according to workload demand. For example, a database configured with 512 ECPUs may be capable of consuming substantially more effective processing capacity when auto scaling is enabled, subject to the service and licensing model.
The 1024-ECPU option therefore does not represent the normal single-instance slider ceiling addressed by this question. The 256- and 128-ECPU values are valid intermediate configurations but are below the maximum.
Consequently, A. 512 ECPUs is correct.
Study Guide reference: Provision and operate Oracle AI Database@Azure Resources - Autonomous AI Database Serverless provisioning, ECPU configuration and compute scaling.

NEW QUESTION: 15
Which two purchase options are available for Oracle Exadata Database Service on Exascale Infrastructure in Oracle AI Database@Azure? (Choose two.)
A. Azure Reserved Instance pricing

B. Private Offer
C. Pay As You Go
D. Free Trial credits
Answer: (SHOW ANSWER)
Oracle Exadata Database Service on Exascale Infrastructure is eligible for both the Private Offer and Public Offer , with the Public Offer representing the Pay As You Go marketplace model.
A Private Offer is negotiated with Oracle and can include custom pricing, contract terms, usage commitments, and enterprise commercial arrangements. After agreement, the customer accepts the private offer through Azure Marketplace.
A Public Offer allows supported Oracle AI Database@Azure services to be purchased directly from Azure Marketplace without first negotiating a private commercial agreement. Oracle describes this model as Pay As You Go, with no predetermined usage commitment. Oracle's current Azure service table explicitly lists Exadata Database Service on Exascale Infrastructure as eligible for both offer types.
Azure Reserved Instance pricing is an Azure compute purchasing construct and is not one of the Oracle AI Database@Azure offer mechanisms. Oracle also specifies that free-trial subscription models are not the purchasing path for these multicloud database resources.
Therefore, B and C are correct.
Study Guide reference: Oracle AI Database@Azure Architecture and Onboarding - Public Offer, Private Offer and Exascale purchasing.

NEW QUESTION: 16
Your team is configuring a ZDM Physical Offline migration to Exadata on Oracle Database@Azure, but the target database host cannot read the RMAN backup files mounted on the Azure Files NFS share. What must you verify to resolve this storage access and permission issue?
A. The Azure Files private endpoint is configured to use Active Directory Domain Services authentication.
B. The Unique Identifier (UID) of the database owner matches on both the source and target hosts.
C. The subnet delegated to Oracle.Database/networkAttachments permits port 443 outbound.
D. The Oracle Net listener on the target Exadata host is running on the default port 1521.
Answer: (SHOW ANSWER)
ZDM Physical Offline migration can use an Azure Files NFS share as the common location for RMAN backup files. Because the NFS share is mounted on both the source and target database hosts, file ownership must be interpreted consistently on both systems.
NFS using Unix-style permissions identifies users primarily by their numeric UID rather than simply by the textual account name. Oracle's ZDM migration guidance explicitly states that the NFS-mounted path is readable by the target database user only when the source database user's UID matches the target database user's UID. Thus, two hosts can both have an account named

oracle, but if those accounts have different numeric UIDs, files created by one host can appear to be owned by a different identity on the other.
Port 443 and SQL*Net listener port 1521 address network/API or database-client connectivity; neither resolves Unix file ownership on an NFS mount. Active Directory Domain Services authentication is likewise not the controlling permission model for the documented NFS migration workflow.
Study Guide reference: Migrate Oracle Databases to Oracle AI Database@Azure - ZDM Physical Offline migration, Azure Files NFS, RMAN backup staging and Linux UID permissions.

Valid 1z0-1147 Dumps shared by EduDump.com for Helping Passing 1z0-1147 Exam! EduDump.com now offer the **newest 1z0-1147 exam dumps**, the EduDump.com 1z0-1147 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com 1z0-1147 dumps with Test Engine here: https://www.edudump.com/exams/Oracle/1z0-1147/premium/ (**103** Q&As Dumps, **35%OFF** **Special Discount Code: freecram**)

NEW QUESTION: 17
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: (SHOW ANSWER)
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.

NEW QUESTION: 18
Your organization has deployed several Exadata VM Clusters within Oracle AI Database@Azure, and a database administrator needs to create custom dashboards and adjust metric collection settings in the Azure Portal. Which built-in Azure role must be assigned to this administrator to enable both viewing metrics and modifying monitoring configurations?
A. Reader
B. Owner
C. Monitoring Reader
D. Monitoring Contributor
Answer: (SHOW ANSWER)
The appropriate built-in Azure RBAC role is Monitoring Contributor . Microsoft defines this role for administrators who need to read monitoring information and create or modify monitoring configuration .
Monitoring Contributor includes the monitoring-read capabilities required to inspect metrics and diagnostic data and adds permissions needed to create or edit diagnostic settings, alert-related configurations, and monitoring resources. Microsoft also identifies creation of private monitoring dashboards as a supported use case for this role.
Monitoring Reader is insufficient because it is primarily a read-only monitoring role. It allows users to inspect monitoring data and settings but does not grant the required configuration-modification permissions. The general Reader role similarly lacks write capabilities. Owner would technically provide sufficient access but violates least privilege because it grants comprehensive resource-management permissions, including access- control authority far beyond monitoring.
Because the administrator's duties are specifically monitoring-oriented but require configuration changes, Monitoring Contributor supplies the correct functional scope without granting unrestricted ownership.
Study Guide reference: Secure Oracle AI Database@Azure - Azure RBAC, monitoring permissions, least privilege and operational administration

NEW QUESTION: 19
Which two Azure network components are required to construct a secure Hub-and-Spoke topology with a central firewall for Oracle AI Database@Azure? (Choose two.)
A. Azure Virtual Machine Scale Sets with local storage mirroring
B. Spoke VNet containing the delegated subnet for Oracle AI Database@Azure
C. Hub VNet hosting an Azure Firewall or third-party Network Virtual Appliance (NVA)

D. OCI FastConnect Shared Virtual Circuit peering

Answer: (SHOW ANSWER)

A Hub-and-Spoke design separates shared connectivity and security services from workload networks. For Oracle AI Database@Azure, Oracle documents a Hub VNet that hosts centralized services such as Azure Firewall or a third-party NVA and a dedicated spoke VNet for Oracle AI Database@Azure containing the delegated subnet .

The Oracle database spoke is connected to the hub through Azure VNet peering. User-Defined Routes are then used to direct traffic through the hub's firewall or NVA when centralized inspection is required.

Application spokes can similarly route database-bound traffic toward the security appliance, while return routes on the database spoke preserve the required symmetric traffic path.

Azure VM Scale Sets are application-compute constructs and do not form the essential Hub-and-Spoke network topology. OCI FastConnect is an OCI/hybrid connectivity capability rather than the Azure-side Hub- and-Spoke component described by this scenario.

Therefore, B and C identify the essential architecture: the Oracle database workload is isolated in a spoke VNet, while the Hub VNet provides centralized traffic inspection and security enforcement.

Study Guide reference: Configure Oracle AI Database@Azure Networking - Hub-and-Spoke architecture, VNet peering, delegated subnet, Azure Firewall/NVA and UDR routing.

NEW QUESTION: 20

When monitoring metrics for Oracle Exadata Database Service on Dedicated Infrastructure within Oracle AI Database@Azure, which specific metric label must you select in the Azure Portal to view the performance of individual Pluggable Databases (PDBs)?

Freecram.net

A. CPU Utilization (Oracle OCI Database)

B. CPU Utilization (OCI Database Cluster)

C. CPU Utilization (OCI VM Host)

D. CPU Utilization (OCI Database)

Answer: A (LEAVE A REPLY)

Oracle AI Database@Azure exposes database and infrastructure metrics through Azure Monitor, but similarly named metrics correspond to different levels of the Oracle resource hierarchy. Correct interpretation of the label is therefore essential.

Oracle's monitoring documentation explicitly differentiates the namespaces. OCI Database Cluster represents VM-cluster-level measurements. OCI Database represents Container Database (CDB) metrics.

Oracle OCI Database identifies Pluggable Database-level metrics. Consequently, to examine CPU utilization for an individual PDB, the administrator selects CPU Utilization (Oracle OCI Database) .

This distinction is particularly important for consolidated Exadata environments. A cluster-level CPU metric can show overall infrastructure utilization but cannot determine which PDB is responsible for a workload spike. Similarly, CDB-level utilization aggregates activity at the

container-database layer. The Oracle OCI Database-labelled metric enables the administrator to work at the PDB scope needed for tenant or application- specific performance analysis.
The Azure Monitor integration therefore provides several metrics with superficially similar names, but their resource dimensions are deliberately different.
Study Guide reference: Provision and operate Oracle AI Database@Azure Resources - Azure Monitor integration, Exadata monitoring, VM Cluster/CDB/PDB metric scopes and operational observability.

NEW QUESTION: 21
Which of the following is NOT one of the documented supported network topologies for Oracle AI Database@Azure?
A. Local VNet peering
B. Hub-and-spoke
C. Azure Front Door single-endpoint routing
D. Azure Virtual WAN
Answer: (SHOW ANSWER)
Azure Front Door single-endpoint routing is not a documented Oracle AI Database@Azure database connectivity topology.
Azure Front Door is primarily a global Layer-7 HTTP/HTTPS application-delivery, acceleration, and load- balancing service. It is designed for web applications and public-facing application endpoints rather than private SQL*Net/database connectivity into an Oracle AI Database@Azure delegated subnet.
Oracle's documented networking patterns instead center on Azure VNets and private routing. These include local/same-region VNet connectivity, VNet peering, Hub-and-Spoke VNet peering, cross-region connectivity, and hybrid on-premises Hub-and-Spoke architectures. Oracle also documents enterprise connectivity patterns involving Azure Virtual WAN.
Hub-and-Spoke is particularly important where centralized firewalls or NVAs must inspect database-bound traffic, while local VNet/VNet peering is useful for direct application-to-database private communication where minimizing network latency is the priority.
Because Azure Front Door does not provide the required private database networking architecture and is absent from Oracle's supported topology set, C is the correct answer.
Study Guide reference: Configure Oracle AI Database@Azure Networking - supported VNet, peering, Hub-and-Spoke and enterprise connectivity topologies.

NEW QUESTION: 22
Which two Microsoft Azure resource providers must be registered in your Azure subscription as prerequisites to both onboard with and provision resources in Oracle AI Database@Azure? (Choose two.)
A. Microsoft.VirtualMachine
B. Microsoft.BareMetal
C. Oracle.Exadata

D. Oracle.Database

Answer: (SHOW ANSWER)

Oracle AI Database@Azure depends on Azure Resource Manager provider namespaces that expose the Oracle service and its underlying infrastructure integration. Among the listed choices, the required providers are Microsoft.BareMetal and Oracle.Database .

Oracle.Database is the Azure resource-provider namespace through which Oracle database multicloud resources are represented and managed in Azure. This includes Oracle AI Database@Azure database-resource lifecycle operations and integration with delegated networking. Microsoft.BareMetal supports the Azure-side integration with specialized Oracle infrastructure physically deployed within Azure facilities.

Current Microsoft onboarding documentation also lists Microsoft.Network and Microsoft.Compute among resource providers required for a fully provisionable subscription. However, those names are not offered as valid alternatives in this question. Microsoft.VirtualMachine is not the correct provider namespace-standard Azure compute uses Microsoft.Compute. Likewise, Oracle.Exadata is not the service namespace; Oracle AI Database@Azure resources are represented under Oracle.Database.

Therefore, among the supplied answer set, B and D are correct.

Study Guide reference: Oracle AI Database@Azure Architecture and Onboarding - Azure subscription prerequisites, resource-provider registration and Oracle.Database integration.

NEW QUESTION: 23

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: (SHOW ANSWER)

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.

NEW QUESTION: 24
Your enterprise is configuring federated single sign-on (SSO) and authentication for administrators operating Oracle AI Database@Azure resources. Which identity provider natively provides federated identity and access management for this co-located database service?
A. Microsoft Entra ID
B. OCI Identity Domains
C. Active Directory Federation Services
D. Okta Identity Cloud
Answer: (SHOW ANSWER)
For Oracle AI Database@Azure's documented Azure-to-OCI federation architecture, Microsoft Entra ID functions as the enterprise identity provider. Oracle's reference architecture explicitly states that Microsoft Entra ID provides federated identity and access management for Oracle AI Database@Azure.
The OCI IAM Identity Domain remains an essential component because it is the Oracle-side service provider and identity domain that trusts the Microsoft Entra identity provider. However, when the question asks which platform provides the federated corporate identity in the Azure-integrated architecture, the answer is Microsoft Entra ID.
Oracle's onboarding instructions configure an Oracle Cloud Infrastructure Console enterprise application in Entra ID and then configure the OCI identity domain with Microsoft Entra ID as its SAML identity provider.
Administrators can subsequently use their corporate Entra credentials for SSO into OCI when Oracle-native management tasks require it.
AD FS represents an older on-premises Microsoft federation service, while Okta is a third-party identity platform rather than the natively documented Azure integration.
Therefore, A is correct.
Study Guide reference: Secure Oracle AI Database@Azure - Entra ID federation, OCI IAM Identity Domains and SSO.

NEW QUESTION: 25
Which two DNS endpoints can be optionally enabled when provisioning a Network Anchor for Oracle AI Database@Azure? (Choose two.)
A. DNS forwarding endpoint
B. DNS replication endpoint
C. DNS peering endpoint

D. DNS listening endpoint

Answer: (SHOW ANSWER)

Oracle AI Database@Azure Network Anchors support DNS integration between Azure-side applications and OCI-side database networking. The two endpoint constructs used by Oracle DNS are the DNS forwarding endpoint and DNS listening endpoint .

For inbound DNS resolution from OCI toward Azure, Oracle's Network Control Plane can create a forwarding DNS endpoint and configure forwarding rules so OCI-hosted resources can forward queries toward the Azure DNS environment. Conversely, outbound resolution from Azure toward Oracle resources relies on an OCI listening DNS endpoint. The listening endpoint receives DNS requests and enables Azure or connected on- premises systems to resolve Oracle AI Database@Azure names.

The terminology is important. Oracle explicitly uses listening and forwarding endpoint constructs. "DNS replication endpoint" is not the endpoint name used for private-zone synchronization, and "DNS peering endpoint" is not an Oracle Network Anchor endpoint type. DNS-record synchronization is a separate behavior from the endpoint resource itself.

Consequently, the correct selections are A and D .

Study Guide reference: Configure Oracle AI Database@Azure Networking - Network Anchors, inbound DNS, outbound DNS, forwarding endpoints and listening endpoints.

NEW QUESTION: 26

You are configuring a Network Anchor in the Azure Portal and want to enable the "Replicate DNS Private zones during DB creation" option to automatically copy private zones from OCI to Azure. Which database service must be associated with the Network Anchor to make this specific DNS option available?

A. Oracle Exadata Database Service on Exascale Infrastructure

B. Oracle Exadata Database Service on Dedicated Infrastructure

C. Oracle Base Database Service

D. Oracle Autonomous AI Database

Answer: (SHOW ANSWER)

The Network Anchor abstraction establishes the mapping between Azure networking resources and corresponding OCI networking constructs. It is particularly important for services that use the newer multicloud anchor model, allowing an Azure VNet and delegated subnet to be associated consistently with Oracle-managed resources.

Oracle's DNS documentation explicitly restricts the Network Anchor DNS-resolution capability in question to Oracle Base Database Service . Network Anchors can provide DNS integration between Azure and OCI, including inbound and outbound resolution paths. Depending on the configuration, DNS listening and forwarding functionality can be enabled so resources on either side of the multicloud boundary can resolve the required private names.

The question's "Replicate DNS Private zones during DB creation" behavior belongs to this Base Database Service Network Anchor DNS workflow. Exadata Database Service on Dedicated Infrastructure and Exadata Database Service on Exascale Infrastructure use their documented

Exadata networking and DNS provisioning mechanisms. Autonomous AI Database likewise has its own networking lifecycle and is not the service to which this Network Anchor-specific private-zone replication setting applies.
Therefore, the database service required to expose the specified option is Oracle Base Database Service .
Study Guide reference: Configure Oracle AI Database@Azure Networking - Network Anchors, private DNS, inbound/outbound DNS resolution, Base Database Service networking.

NEW QUESTION: 27
Which network port must be open and tested to ensure SQL*Net connectivity between the source database and the ZDM Service Host during a ZDM Physical Online migration to Exadata on Oracle Database@Azure?
A. 2049
B. 443
C. 22
D. 1521
Answer: (SHOW ANSWER)
The standard Oracle listener and SQL Net port used in the documented ZDM Physical Online migration workflow is TCP 1521 . Oracle's Oracle AI Database@Azure ZDM migration procedure explicitly instructs administrators to test SQL Net connectivity by running tnsping between source and target database hosts using port 1521.
This connectivity is required because ZDM and its underlying migration technologies must establish database- level communication for prechecks, configuration, target instantiation, Data Guard configuration, synchronization, and migration orchestration. If the source or target listener has been intentionally configured to use a nondefault port, that configured listener port must be permitted instead.
Port 22 is used for SSH administration between the ZDM Service Host and managed hosts. Port 443 is associated with HTTPS-based cloud service and API communication. Port 2049 is conventionally associated with NFS and is relevant when migration workflows use NFS-based staging, but it is not SQL*Net.
Therefore, D. 1521 is the correct answer.
Study Guide reference: Migrate Oracle Databases to Oracle AI Database@Azure - ZDM network prerequisites, SQL*Net and physical online migration.

NEW QUESTION: 28
Which Recovery Manager (RMAN) method is recommended by Oracle Maximum Availability Architecture (MAA) and used by Zero Downtime Migration (ZDM) when direct data transfer is enabled?
A. Restore From Service
B. Compressed Backupset Streaming
C. Image Copy Incremental Push

D. Active Database Duplication

Answer: (SHOW ANSWER)

For ZDM physical online migrations using direct data transfer , Oracle's Maximum Availability Architecture pattern uses RMAN's Restore From Service capability. Rather than first creating backup pieces in a staging location, transporting them, and then restoring them on the target, RMAN restores the target database directly across the network from the source service.

This design reduces intermediate storage requirements and simplifies the migration pipeline. Once the target database has been instantiated, ZDM automatically configures Oracle Data Guard and uses redo transport to maintain synchronization until the planned switchover. The combination of Restore From Service and Data Guard is the central technical pattern Oracle documents for physical online migration to Exadata Database Service on Oracle AI Database@Azure.

Although DUPLICATE ... FROM ACTIVE DATABASE is a legitimate RMAN cloning capability in other contexts, it is not the named method documented for this ZDM direct-data-transfer workflow. Compressed backup-set streaming and incremental image-copy pushing also imply different backup/staging mechanisms.

Therefore, Restore From Service is the MAA-aligned RMAN method represented by the scenario.

Study Guide reference: Migrate Oracle Databases to Oracle AI Database@Azure - MAA, ZDM Physical Online migration, RMAN direct data transfer and Restore From Service.

NEW QUESTION: 29

You have selected the custom FQDN option and specified your own domain during Exadata VM Cluster provisioning in Oracle AI Database@Azure. What must you do afterward to ensure name resolution works from your production DNS zone?

A. Disable the OCI private view associated with the VM cluster.

B. Wait for Oracle to automatically create the oraclevcn.com zone in Azure.

C. Manually duplicate the SCAN and VIP DNS records into your production DNS zone.

D. Reconfigure the VM cluster to use Default DNS instead.

Answer: (SHOW ANSWER)

When an architect selects Custom DNS/FQDN , the database cluster is created using the organization's chosen DNS namespace instead of relying exclusively on Oracle's default automatically managed oraclevcn.

com private-zone behavior.

The important operational requirement is that application clients must be able to resolve the Exadata cluster's connection endpoints. That includes the Single Client Access Name (SCAN) and the relevant Virtual IP (VIP) records. Because the organization controls the custom production DNS namespace, those records must be represented in that DNS environment so client applications can resolve the cluster correctly.

Disabling OCI private DNS would damage internal Oracle name resolution rather than extending resolution to the enterprise DNS system. Waiting for Oracle to create a default oraclevcn.com

zone does not satisfy a custom-domain deployment, and reverting to default DNS would defeat the original custom-FQDN design.
Therefore, the required action is to manually duplicate or publish the SCAN and VIP records in the production DNS zone .
Study Guide reference: Configure Oracle AI Database@Azure Networking - custom FQDN, SCAN
/VIP records and enterprise DNS integration.

NEW QUESTION: 30

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?

Freecram.net

A. OCI Object Storage
B. OCI File Storage Service
C. Azure Files NFS share
D. Azure Blob Storage API
Answer: C (LEAVE A REPLY)
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.

NEW QUESTION: 31
A mid-size retailer wants to move a moderately sized seasonal transactional workload to Oracle AI Database@Azure. They want to avoid provisioning dedicated database and storage servers, want to start with a small committed footprint and elastically scale ECPUs and storage online as holiday sales volume grows, and want the ability to create instant, space-efficient thin clones of

production for their developers. As the architect, you must select the Oracle AI Database@Azure deployment option that meets all of these needs.
Choose ONE.
A. Oracle Autonomous AI Database Serverless
B. Oracle Exadata Database Service on Exascale Infrastructure
C. Oracle Exadata Database Service on Dedicated Infrastructure
D. Oracle Base Database Service
Answer: B (LEAVE A REPLY)
Oracle Exadata Database Service on Exascale Infrastructure precisely matches the scenario. Exascale uses Oracle-managed, multitenant physical database-server infrastructure, allowing customers to provision relatively small VM Clusters rather than committing to dedicated Exadata database and storage servers.
A VM Cluster can begin at 8 ECPUs per VM and scale its compute allocation in 4-ECPU increments.
Exascale storage is also elastic and can be expanded without requiring customers to add or manage physical Exadata Storage Servers.
A distinctive Exascale feature is redirect-on-write thin cloning . Oracle documents the ability to create space-efficient PDB clones in which unchanged blocks remain shared between the parent and clone. This makes development and test copies fast to provision and economical in storage consumption.
Dedicated Infrastructure contradicts the requirement to avoid dedicated server capacity. Base Database Service does not offer the same Exascale thin-clone architecture, while Autonomous Serverless represents a different autonomous operating model.
Therefore, B is correct.
Study Guide reference: Provision and operate Oracle AI Database@Azure Resources - Exascale elasticity and thin clones.

Valid 1z0-1147 Dumps shared by EduDump.com for Helping Passing 1z0-1147 Exam! EduDump.com now offer the **newest 1z0-1147 exam dumps**, the EduDump.com 1z0-1147 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com 1z0-1147 dumps with Test Engine here: https://www.edudump.com/exams/Oracle/1z0-1147/premium/ (**103** Q&As Dumps, **35%OFF** **Special Discount Code: freecram**)

NEW QUESTION: 32
To achieve continuous availability and automated database failover using Data Guard Fast-Start Failover (FSFO) across availability zones in Oracle AI Database@Azure, where is it recommended to host the client- side Observer process?
A. On one of the primary Exadata VM cluster database nodes

B. On an OCI utility server deployed in the OCI parent region control plane
C. On a separate Azure virtual machine deployed in a different availability zone
D. On one of the standby Exadata VM cluster database nodes
Answer: C (LEAVE A REPLY)

NEW QUESTION: 33
Your database administrator enabled a backup retention lock on an Autonomous AI Database in Oracle AI Database@Azure but now needs to reduce the retention period. What must you do to disable this retention lock or alter the retention period?
A. File a Service Request at Oracle Cloud Support
B. Modify the backup retention slider in the Azure Portal
C. Run an ALTER DATABASE system command in the OCI Cloud Shell
D. Update the key rotation settings in the Azure Key Vault
Answer: (SHOW ANSWER)
Backup Retention Lock is a protection mechanism intended to prevent administrators or compromised credentials from prematurely reducing the recoverability window of Autonomous AI Database backups. Once the lock is established, normal administrative interfaces cannot simply disable it or reduce the protected backup-retention period.
Oracle's Autonomous AI Database documentation explicitly states that after the backup retention lock has been set, the option cannot be disabled or the retention period changed through ordinary database administration. To disable the lock or modify the locked retention setting, the administrator must file a Service Request with Oracle Cloud Support .
This design provides an administrative separation between ordinary database management and an action that could materially weaken backup protection. Changing an Azure Key Vault key has no relationship to the retention-policy lock. An SQL ALTER DATABASE command cannot override a service-level backup governance control, and the Azure Portal does not expose a supported bypass to reduce the locked period.
Accordingly, A is the only supported escalation path in the choices.
Study Guide reference: Secure Oracle AI Database@Azure - Autonomous AI Database backup protection, retention policy, retention lock and operational security controls.

NEW QUESTION: 34
An administrator is preparing to migrate an on-premises database to Oracle Database@Azure and runs the zdmcli migrate database command with the -eval option. What is the outcome of executing this command with this specific option?
A. It schedules a job of type EVAL to execute only migration prechecks and validation phases.
B. It starts a physical offline migration that pauses automatically before copying the database files.
C. It generates a mock standby database and performs a dry-run Data Guard switchover.
D. It downloads the Cloud Premigration Advisor Tool (CPAT) and installs it on the target Exadata host.

Answer: (SHOW ANSWER)

The -eval option invokes the evaluation workflow of Oracle Zero Downtime Migration rather than executing the production migration. Oracle specifically directs administrators to use ZDMCLI migrate database ... -eval to evaluate the migration process, including prerequisite checks, source/target connectivity verification, configuration validation, and applicable pre-migration assessment phases.

The resulting job is an evaluation job rather than a migration job that performs database movement. This distinction is reflected throughout ZDM command handling: ZDM can query evaluation jobs separately from actual migration jobs, and the evaluation workflow is intended to identify prerequisite or configuration failures before production data movement begins.

The option does not create a mock standby or execute a simulated Data Guard switchover. It also does not begin an offline migration and stop before the copy phase. CPAT may participate in migration-readiness assessment depending on ZDM version and command options, but -eval is not an installation command for CPAT.

Consequently, A correctly represents the effect: an EVAL-type workflow performs migration prechecks and validation without executing the actual migration.

Study Guide reference: Migrate Oracle Databases to Oracle AI Database@Azure - ZDM CLI, evaluation mode, migration validation and prechecks.

NEW QUESTION: 35

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: (SHOW ANSWER)

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.

NEW QUESTION: 36
Your organization operates a 20 TB on-premises Oracle Database 19c on Oracle ASM storage that must be migrated to Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D) on Oracle AI Database@Azure. The application owner requires that database and application downtime be limited to only the time needed to apply the final increment of redo before switchover. The source and target must remain on the exact same Oracle Database version and character set, and the migration must use RMAN-based backup and restore rather than a logical export/import. As the migration architect, you must select the ZDM workflow that satisfies all of these requirements. Choose ONE.
A. ZDM Logical Offline Migration
B. ZDM Physical Online Migration
C. ZDM Logical Online Migration
D. ZDM Physical Offline Migration
Answer: B (LEAVE A REPLY)
The requirements map directly to ZDM Physical Online Migration . A physical workflow uses RMAN- based database instantiation rather than Data Pump logical export/import. For online migration, ZDM combines physical target creation with Oracle Data Guard synchronization so redo generated after the initial database copy continues to be transported and applied to the target.
Because physical migration preserves the database at the block/file level, source and target compatibility requirements-including database version, platform characteristics, and character-set considerations-are more restrictive than with logical migration. The scenario explicitly states that the version and character set remain identical, which makes the physical workflow appropriate.
The online variant minimizes downtime because users continue operating against the source while target synchronization proceeds. Application interruption is concentrated around the final synchronization and switchover.
Logical Online instead uses Data Pump and GoldenGate. Physical Offline satisfies the RMAN requirement but introduces substantially longer downtime because the source cannot remain operational throughout migration.
Therefore, B is correct.
Study Guide reference: Migrate Oracle Databases to Oracle AI Database@Azure - ZDM Physical Online, RMAN and Data Guard.

NEW QUESTION: 37
What is the allowed range of days for the automatic backup retention period when configuring an instance of Oracle Autonomous AI Database Serverless in Oracle AI Database@Azure?
A. 1 to 60 days

B. 7 to 60 days
C. 1 to 30 days
D. 7 to 30 days
Answer: (SHOW ANSWER)

NEW QUESTION: 38
For Oracle Autonomous AI Database@Azure cross-region disaster recovery using Autonomous Data Guard, what is the approximate recovery point objective (RPO)?
A. Zero data loss (fully synchronous)
B. 5 minutes
C. 15 minutes
D. 1 minute
Answer: (SHOW ANSWER)
For a cross-region Autonomous Data Guard standby , Oracle specifies a maximum possible data-loss objective of approximately 1 minute . Accordingly, the intended Recovery Point Objective is one minute.
Cross-region protection differs from a local Autonomous Data Guard standby because geographical distance makes fully synchronous zero-data-loss transport impractical for normal cross-region operation. Oracle therefore uses an asynchronous protection model that provides a bounded RPO while maintaining acceptable database performance.
A useful distinction is between RPO and RTO. RPO measures potential data loss, whereas RTO measures how long restoration of service is targeted to take. Oracle's current Autonomous AI Database Serverless SLO documentation continues to specify a 1-minute cross-region RPO . Current documentation has improved the cross-region RTO target to less than 10 minutes ; older material cited approximately 15 minutes. That RTO evolution does not affect the answer to this question because only RPO is being tested.
Zero data loss therefore applies to different protection scenarios, while five and fifteen minutes do not represent the documented cross-region RPO.
Study Guide reference: HA and DR - Autonomous Data Guard local and cross-region recovery objectives.

NEW QUESTION: 39
Who performs infrastructure patching and software updates for the Exadata infrastructure underlying Oracle AI Database@Azure?
A. The customer's Azure administrators
B. A third-party managed service provider selected by the customer
C. Oracle's Cloud Infrastructure operations team
D. Microsoft's Azure support engineers
Answer: (SHOW ANSWER)

Oracle owns and operates the underlying Exadata infrastructure layer used by Oracle AI Database@Azure, even though that infrastructure is physically located in Microsoft Azure data centers.
Consequently, Oracle performs infrastructure-level maintenance, patching, security updates, monitoring, and lifecycle operations.
Oracle-managed Exadata infrastructure components include the physical database-server hosts, storage servers, network fabric, hypervisor or virtualization infrastructure, firmware, management components, and control-plane systems. Microsoft supplies the Azure data-center environment and associated Azure platform capabilities, but Microsoft support engineers do not maintain Oracle's Exadata hardware/software stack.
The customer retains responsibility for customer-controlled portions of co-managed database environments, including database configuration and, depending on service and maintenance model, guest VM, Grid Infrastructure, database homes, and database software lifecycle activities.
Oracle has also introduced optional Oracle-managed database software update capabilities for applicable newer Exadata services, but that does not change responsibility for the underlying infrastructure.
Therefore, C correctly identifies Oracle's Cloud Infrastructure operations organization.
Study Guide reference: Provision and operate Oracle AI Database@Azure Resources - shared operational responsibility and Exadata infrastructure maintenance.

NEW QUESTION: 40

Your organization has already deployed five Oracle Database@Azure delegated subnets across different virtual networks in a single Azure region. What is required if you need to provision a sixth delegated subnet in the same region?

Freecram.net

A. Deploy a transit virtual network utilizing an Azure Route Server to bypass the local regional limit.
B. Create a peer virtual network and route traffic through a separate OCI Parent Site.
C. Onboard a secondary Microsoft Entra tenant and establish cross-tenant VNet peering.
D. Submit a support ticket with Oracle to request a service limit increase.
Answer: (SHOW ANSWER)
Oracle AI Database@Azure applies a regional service limit to delegated subnet usage. The documented default is five delegated subnets per OCI tenancy in a specific region . This is a service quota rather than a limitation that can be bypassed through additional Azure routing constructs.
To exceed the default, the customer must request a service limit increase . Oracle identifies the relevant OCI service as Multicloud, the resource as Delegated Subnet Multicloud Links , and the corresponding service- limit name as azure-delegated-subnet-count. The limit is regional, so introducing another VNet, Route Server, VNet peering relationship, or Microsoft Entra tenant does not increase the permitted count.

Current documentation phrases the operation as requesting a service limit increase through Oracle's service- limit process. Among the supplied choices, option D is the only one that represents that required Oracle-side quota increase.
Architecturally, this distinction matters because routing topology and service capacity are independent control planes: networking objects determine reachability, whereas Oracle-enforced service limits determine how many delegated-subnet links may exist.
Study Guide reference: Configure Oracle AI Database@Azure Networking - delegated subnet limits and Multicloud service-limit management.

NEW QUESTION: 41
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 (LEAVE A REPLY)
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.

NEW QUESTION: 42
Which of the following is NOT a valid managed backup destination for an Exadata Database in Oracle AI Database@Azure?

A. Autonomous Recovery Service
B. Zero Data Loss Recovery Appliance
C. Object Storage
D. Long-term retention via Recovery Service
Answer: (SHOW ANSWER)
The invalid managed backup destination is the physical Zero Data Loss Recovery Appliance (ZDLRA) .
Oracle AI Database@Azure currently provides two primary managed automatic backup destinations for Exadata databases: Oracle Database Autonomous Recovery Service and OCI Object Storage . Oracle recommends Autonomous Recovery Service because it provides database-aware incremental-forever protection, policy-based retention, recoverability validation, optional real-time redo protection, and retention- lock capabilities.
Long-Term Retention is a supported capability of Recovery Service and allows selected backups to be preserved for extended compliance periods. It therefore represents a valid managed backup lifecycle option, even though technically it is a Recovery Service feature rather than an independent storage system.
The Zero Data Loss Recovery Appliance is a physical Oracle engineered system primarily used in customer data centers and appropriate Cloud@Customer architectures. Autonomous Recovery Service incorporates technologies derived from the Recovery Appliance but is delivered as a managed cloud service; these two should not be confused.
Therefore, the answer to the "NOT valid" question is B .
Study Guide reference: Provision and operate Oracle AI Database@Azure Resources - Exadata automatic backups, Object Storage and Autonomous Recovery Service.

NEW QUESTION: 43
Which two workflow categories does Oracle Zero Downtime Migration (ZDM) support for migrating databases to Oracle AI Database@Azure? (Choose two.)
A. Container-based migration
B. Logical migration
C. Snapshot-based migration
D. Physical migration
Answer: (SHOW ANSWER)
Oracle Zero Downtime Migration organizes its Oracle database migration methods into two fundamental workflow families: Physical Migration and Logical Migration . Each family can then have online and offline variants depending on compatibility requirements and acceptable downtime.
Physical migrations operate on the physical database representation and rely on technologies such as RMAN and Oracle Data Guard. Physical Online migration can instantiate a target using Restore From Service and maintain synchronization using Data Guard, while Physical Offline uses RMAN backup and restore.

Logical migrations operate through logical database objects and data. Oracle Data Pump provides the initial export/import mechanism, while Logical Online migration integrates Oracle GoldenGate to capture and replicate ongoing transactions before cutover. Logical workflows are also suitable for many cross-version or cross-platform scenarios that physical workflows cannot support.
"Container-based migration" and "snapshot-based migration" are not ZDM's top-level migration workflow categories.
Oracle's current ZDM documentation explicitly states that ZDM supports both Physical and Logical Migration workflows.
Therefore, B and D are correct.
Study Guide reference: Migrate Oracle Databases to Oracle AI Database@Azure - ZDM workflow categories.

NEW QUESTION: 44
Your team has configured a local standby database for an Oracle Autonomous Database@Azure deployment within the same region to protect against instance-level failures. What are the Service Level Objectives (SLOs) for Recovery Time Objective (RTO) and Recovery Point Objective (RPO) when utilizing this local standby configuration with automatic failover?
A. RTO of two minutes and RPO of 10 seconds
B. RTO of one hour and RPO of zero seconds
C. RTO of less than ten seconds and RPO of five minutes
D. RTO of 15 minutes and RPO of 30 seconds
Answer: (SHOW ANSWER)
Oracle Autonomous AI Database Serverless documents a local Autonomous Data Guard standby RTO SLO of two minutes for failure events that require failover to the standby.
For the Recovery Point Objective, current Oracle documentation is slightly more precise than the wording in option A: the local standby's maximum possible data-loss objective is less than 10 seconds . Therefore, the exam option stating "RPO of 10 seconds" represents the intended and closest supported value.
These SLOs apply to events such as full database failure, complete storage failure, corruption scenarios requiring standby transition, and availability-domain failures when Autonomous Data Guard is enabled. Local standby protection provides materially faster recovery and a substantially smaller possible data-loss window than cross-region standby protection because replication remains geographically close.
An hour-long RTO is inconsistent with active standby protection, while five-minute RPO and thirty-second RPO values do not match Oracle's stated local Autonomous Data Guard objectives.
Therefore, A is correct, interpreted technically as 2-minute RTO and < 10-second RPO .
Study Guide reference: HA and DR - Autonomous Data Guard local standby SLOs.

NEW QUESTION: 45

Which two key management options are available for protecting transparent data encryption (TDE) master keys on Exadata databases in Oracle AI Database@Azure? (Choose two.)

A. Oracle Wallet with Oracle-managed keys

B. AWS Key Management Service

C. Google Cloud KMS

D. Azure Key Vault with customer-managed keys

Answer: (SHOW ANSWER)

Oracle AI Database@Azure supports two broad TDE master-key management approaches: Oracle-managed keys and customer-managed keys .

With the default Oracle-managed model, the TDE master encryption key is generated and stored in an Oracle Wallet within the database environment. Oracle manages key generation, storage, rotation, backup, and restoration activities associated with this model.

For customer-controlled key management, Oracle supports several external key stores. For Exadata Database Service on Dedicated Infrastructure, these include Azure Key Vault , OCI Vault, and Oracle Key Vault.

Azure Key Vault integration is particularly useful for organizations that want their database encryption-key governance aligned with existing Azure key-management, auditing, and separation-of-duties practices.

AWS KMS and Google Cloud KMS are not native key-management choices for an Oracle AI Database@Azure Exadata deployment.

Therefore, the two correct selections from this answer set are A and D .

Study Guide reference: Secure Oracle AI Database@Azure - Transparent Data Encryption, Oracle Wallet, customer-managed keys and Azure Key Vault.

Valid 1z0-1147 Dumps shared by EduDump.com for Helping Passing 1z0-1147 Exam! EduDump.com now offer the **newest 1z0-1147 exam dumps**, the EduDump.com 1z0-1147 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com 1z0-1147 dumps with Test Engine here: https://www.edudump.com/exams/Oracle/1z0-1147/premium/ (**103** Q&As Dumps, **35%OFF** **Special Discount Code: freecram**)