Up-to-date Practice Test with Latest Questions and Answers covering latest syllabus and topics of the exam. Makes you ready to face actual exam.

AZ-700 Practice Questions
AZ-700 Practice Test
AZ-700 Practice Exam
AZ-700 Exam Questions
AZ-700 Study Guide

killexams.com

Microsoft

AZ-700

Designing and Implementing Microsoft Azure Networking Solutions

ORDER FULL VERSION

https://killexams.com/pass4sure/exam-detail/AZ-700

Question: 1412

You need to deploy Azure Application Gateway v2 with private frontend only, integrated with Private Link for backend services, and ensure isolation. The subnet must support up to 50 instances with autoscaling. Which subnet planning actions should you take?

A. Avoid UDRs on the Application Gateway subnet to maintain health probe and metric functionality
B. Use the same subnet for Application Gateway and Private Link configuration to simplify IP management
C. Deploy a separate dedicated subnet for Private Link IP configurations that contains no Application Gateway instances
D. Enable subnet delegation to Microsoft.Network/applicationGateways for the Application Gateway subnet
E. Create a dedicated subnet for Application Gateway with no name reservation and size at least /24 to accommodate instances plus private frontend IP

Answer: A,C,E
Explanation: Application Gateway requires a dedicated subnet (no specific name needed) sized appropriately (e.g., /24 for scaling to many instances plus frontend IP). Private Link configuration needs its own separate dedicated subnet exclusively for Private Link IPs, isolated from the gateway instances. UDRs should not be applied to the Application Gateway subnet to prevent issues with backend health, logs, and metrics visibility.

Question: 1413

Case Study
CloudServices Provider (CSP) manages several small clients who need to extend their on-premises subnets into Azure without changing their IP addresses. One client, "LegacySoft," has an on-premises Windows Server 2022 cluster running an old application that hardcodes the IP address 192.168.10.50. This server must be moved to Azure.
CSP decides to implement the Azure Extended Network (AEN) feature. They have set up a VNet in Azure with two subnets: a "Managed" subnet for standard Azure VMs and an "Extended" subnet that mirrors the on-premises IP space. CSP is using an Azure VPN Gateway (VpnGw2) to facilitate this. They also need to provide Point-to-Site VPN access for the client's administrators. The administrators use a mix of Windows 11 and macOS devices. The CSP wants to use a single gateway to support both the S2S Extended Network and the P2S administrative access.

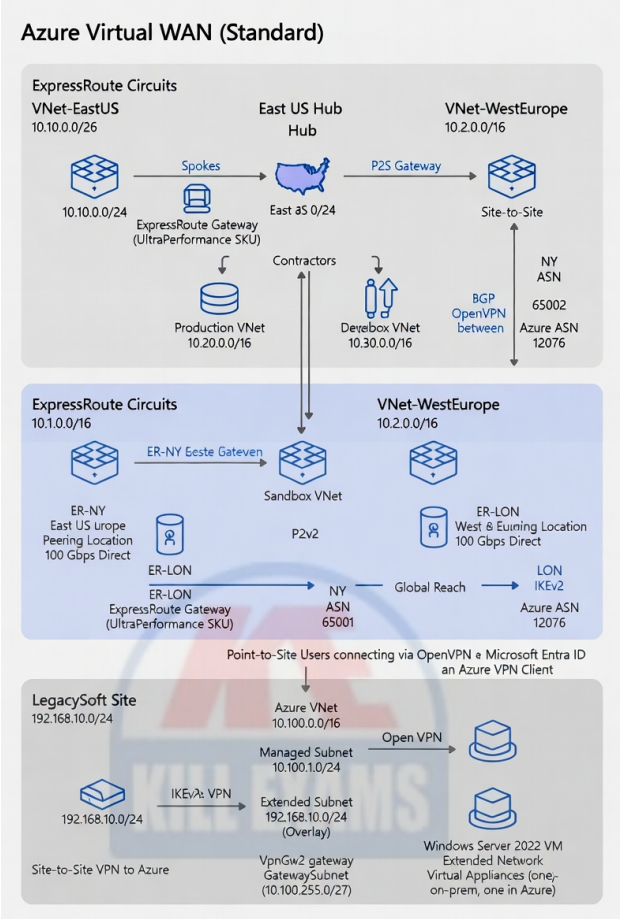

Which components are required to be installed and configured to support the Azure Extended Network for LegacySoft? (Select All that Apply)

A. Windows Admin Center (WAC)
B. Two Windows Server VMs (one on-premises, one in Azure) to act as appliances
C. Azure Extended Network tool extension in WAC
D. An ExpressRoute circuit with Microsoft Peering
E. A Layer 3 switch with VXLAN support on-premises

Answer: A,B,C
Explanation: Azure Extended Network (AEN) is a specialized solution that allows you to stretch an on-premises Layer 2 subnet into Azure. The primary management interface for setting up and maintaining AEN is Windows Admin Center (WAC), which requires the "Azure Extended Network" tool extension. The architecture depends on two "appliance" VMs: one running on-premises (typically as a Hyper-V VM) and one running in Azure. These appliances establish a VXLAN tunnel between them over a standard S2S VPN or ExpressRoute connection, effectively "bridging" the two locations at Layer 2. This allows the LegacySoft server to retain its 192.168.10.50 address in Azure while communicating with other servers on the same 192.168.10.0/24 subnet still located on-premises. No specific hardware VXLAN support is required on the physical switches because the encapsulation is handled by the Windows Server appliances.

Question: 1414

You need to update an existing Site-to-Site VPN connection to use a new custom IPsec policy. Which Azure PowerShell cmdlet should you use to apply the policy to the existing connection object?

A. Set-AzVirtualNetworkGatewayConnection
B. New-AzIpsecPolicy
C. Set-AzVirtualNetworkGateway
D. Update-AzLocalNetworkGateway

Answer: A
Explanation: The `Set-AzVirtualNetworkGatewayConnection` cmdlet is used to update the properties of an existing connection, including the assignment of a custom IPsec/IKE policy. The `New-AzIpsecPolicy` cmdlet only creates the policy object in memory; it must then be applied to the connection using the Set command.

Question: 1415

A company is implementing Azure Load Balancer within a new infrastructure setup. What practices should be followed to maximize load balancing efficiency?

A. Establish multiple health probes for redundancy.
B. Restrict load balancing to only web applications.
C. Utilize both layer 4 and layer 7 load balancing for flexibility.
D. Plan for potential scale by choosing an appropriate SKU.
E. Set a session affinity for specific applications that require it.

Answer: C,D,E
Explanation: Utilizing both layer 4 and layer 7 load balancing offers flexibility and efficiency in directing traffic according to service requirements. Setting session affinity ensures specific user sessions are maintained effectively. Planning for potential scaling by choosing an appropriate SKU helps accommodate growing traffic.

Question: 1416

Your organization is designing a high-availability hub-and-spoke architecture. The hub virtual network (10.10.0.0/16) must support an Azure Firewall Premium instance, an Azure Bastion (Standard SKU) with IP-based connection enabled, and a VPN Gateway. You need to ensure the subnets meet the minimum size requirements and naming conventions for these services. Which of the following configurations must be implemented?

A. Use a /28 prefix for the AzureBastionSubnet to minimize address waste
B. Allocate a subnet named AzureBastionSubnet with at least a /26 prefix
C. Allocate a subnet named GatewaySubnet with at least a /27 prefix
D. Create a subnet named AzureFirewallManagementSubnet with a /28 prefix for the firewall's management traffic

E. Allocate a subnet named AzureFirewallSubnet with at least a /26 prefix

Answer: B,C,E
Explanation: Azure Bastion requires a dedicated subnet named AzureBastionSubnet with a minimum size of /26 to support features like scaling and IP-based connection. Azure Firewall requires a dedicated subnet named AzureFirewallSubnet with a minimum size of /26 for its operational nodes. The VPN Gateway requires a dedicated subnet named GatewaySubnet; while a /29 is the absolute minimum, Microsoft recommends at least a /27 to ensure enough addresses for gateway instances and to avoid future re-addressing if an ExpressRoute gateway is added.

Question: 1417

An organization requires their Public IP addresses to be "Global" so they can be associated with a Cross-region Load Balancer. Which condition must be met when creating a Public IP address for this purpose?

A. The IP must be created with the SKU set to Standard and the Tier set to Global.
B. The IP must be created in a VNet that is peered across regions.
C. The IP must be created with the SKU set to Premium and the Tier set to Regional.
D. The IP must be created in the "Global" region and use the Basic SKU.

Answer: A
Explanation: For Cross-region Load Balancers, the frontend Public IP must be of the Standard SKU and have its Tier property set to Global. This allows the IP address to be advertised globally and route traffic to regional load balancer backends.

Question: 1418

You are monitoring a production application using Azure Monitor Network Insights. You observe a sudden spike in "Hyper-V Network Virtualization" drops on several VMs in the same availability set. The VMs are healthy and there are no NSG deny logs. What is the most probable cause of these drops?

A. The Azure Host is experiencing transient failures
B. The VM's accelerated networking feature is disabled on the subnet
C. The VNet address space overlaps with an on-premises network causing encapsulation conflicts
D. The VMs are experiencing a Distributed Denial-of-Service (DDoS) attack

Answer: C
Explanation: "Hyper-V Network Virtualization" drops in Azure Network Insights specifically indicate issues with the virtual network encapsulation layer. One of the most common causes is overlapping IP address spaces between a VNet and an on-premises network or another VNet connected via peering or VPN. When Azure attempts to encapsulate traffic, address conflicts cause packet drops at the hypervisor level. NSG denies would not show here as they are a separate control plane.

Question: 1419

Case study
Proseware Inc. is implementing a secure hybrid network for its e-commerce site. The architecture uses Azure Application Gateway v2 (WAF) as the ingress controller for AKS clusters hosting microservices. Specifics:
- AKS cluster in `aks-subnet` (10.1.1.0/24).
- App Gateway v2 in `appgw-subnet` (10.1.2.0/24).
- A backend microservice for checkout requires mutual TLS (mTLS) where the App Gateway must present a client certificate to the backend pod.
- A routing rule must direct traffic for `www.proseware.com/checkout/*` to this mTLS-enabled backend pool.
- The WAF policy must be in Prevention mode and use the OWASP 3.2 rule set.

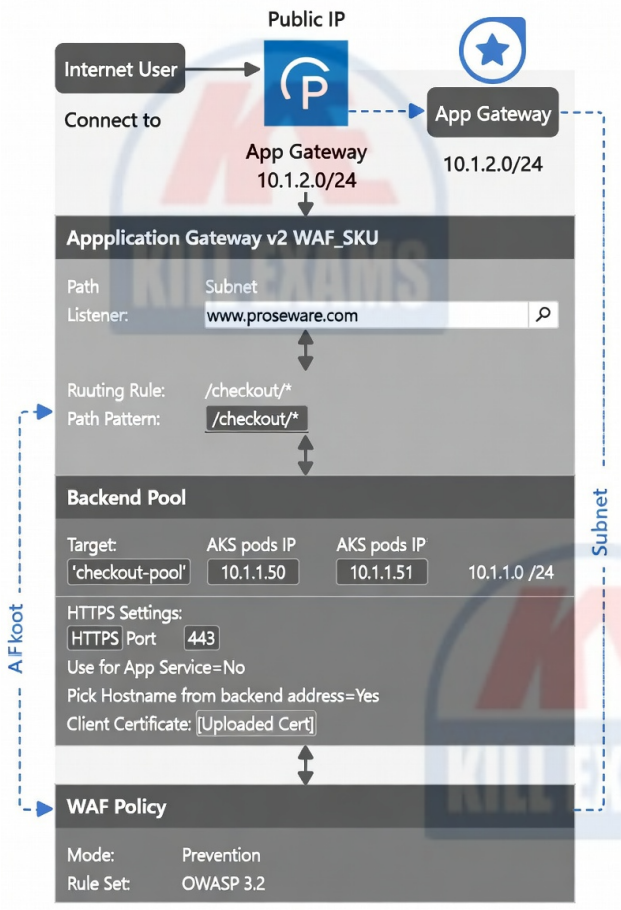

The checkout microservice pods are scaled by AKS Horizontal Pod Autoscaler. How should the Application Gateway backend pool be configured to dynamically discover these pod IPs? (Select all that apply)

A. Configure the backend pool with an FQDN that resolves to the Kubernetes Service's ClusterIP.
B. Use a Virtual Network type backend pool and target the AKS subnet (`10.1.1.0/24`).
C. Manually add the IP addresses of each AKS node to the backend pool.
D. Integrate Application Gateway with AKS using the Application Gateway Ingress Controller (AGIC) add-on, which manages the backend pool membership.
E. Create the backend pool with target type "App Service" and select the AKS cluster resource.

Answer: D

Explanation: For dynamic discovery of AKS pod IPs in a scalable manner, the recommended pattern is to use the Application Gateway Ingress Controller (AGIC). AGIC is a Kubernetes pod that watches the Kubernetes API for Ingress resource definitions and automatically configures the associated Application Gateway (including backend pool membership, HTTP settings, and routing rules) to reflect the state of the AKS cluster. It dynamically updates the backend pool with the private IPs of the pods matching the Kubernetes service selectors, handling pod creation and deletion seamlessly.

Question: 1420

After provisioning an inbound endpoint (IPs 10.40.0.4 and 10.40.0.5) in subnet dns-in/28, on-premises conditional forwarders point to these IPs for corp.internal but queries time out sporadically. Which configurations resolve the issue?

A. Link the corp.internal private DNS zone to the resolver VNet
B. Add outbound ruleset forwarding corp.internal back to on-premises
C. Configure NSG on the inbound endpoint subnet to allow UDP/TCP 53 from on-premises CIDR
D. Enable auto-registration on the corp.internal zone link
E. Increase inbound endpoint IP count to 10

Answer: A,C
Explanation: Linking corp.internal to the resolver VNet enables authoritative resolution via inbound endpoint. NSG rules permitting UDP/TCP port 53 inbound from on-premises sources prevent query blocking and timeouts.

Question: 1421

Case study
HealthData Systems uses Azure to process sensitive patient records. They have a "Locked-Down" VNet (10.80.0.0/16) with no internet access (0.0.0.0/0 next hop is None). They need to use Azure App Service for their API, but the App Service must access an Azure SQL database that is restricted to private access only. The security team insists on using Service Endpoints for the App Service to reach SQL, but they are concerned about "data exfiltration" where a compromised App Service could send data to a malicious SQL server in a different Azure AD tenant.

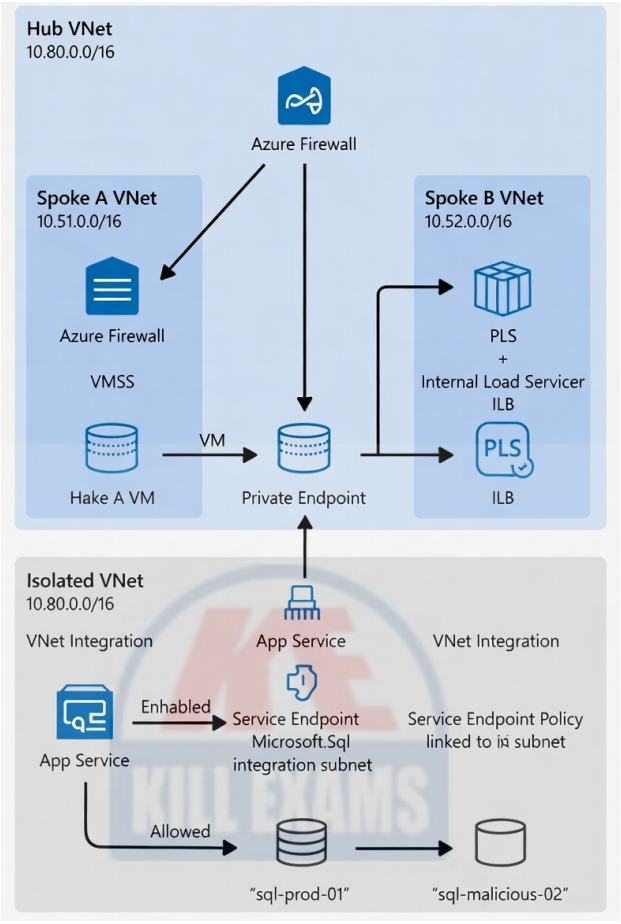

The App Service requires "Regional VNet Integration" to access the SQL database privately. Which three statements are true regarding Regional VNet Integration requirements? (Select three)

A. The App Service and the VNet must be in the same Azure region
B. The integration subnet can host other resources like VMs or Gateways
C. The integration subnet must have at least a /28 prefix
D. The integration subnet must be delegated to 'Microsoft.Web/serverFarms'
E. The App Service plan must be in the Standard, Premium, or Isolated tier

Answer: A,D,E
Explanation: Regional VNet Integration requires a supported App Service Plan (Standard or higher) and the App Service must be in the same region as the VNet. The dedicated subnet must be delegated specifically to `Microsoft.Web/serverFarms` and cannot contain other non-App Service resources. (Note: A /28 is a recommended minimum, but not a strict platform requirement for the functionality itself, whereas delegation and regional parity are).

Question: 1422

An Azure Administrator creates a VNet with the address space 10.0.0.0/16 and a subnet 10.0.0.0/24. They attempt to create an Azure Firewall in that subnet but receive an error. What is the most likely cause?

A. The subnet is too small for Azure Firewall.
B. The VNet is in a region that doesn't support firewalls.
C. The subnet is not named AzureFirewallSubnet.
D. There is already a VM in the subnet.

Answer: C
Explanation: Azure Firewall has a strict requirement that it must be deployed into a subnet named exactly "AzureFirewallSubnet". If the subnet has any other name, the deployment of the Azure Firewall resource will fail.

Question: 1423

Can a point-to-site VPN connection leverage both Microsoft Entra ID and RADIUS simultaneously for user authentication? If so, what key configuration aspect should be noted for this setup?

A. Both authentication methods must use the same user credential store
B. Configuration of a custom policy in Azure AD for dual authentication
C. Explicit setting in the VPN client to choose the authentication method
D. Clear distinction in user groups with different authentication requirements
E. The virtual network gateway cannot support simultaneous methods

Answer: A,B
Explanation: While it's possible to set up both Microsoft Entra ID and RADIUS for authentication, they must share the same user credential store to function correctly. Configuring a custom policy in Azure AD may also facilitate this dual-authentication setup.

Question: 1424

You are deploying a Virtual Network Gateway in a VNet with the address space 10.0.0.0/16. You need to create the GatewaySubnet. What is the minimum recommended prefix length for the GatewaySubnet to ensure future-proofing for co-existing ExpressRoute and VPN gateways?

A. /29
B. /28
C. /24
D. /27

Answer: D
Explanation: Microsoft recommends a minimum prefix length of /27 for the GatewaySubnet. While a /28 is the absolute minimum allowed, a /27 (or /26) is preferred to ensure there are enough IP addresses for the gateway instances, especially when running in Active-Active mode or when co-existing with an ExpressRoute gateway.

Question: 1425

You need a static outbound IP for compliance in a subnet with NAT gateway already attached. Additional static IPs are required for different workloads. What approach supports this?

A. Assign instance-level public IPs
B. Configure UDR to different NAT gateways
C. Deploy Standard Load Balancer with outbound rules per workload
D. Use multiple NAT gateways per subnet
E. Attach multiple public IPs to the NAT gateway

Answer: C,E
Explanation: Attach multiple public IPs to the NAT gateway provides additional static source IPs for SNAT, supporting multiple allow-listed ranges from the same subnet. Deploy Standard Load Balancer with outbound rules per workload offers workload-specific static IPs as an alternative.

Question: 1426

An ExpressRoute Direct deployment supports 100 Gbps aggregate bandwidth across ports. The gateway must handle 1 million packets per second.
Which TWO gateway configurations are appropriate?

A. Configure private peering with large route advertisements
B. Deploy ErGwScale with sufficient scale units for packet rate
C. Implement route summarization aggressively
D. Use ErGw3Az for fixed high-performance baseline
E. Enable FastPath to offload gateway processing

Answer: B,E
Explanation: ErGwScale allows scaling to meet high packet-per-second demands beyond fixed SKUs. FastPath reduces gateway involvement for compatible flows, preserving capacity for other traffic.

Question: 1427

In your Azure environment, you are planning to implement a highly segmented network architecture. What strategies can you adopt to manage IP addresses effectively?

A. Use Azure's IP Address Manager for dynamic allocation
B. Maintain a manual spreadsheet for tracking IP addresses
C. Create multiple address spaces and associate them with subnets
D. Implement subnet CIDR notations based on departmental needs
E. Consolidate all VMs to a single subnet for easier management

Answer: B,C,D

Explanation: Creating different address spaces for various segments allows greater control and better organization. CIDR notations based on departmental needs facilitate clear usage. A spreadsheet aids in tracking allocations, despite it being less dynamic than other methods.

Question: 1428

Case Study

Contoso Pharma is building a highly secure research environment in Azure. The design includes multiple perimeter networks (DMZs) using Azure Firewall and network virtual appliances (NVAs) from a partner. They use a hub-spoke with a dedicated "Inspection" VNet hosting the NVA cluster. All traffic from on-premises (via ExpressRoute) to the internet, and from Azure spokes to on-premises, must pass through this NVA cluster. They also implement service endpoints for PaaS services but need to ensure data exfiltration is prevented.

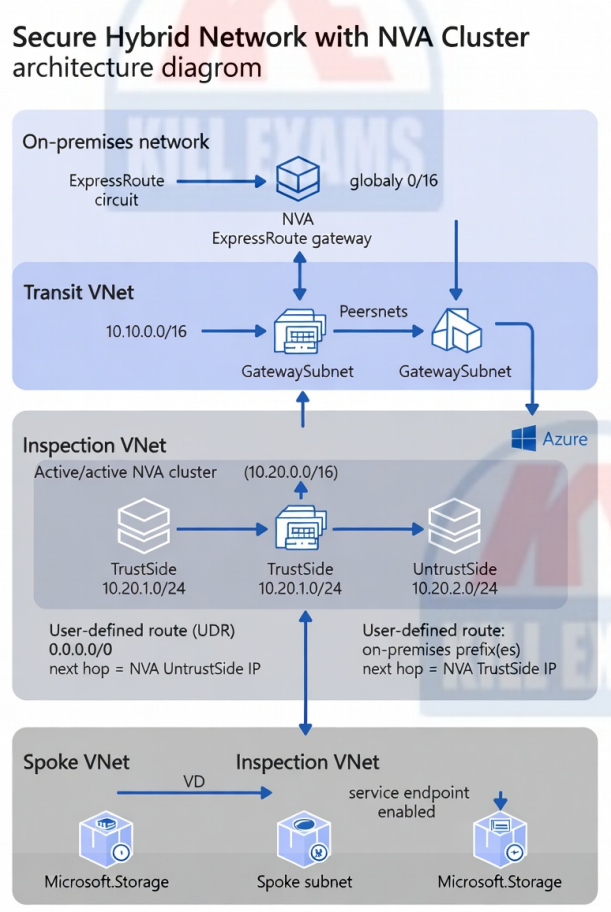

You need to configure routing so that traffic from the spoke VNets to on-premises prefixes is sent to the NVA cluster's trust-side interface. The Inspection VNet is peered to the spokes. Which action achieves this?

A. Enable "Use remote gateways" on the spoke VNet peerings to the Inspection VNet.

B. Propagate on-premises routes (learned via BGP from ExpressRoute) to the spoke VNets via peering.

C. Add a user-defined route table to the spoke subnets with a route for on-premises prefixes, next hop = NVA trust-side IP.

D. Disable "Allow gateway transit" on the Inspection VNet's peering to the spokes.

Answer: C
Explanation: In this hub-spoke design with a dedicated inspection VNet (acting as the hub), you want spoke-to-on-premises traffic to go through the NVA) The on-premises routes are learned by the ExpressRoute gateway in the Transit VNet. For the spokes to reach on-premises, they need a route. Since the Inspection VNet is not the one with the gateway (Transit VNet is), gateway transit settings are not the primary mechanism. The most direct method is to apply a UDR to the spoke subnets. This UDR should contain a route for the on-premises address prefixes (or a summary) with a next hop type of 'Virtual Appliance' and the address set to the NVA's trust-side IP address in the Inspection VNet.

Question: 1429

In a multi-hub dual-homed spoke scenario with Route Server in the spoke, what enables dynamic route selection from multiple hubs?

A. Advertise hub routes with different MED values
B. Use UDRs pointing to each hub
C. Configure Virtual Network Manager mesh
D. Deploy Route Server in spoke and peer with hub NVAs
E. Enable global VNet peering to both hubs

Answer: A,D
Explanation: Deploy Route Server in spoke and peer with hub NVAs allows the spoke Route Server to learn routes dynamically from multiple hubs. Advertise hub routes with different MED values influences BGP best-path selection for dynamic preference in dual-homed connectivity.

Question: 1430

A project requires a temporary ExpressRoute circuit for a data migration that will last 48 hours. You want to minimize costs while maintaining high performance. Which billing model should you choose?

A. Pay-As-You-Go Gateway SKU
B. Metered Data
C. Reserved Capacity
D. Unlimited Data

Answer: B
Explanation: For a short-term, high-bandwidth migration, the Metered Data billing model is the most appropriate. With Metered Data, you pay a lower monthly port fee and are charged based on the amount of data transferred. This is usually cheaper than the Unlimited model unless you are transferring massive amounts of data consistently over a long period.

KILLEXAMS.COM

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including Exam Questions, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.

Practice Exam Questions Based on Current Exam Objectives

Killexams.com provides practice exam questions aligned with the latest official exam objectives and latest syllabus. Our content is reviewed and updated regularly to reflect recent changes announced by certification vendors. By studying these practice questions, candidates will cover the structure, difficulty level, and topics of the actual exam, helping them prepare more effectively and efficiently.

Comprehensive Practice Exams (PDF Format)

Killexams.com offers multiple-choice questions (MCQs) in easy-to-read PDF format, covering all major domains of the exam. Each PDF contains a structured collection of practice questions and verified answers designed to support focused study. These MCQs help candidates reinforce key concepts, identify knowledge gaps, and improve exam readiness through consistent practice.

Realistic Practice Tests (Online Test Engine & Desktop Test Engine)

To support hands-on preparation, Killexams.com provides practice tests through both an Online Test Engine and a Desktop Test Engine. These tools are designed to simulate a real exam environment, allowing candidates to practice under exam-like conditions, with latest syllbus and topics of the exam. Performance tracking, test history, and result analysis help users evaluate their progress and focus on areas that need improvement.

Risk-Free Purchase Policy

Killexams.com follows a transparent and customer-friendly purchase policy. If users are not satisfied with the study materials, they may request assistance or a refund in accordance with our published terms and conditions. This policy reflects our commitment to customer satisfaction, fairness, and confidence in our preparation resources.

Regularly Updated Content

Our practice question bank is reviewed and updated on an ongoing basis to stay aligned with the latest exam outlines and vendor updates. This ensures candidates are studying up-to-date, relevant material, and preparing with content that reflects current exam expectations, helping them stay confident and well-prepared.