ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

What is the VMware recommended way to deploy a virtual NSX Edge Node?

A. Through the NSX Ul
B. Through automated or interactive mode using an ISO
C. Through the vSphere Web Client
D. Through the OVF command line tool

Answer: B

Explanation:

VMware recommends deploying a virtual NSX Edge Node using an ISO in either automated or interactive mode. This method provides flexibility and ensures that the NSX Edge node is deployed properly with all the necessary configurations. Using an ISO allows for a more streamlined and controlled deployment process, especially in larger environments.

Question: 2

Which three selections are capabilities of Network Topology? (Choose three.)

A. Display how the different NSX components are interconnected.
B. Display the VMs connected to Segments.
C. Display how the Physical components are interconnected.
D. Display the uplinks configured on the Tier-1 Gateways.
E. Display the uplinks configured on the Tier-0 Gateways.

Answer: A, B, C

Explanation:

Display how the different NSX components are interconnected.
Network Topology in NSX provides a visual representation of how different NSX components (like Edge nodes, Logical Routers, and other NSX components) are interconnected.
Display the VMs connected to Segments.
It also allows you to see which VMs are connected to specific segments (logical switches).
Display how the Physical components are interconnected.
The Network Topology view includes information about how physical network components are connected, providing a comprehensive overview of both the virtual and physical networking infrastructure.

Question: 3

An NSX administrator has deployed a single NSX Manager node and will be adding two additional nodes to form a 3-node NSX Management Cluster for a production environment. The administrator will deploy these two additional nodes and Cluster VIP using the NSX UI.
What two are the prerequisites for this configuration? (Choose two.)

A. The cluster configuration must be completed using API.
B. All nodes must be in the same subnet.
C. All nodes must be in separate subnets.
D. A compute manager must be configured.
E. NSX Manager must reside on a Windows Server.

Answer: B, D

Explanation:

For a 3-node NSX Manager cluster, all nodes must be within the same subnet to ensure proper communication and functionality between them.
A compute manager must be configured before adding nodes to the cluster, as it provides the necessary integration between the NSX Manager and the underlying virtualization infrastructure (such as vSphere or vCenter).

Question: 4

Which two commands does an NSX administrator use to check the IP address of the VMkernel port for the Geneve protocol on the ESXi transport node? (Choose two.)

A. net-dvs
B. esxcfg-nics -l
C. esxcli network ip interface ipv4 get
D. esxcfg-vmknic -l
E. esxcli network nic list

Answer: C

Explanation:

The esxcli network ip interface ipv4 get command is used to display the IP address configuration of the VMkernel network interfaces, including those used for the Geneve protocol.
The esxcfg-vmknic -l command lists all VMkernel network interfaces, including their IP addresses, which can

help identify the VMkernel port for the Geneve protocol.

Question: 5

Which two are supported by L2 VPN clients? (Choose two.)

A. NSX Autonomous Edge
B. NSX Edge
C. NSX for vSphere Edge
D. 3rd party Hardware VPN Device

Answer: B, D

Explanation:

The NSX Edge supports L2 VPN (Layer 2 VPN) functionality, which allows it to connect different Layer 2 networks over an IP transport.
Third-party hardware VPN devices can also be used as L2 VPN clients, providing connectivity between different Layer 2 networks through an external device.

Question: 6

As part of an organization's IT security compliance requirement, NSX Manager must be configured for 2FA (two-factor authentication).
What should an NSX administrator have ready before the integration can be configured?

A. Active Directory LDAP integration with ADFS
B. VMware Identity Manager with NSX added as a Web Application
C. VMware Identity Manager with an OAuth Client added
D. Active Directory LDAP integration with OAuth Client added

Answer: B

Explanation:

To enable two-factor authentication (2FA) for NSX Manager, VMware Identity Manager must be configured and integrated with NSX. The NSX Manager should be added as a web application in VMware Identity Manager, which will allow 2FA to be applied during the authentication process. VMware Identity Manager supports 2FA methods, including integration with external identity providers, and it can manage access to NSX with additional security layers.

Question: 7

What should an NSX administrator check to verify that VMware Identity Manager integration is successful?

A. From the NSX Ul the status of the VMv/are Identity Manager Integration must be Enabled'
B. From the NSX CLI the status of the VMware Identity Manager Integration must be Configured'
C. From VMware Identity Manager the status of the remote access application must be green
D. From the NSX Ul the URI in the address bar must have locaMalstf part of it.

Answer: B

Explanation:

To verify that VMware Identity Manager integration is successful with NSX, the administrator should check the NSX UI for the integration status. If it is configured correctly, the status should be marked as "Enabled," indicating that the integration is active and functioning.

Question: 8

An administrator has been tasked with implementing the SSL certificates for the NSX Manager Cluster VIP. Which is the correct way to implement this change?

A. Send an API call to https://<nsx-mgr>/api/vl/cluster/api-certificate?action=set_cluster_certificate&certificate_id=<certificate_id>
B. Send an API call to https://<nsx-mgr>/api/vl/node/services/http?action=apply_certificate&certificate_id=<certificate_id>
C. SSH as admin into the NSX manager with the cluster VIP IP and run nsxcli cluster certificate node install <certificate_id>
D. SSH as admin into the NSX manager with the cluster VIP IP and run nsxcli cluster certificate vip install <certificate_id>

Answer: D

Explanation:

To implement SSL certificates for the NSX Manager Cluster VIP, the correct method is to SSH into the NSX Manager (using the Cluster VIP IP) and run the nsxcli cluster certificate vip install <certificate_id> command. This command installs the SSL certificate for the VIP, ensuring that the cluster's SSL certificate is properly configured for secure communications.

Question: 9

An administrator wants to validate the BGP connection status between the Tier-0 Gateway and the upstream physical router.
What sequence of commands could be used to check this status on NSX Edge node?
A.

- enable <LR-D>
- get vrf <ID>
- show bgp neighbor

B.

- get gateways
- vrf <number>
- get bgp neighbor

C.

- set vrf <ID>
- show logical-routers
- show <LR-D> bgp

D.

- show logical-routers
- get vrf
- show ip route bgp

Answer: A

Explanation:

To validate the BGP connection status between the Tier-0 Gateway and the upstream physical router on an NSX Edge node, the correct sequence involves enabling the specific logical router (Tier-0 Gateway), checking the VRF (Virtual Routing and Forwarding) context, and then using the show bgp neighbor command to view the BGP session status.

enable <LR-D>: This command enables the logical router interface (Tier-0 Gateway) to access its configuration.

get vrf <ID>: This command checks the specific VRF (used for routing separation) to see the associated routing table.

show bgp neighbor: This command displays the status of the BGP connection, including details about the neighbor relationships and their state.

Question: 10

What is VMware's recommendation for the minimum MTU requirements when planning an NSX deployment?

A. MTU should be set to 1700 or greater across the data center network including inter-data center connections.

B. MTU should be set to 1500 or less only on inter-data center connections.

C. Configure Path MTU Discovery and rely on fragmentation.

D. MTU should be set to 1550 or less across the data center network including inter-data center connections.

Answer: A

Explanation:

VMware recommends setting the MTU (Maximum Transmission Unit) to 1700 or greater for NSX deployments. This is to ensure that the VXLAN encapsulation, which adds overhead to the original Ethernet frame, can be accommodated without fragmentation. This MTU requirement includes the entire data center network, including inter-data center connections, to ensure consistent communication across all network components involved in the NSX deployment.

Question: 11

In which VPN type are the Virtual Tunnel interfaces (VTI) used?

A. SSL-based VPN
B. Route & SSL based VPNs
C. Policy & Route based VPNs
D. Route-based VPN

Answer: D

Explanation:

Virtual Tunnel Interfaces (VTI) are used in route-based VPNs. In this type of VPN, the tunnel is treated like a regular interface on the router. This allows for the configuration of routing protocols and the application of routing decisions to the traffic flowing through the VPN tunnel. VTIs simplify the management of routing and make it more flexible in VPN scenarios.

Question: 12

In an NSX environment, an administrator is observing low throughput and congestion between the Tier-0 Gateway and the upstream physical routers.
Which two actions could address low throughput and congestion? (Choose two.)

A. Configure ECMP on the Tier-0 gateway.
B. Configure a Tier-1 gateway and connect it directly to the physical routers.
C. Deploy Large size Edge node/s.
D. Configure NAT on the Tier-0 gateway.
E. Add an additional vNIC to the NSX Edge node.

Answer: A, C

Explanation:

Configure ECMP on the Tier-0 gateway: ECMP (Equal-Cost Multi-Path) allows multiple paths for traffic between the Tier-0 Gateway and the upstream physical routers, effectively distributing the traffic load and improving throughput. By enabling ECMP, you can reduce congestion and increase bandwidth utilization, thus addressing performance issues.
Deploy Large size Edge node/s: Deploying larger Edge nodes can provide more resources (CPU, memory, and network interfaces) to handle higher throughput and reduce congestion. This is especially important if the existing Edge node is overwhelmed by the amount of traffic.

Question: 13

A company security policy requires all users to log into applications using a centralized authentication system.

Which two authentication, authorization, and accounting (AAA) systems are available when integrating NSX with VMware Identity Manager? (Choose two.)

A. RSA SecureID
B. SecureDAP
C. RADII 2.0
D. LDAP and OpenLDAP based on Active Directory (AD)
E. Keygen Enterprise

Answer: A, D

Explanation:

RSA SecureID: RSA SecureID is a commonly used two-factor authentication (2FA) system that can integrate with VMware Identity Manager for enhanced security during authentication, making it a suitable AAA system for user authentication.
LDAP and OpenLDAP based on Active Directory (AD): VMware Identity Manager can integrate with LDAP and OpenLDAP directories, including Active Directory (AD), for centralized user authentication. This allows users to authenticate against an organization's directory service.

Question: 14

An NSX administrator would like to export syslog events that capture messages related to NSX host preparation events.
Which message ID (msgid) should be used in the syslog export configuration command as a filter?

A. FABRIC
B. SYSTEM
C. GROUPING
D. MONITORING

Answer: A

Explanation:

In NSX, the FABRIC message ID is used to capture and export syslog events related to host preparation and other fabric-related activities. These events are important for tracking and troubleshooting the setup and configuration of NSX components across the fabric, including host preparation events.

Question: 15

An NSX administrator wants to create a Tier-0 Gateway to support equal cost multi-path (ECMP) routing.
Which failover detection protocol must be used to meet this requirement?

A. Host Standby Router Protocol (HSRP)
B. Beacon Probing (BP)

C. Virtual Router Redundancy Protocol (VRRP)
D. Bidirectional Forwarding Detection (BFD)

Answer: D

Explanation:

To support Equal-Cost Multi-Path (ECMP) routing in an NSX environment, Bidirectional Forwarding Detection (BFD) must be used for failover detection. BFD is a rapid failure detection protocol that works with ECMP to provide fast failure detection between routers. It helps in detecting link failures more quickly than traditional protocols, ensuring that traffic is routed through available paths as quickly as possible.

Question: 16

An administrator has connected two virtual machines on the same overlay segment. Ping between both virtual machines is successful.
What type of network boundary does this represent?

A. Layer 2 bridge
B. Layer 2 broadcast domain
C. Layer 2 VPN
D. Layer 3 route

Answer: B

Explanation:

When two virtual machines are connected on the same overlay segment, they are part of the same Layer 2 broadcast domain. In this case, the communication between the two VMs is happening within the same broadcast domain, which means that broadcast traffic can be sent to all devices on the segment. Since the ping is successful, the two VMs can communicate directly over Layer 2 without needing routing.

Question: 17

What are two supported host switch modes? (Choose two.)

A. Overlay Datapath
B. Secure Datapath
C. Standard Datapath
D. Enhanced Datapath
E. DPDK Datapath

Answer: C, D

Explanation:

Standard Datapath: This is the traditional mode used by the NSX host switch. It is typically used in environments where performance requirements are standard and no special acceleration techniques are

needed.

Enhanced Datapath: This mode is designed to improve performance and provide better scalability, especially for environments with higher traffic loads or more demanding applications. It can provide better performance in certain scenarios by improving packet processing efficiency.

Question: 18

Which is an advantage of an L2 VPN in an NSX 4.x environment?

A. Achieve better performance
B. Use the same broadcast domain
C. Enables Multi-Cloud solutions
D. Enables VM mobility with re-IP

Answer: B

Explanation:

An L2 VPN (Layer 2 VPN) in an NSX 4.x environment allows you to extend a Layer 2 network across different sites or data centers. This enables the connected environments to share the same broadcast domain, meaning that broadcast traffic can be transmitted between sites as if they were on the same local network. This is particularly useful for scenarios where you need to maintain Layer 2 connectivity across geographically dispersed locations.

Question: 19

Which two steps must an NSX administrator take to integrate VMware Identity Manager in NSX to support role-based access control? (Choose two.)

A. Create a SAML authentication in VMware Identity Manager using the NSX Manager FQDN.
B. Add NSX Manager as a Service Provider (SP) in VMware Identity Manager.
C. Enter the Identity Provider (IdP) metadata URL in NSX Manager.
D. Enter the service URL, Client Secret, and SSL thumbprint in NSX Manager.
E. Create an OAuth 2.0 client in VMware Identity Manager.

Answer: B, C

Explanation:

Adding NSX Manager as a Service Provider (SP) in VMware Identity Manager is necessary to enable SAML-based single sign-on (SSO), which allows VMware Identity Manager to manage and authenticate users accessing NSX.

Entering the Identity Provider (IdP) metadata URL in NSX Manager is required to establish a connection between NSX and VMware Identity Manager, enabling NSX to use VMware Identity Manager as the IdP for authentication.

Question: 20

Which of the two following characteristics about NAT64 are true? (Choose two.)

A. NAT64 requires the Tier-1 gateway to be configured in active-active mode.
B. NAT64 is stateless and requires gateways to be deployed in active-standby mode.
C. NAT64 is supported on Tier-0 and Tier-1 gateways.
D. NAT64 is supported on Tier-1 gateways only.
E. NAT64 requires the Tier-1 gateway to be configured in active-standby mode.

Answer: C, E

Explanation:

NAT64 is supported on both Tier-0 and Tier-1 gateways, allowing for IPv6-to-IPv4 address translation at different gateway levels within NSX.
NAT64 requires the Tier-1 gateway to be configured in active-standby mode, as this configuration ensures stateful translation and consistency for IPv6-to-IPv4 traffic handling.

Question: 21

Which VMware GUI tool is used to identify problems in a physical network?

A. VMware Aria Operations Networks
B. VMware Aria Automation
C. VMware Site Recovery Manager
D. VMware Aria Orchestrator

Answer: A

Explanation:

VMware Aria Operations Networks (formerly known as vRealize Network Insight) is a tool specifically designed for network visibility and troubleshooting. It provides insights into both virtual and physical network infrastructures, making it ideal for identifying problems in a physical network.

Question: 22

Which three protocols could an NSX administrator use to transfer log messages to a remote log server? (Choose three.)

A. HTTPS
B. SSH
C. TCP
D. UDP
E. SSL
F. TLS

Answer: C, D, F

Explanation:

Both TCP and UDP are commonly used protocols for transferring log messages in syslog configurations. TCP is preferred when reliability is needed, while UDP is used for faster, connectionless transmission.
TLS can be used to secure the log messages being sent over TCP, ensuring encrypted transmission to the remote log server.

Question: 23

Where does an administrator configure the VLANs used in VRF Lite? (Choose two.)

A. uplink interface of the VRF gateway
B. uplink interface of the default Tier-0 gateway
C. uplink trunk segment
D. segment connected to the Tier-1 gateway

Answer: A, D

Explanation:

The VLANs used in VRF Lite are configured on the uplink interface of the VRF gateway, which enables traffic segmentation and routing within the VRF context.
The uplink trunk segment is where multiple VLANs can be configured and tagged, allowing them to be used by the VRF Lite setup for routing and segmentation across the network.

Question: 24

Which two logical router components span across all transport nodes? (Choose two.)

A. SERVICE_ROUTER_TIER0
B. TIER0_DISTRIBUTED_ROUTER
C. DISTRIBUTED_ROUTER_TIER0
D. DISTRIBUTED_ROUTER_TIER1
E. SERVICE_ROUTER_TIER1

Answer: B, D

Explanation:

TIER0_DISTRIBUTED_ROUTER: The Tier-0 Distributed Router spans all transport nodes, providing distributed routing capabilities across the NSX environment at the Tier-0 level.
DISTRIBUTED_ROUTER_TIER1: Similarly, the Tier-1 Distributed Router spans all transport nodes, enabling distributed routing at the Tier-1 level, which allows routing functions to occur closer to the workload VMs across the transport nodes.

Question: 25

What must be configured on Transport Nodes for encapsulation and decapsulation of Geneve protocol?

A. TEP
B. STT
C. VXLAN
D. UDP

Answer: A

Explanation:

TEP (Tunnel Endpoint): TEPs (Tunnel Endpoints) are configured on transport nodes to handle the encapsulation and decapsulation of the Geneve protocol. TEPs are responsible for creating the overlay network by encapsulating traffic in the Geneve protocol when it moves between transport nodes and decapsulating it upon arrival.

Question: 26

A customer is preparing to deploy a VMware Kubernetes solution in an NSX environment. What is the minimum MTU size for the UPLINK profile?

A. 1700
B. 1500
C. 1550
D. 1650

Answer: A

Explanation:

For a VMware Kubernetes deployment in an NSX environment, the minimum recommended MTU size for the UPLINK profile is 1700. This allows sufficient space for the additional overhead introduced by encapsulation protocols, such as Geneve, used in NSX-T Data Center, ensuring optimal performance and avoiding fragmentation.

Question: 27

What are three NSX Manager roles? (Choose three.)

A. master
B. manager
C. controller

D. cloud
E. policy
F. zookeeper

Answer: A, C, F

Explanation:

master: The master role in NSX Manager is responsible for managing and coordinating the other NSX Manager nodes in the cluster.
policy: The policy role handles the policy-driven API and configuration, allowing administrators to define and manage network and security policies.
controller: The controller role in NSX Manager manages control plane functions and handles routing, switching, and other network state information required for NSX operations.

Question: 28

Which two CLI commands could be used to see if vmnic link status is down? (Choose two.)

A. esxcfg-nics -l
B. esxcli network nic list
C. esxcfg-vmknic -l
D. esxcfg-vmsvc/get.networks
E. esxcli network vswitch dvs vmware list

Answer: A, B

Explanation:

esxcfg-nics -l: This command lists all physical NICs on the ESXi host along with their link status, allowing you to check if any vmnic link status is down.
esxcli network nic list: This command provides a list of network interfaces with their details, including link status, making it useful for verifying if the link status of a vmnic is down.

Question: 29

Which VMware NSX Portfolio product can be described as a distributed analysis solution that provides visibility and dynamic security policy enforcement for NSX environments?

A. NSX Manager
B. NSX Distributed IDS/IPS
C. NSX Intelligence
D. NSX Cloud

Answer: C

Explanation:

NSX Intelligence is a distributed analytics solution within the VMware NSX Portfolio that provides visibility and dynamic security policy enforcement in NSX environments. It enables detailed traffic analysis, identifies security threats, and helps in the automated creation and enforcement of security policies based on observed network traffic patterns and behaviors.

Question: 30

An administrator has a requirement to have consistent policy configuration and enforcement across NSX instances.
What feature of NSX fulfills this requirement?

A. Multi-hvpervisor support
B. Federation
C. Load balancer
D. Policy-driven configuration

Answer: B

Explanation:

NSX Federation allows consistent policy configuration and enforcement across multiple NSX instances or environments. It provides a unified framework to manage security and networking policies across different NSX deployments, enabling centralized control and consistent application of policies across multiple sites or data centers.

Question: 31

When collecting support bundles through NSX Manager, which files should be excluded for potentially containing sensitive information?

A. Core Files
B. Controller Files
C. Audit Files
D. Management Files

Answer: A

Explanation:

Core Files should be excluded when collecting support bundles through NSX Manager because they may contain sensitive information, such as memory dumps that could reveal sensitive data from processes at the time of an issue. Excluding core files helps ensure that potentially sensitive data is not unintentionally

shared.

Question: 32

What can the administrator use to identify overlay segments in an NSX environment if troubleshooting is required?

A. Geneve ID
B. VMI ID
C. Segment ID
D. VLANID

Answer: B

Explanation:

In an NSX environment, each overlay segment is uniquely identified by a VNI ID (Virtual Network Identifier). The VNI is used to distinguish different overlay networks within the NSX environment and is essential for troubleshooting, as it helps administrators identify specific segments where traffic is encapsulated and isolated.

Question: 33

How does the Traceflow tool identify issues in a network?

A. Compares intended network state in the control plane with Tunnel End Point (TEP) keepalives in the data plane.
B. Compares the management plane configuration states containing control plane traffic and error reporting from transport node agents.
C. Injects ICMP traffic into the data plane and observes the results in the control plane.
D. Injects synthetic traffic into the data plane and observes the results in the control plane.

Answer: D

Explanation:

The Traceflow tool in NSX injects synthetic traffic into the data plane and monitors the traffic flow through the network, allowing administrators to observe how the traffic is handled at each hop. This approach helps identify issues such as dropped packets, routing errors, or misconfigurations by providing visibility into the path taken by the traffic and any potential disruptions.

Question: 34

Where is the insertion point for East-West network introspection?

A. Tier-0 router

B. Guest VM vNIC
C. Partner SVM
D. Host Physical NIC

Answer: B

Explanation:

The insertion point for East-West network introspection in NSX is at the Guest VM vNIC (virtual Network Interface Card). By inspecting traffic at the vNIC level, NSX can monitor and apply security policies to traffic between virtual machines (East-West traffic) within the same network segment or data center, providing detailed security controls for VM-to-VM communication.

Question: 35

Which is the only supported mode in NSX Global Manager when using Federation?

A. Proxy
B. Policy
C. Controller
D. Proton

Answer: B

Explanation:

When using NSX Federation, Policy mode is the only supported mode in NSX Global Manager. This

mode allows centralized management and consistent policy enforcement across multiple NSX environments, providing a unified approach to managing network and security policies in federated deployments.

Question: 36

When a stateful service is enabled for the first time on a Tier-0 Gateway, what happens on the NSX Edge node?

A. DR is instantiated and automatically connected with SR.
B. SR is instantiated and automatically connected with DR.
C. SR and DR doesn't need to be connected to provide any stateful services.
D. SR and DR is instantiated but requires manual connection.

Answer: B

Explanation:

When a stateful service (such as NAT or firewall) is enabled for the first time on a Tier-0 Gateway, the Service Router (SR) is instantiated on the NSX Edge node and automatically connected with the Distributed Router (DR). This connection enables the Tier-0 Gateway to handle stateful services by routing traffic through the SR,

which manages stateful packet processing, while the DR provides distributed routing functionality.

Question: 37

An NSX administrator is creating a Tier-1 Gateway configured in Active-Standby High Availability Mode. In the event of node failure, the failover policy should not allow the original failed node to become the Active node upon recovery.
Which failover policy meets this requirement?

A. Enable Preemptive
B. Non-Preemptive
C. Preemptive
D. Disable Preemptive

Answer: B

Explanation:

In Non-Preemptive failover policy, once a failover occurs and a new Active node is designated, the original failed node will not automatically become the Active node upon recovery. This setting ensures that the failover does not revert to the original node after it comes back online, maintaining the stability of the network by keeping the current Active node as is.

Question: 38

Which CLI command is used for packet capture on the ESXi Node?

A. tcpdump
B. set capture
C. pktcap-uw
D. debug

Answer: C

Explanation:

The pktcap-uw command is specifically used on ESXi hosts for packet capture. It provides a detailed packet capture utility that allows administrators to capture traffic at various points on the ESXi host, such as virtual switches, uplinks, and VMkernel interfaces, making it a powerful tool for network troubleshooting on ESXi nodes.

Question: 39

Which command on ESXi is used to verify the Local Control Plane connectivity with Central Control Plane?

A. esxcli network ip connection list | grep netcpa
B. esxcli network ip connection list | grep ccpd

C. esxcli network ip connection list | grep 1234
D. esxcli network ip connection list | grep 1235

Answer: A

Explanation:

The netcpa process is responsible for Local Control Plane (LCP) connectivity with the Central Control Plane (CCP) in NSX. Using the command esxcli network ip connection list | grep netcpa, administrators can verify the connectivity status between the LCP on the ESXi host and the CCP, ensuring proper communication for NSX operations.

Question: 40

Which two statements describe the characteristics of an Edge Cluster in NSX? (Choose two.)

A. Must have only active-active edge nodes
B. Can contain multiple types of edge nodes (VM or bare metal)
C. Must contain only one type of edge nodes (VM or bare metal)
D. Can have a maximum of 10 edge nodes
E. Can have a maximum of 8 edge nodes

Answer: B, E

Explanation:

An NSX Edge Cluster can contain a mix of edge node types, meaning it can have both virtual machine (VM) and bare-metal edge nodes within the same cluster.

An NSX Edge Cluster supports a maximum of 8 edge nodes, allowing for scalability while adhering to the NSX design limitations for edge clusters.

Question: 41

An NSX administrator is using ping to check connectivity between VM1 running on ESXi1 to VM2 running on ESXi2. The ping tests fail. The administrator knows the maximum transmission unit size on the physical switch is 1600.
Which command does the administrator use to check the VMware kernel ports for tunnel end point communication?

A. vmkping ++netstack=geneve -d -s 1572 <destination IP address>
B. vmkping ++netstack=vxlan -d -s 1572 <destination IP address>
C. esxcli network diag ping –H <destination IP address>
D. esxcli network diag ping -I vmk0 -H <destination IP address>

Answer: A

Explanation:

The vmkping ++netstack=geneve -d -s 1572 <destination IP address> command is used to check connectivity for VMware kernel ports specifically for Geneve tunnel endpoints (TEPs). The -s 1572 option sets the packet size to test within the 1600 MTU limit, accounting for the Geneve encapsulation overhead. The -d option enables the "Don't Fragment" bit, ensuring the packet isn't fragmented along the path, which is essential for verifying MTU consistency across the network.

Question: 42

An NSX administrator is creating a NAT rule on a Tier-0 Gateway configured in active-standby high availability mode.
Which two NAT rule types are supported for this configuration? (Choose two.)

A. Port NAT
B. 1:1 NAT
C. Destination NAT
D. Reflexive NAT
E. Source NAT

Answer: C

Explanation:

In an NSX environment with a Tier-0 Gateway configured in active-standby high availability mode, Destination NAT (DNAT) and Source NAT (SNAT) are supported NAT rule types. These allow for traffic redirection by modifying the destination or source IP addresses as needed, which is commonly used in configurations involving external access and internal IP address translation.

Question: 43

When deploying an NSX Edge Transport Node, what two valid IP address assignment options should be specified for the TEP IP addresses? (Choose two.)

A. Use an IP Pool
B. Use RADIUS
C. Use a Static IP List
D. Use BootP
E. Use a DHCP Server

Answer: A, E

Explanation:

IP Pool: This allows you to define a range of IP addresses within NSX that the TEPs can use. DHCP Server: This enables the TEPs to automatically obtain IP addresses from a DHCP server configured in the network.

Question: 44

DRAG DROP
Match the NSX Intelligence recommendations with their correct purpose.

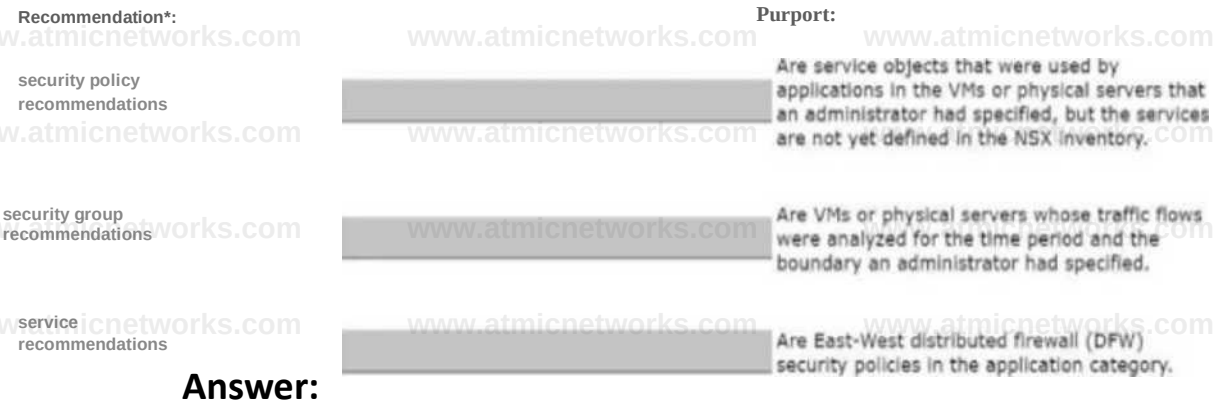

Answer:

Explanation:

Security policy recommendations: Are East-West distributed firewall (DFW) security policies in the application category12.

Security group recommendations: Are VMs or physical servers whose traffic flows were analyzed for the time period and the boundary you had specified12.

Service recommendations: Are service objects that were used by applications in the VMs or physical servers that you had specified, but the services are not yet defined in the NSX inventory12.

https://docs.vmware.com/en/VMware-NSX-Intelligence/4.1/user-guide/GUID-BA3B0D67-4AA8-439E-A845-4598DAD6B9D0.html

Question: 45

Which two built-in VMware tools will help identify the cause of packet loss on VLAN Segments? (Choose two.)
Which two built-in VMware tools will help identify the cause of packet loss on VLAN Segments? (Choose two.)

A. Flow Monitoring
B. Traceflow
C. Live Flow
D. Packet Capture
E. Activity Monitoring

Answer: B, D

Explanation:

Traceflow: This tool helps in troubleshooting network issues by injecting synthetic packets into the network and observing their path. It allows administrators to trace the packet flow across various network segments, making it easier to identify points of packet loss.
Packet Capture: This tool enables detailed inspection of traffic by capturing packets at specific points in the network. It allows administrators to analyze packet headers and payloads to determine if packet loss is occurring and to identify possible causes.

Question: 46

A customer has a network where BGP has been enabled and the BGP neighbor is configured on the Tier-0 Gateway. An NSX administrator used the get gateways command to retrieve this information:

```
sa-nsxedge-01> get gateways

Logical Router

UUID                                   VRF    GW_ID     Name                 Type
Ports

736a80e3-23f6-5a2d-81d6-bbefb2786666   0      0                              TUNNEL

B10ef54e-d5f3-49e5-99b7-8a51366d0S92   1      1C25      SR-T1-LR-01          SERVICE_ROUTER_TIER1

5a5ddd63-3764-4d28-b82e-ee4c964a0dfd   3      2049      SR-T0-LR-01          SERVICE ROUTER TTERC

0E0784db-511f-fa72-ae0b-lccaa0262ad2   4      7         DR-T0-LR-01          DISTRIBUTED_ROUTER_TIERO
```

Which two commands must be executed to check BGP neighbor status? (Choose two.)

A. vrf 3
B. sa-nsxedge-01(tier0_dr)> get bgp neighbor
C. vrf 1
D. sa-nsxedge-01(tier1_sr)> get bgp neighbor
E. sa-nsxedge-01(tier0_sr)> get bgp neighbor
F. vrf 4

Answer: A, E

Explanation:

vrf 3: The VRF ID for the Tier-0 Service Router (SR) is 3, as indicated in the output. To check the BGP neighbor status, you need to enter the correct VRF context.

sa-nsxedge-01(tier0_sr)> get bgp neighbor: This command retrieves the BGP neighbor status on the Tier-0 Service Router, which is where BGP neighbors are configured in NSX environments.

Question: 47

Which two of the following will be used for ingress traffic on the Edge node supporting a Single Tier topology? (Choose two.)

A. Tier-1 SR Router Port
B. Tier-0 Uplink interface
C. Downlink Interface for the Tier-0 DR
D. Downlink Interface for the Tier-1 DR
E. Inter-Tier interface on the Tier-0 gateway

Answer: A, B

Explanation:

Tier-1 SR Router Port: This port is used for ingress traffic on the Tier-1 Service Router (SR), which handles traffic as it enters the Tier-1 gateway.

Tier-1 SR Router Port: This port is used for ingress traffic on the Tier-1 Service Router (SR), which handles traffic as it enters the Tier-1 gateway.

Question: 48

Which TraceFlow traffic type should an NSX administrator use for validating connectivity between App and DB virtual machines that reside on different segments?

A. Anycast
B. Multicast
C. Broadcast
D. Unicast

Answer: B

Explanation:

In NSX, Unicast traffic type should be used in TraceFlow when validating connectivity between two specific virtual machines, such as App and DB VMs, that reside on different segments. Unicast traffic is directed from one source to a single

destination, making it suitable for testing direct connectivity between two VMs.

Question: 49

Which two tools are used for centralized logging in VMware NSX? (Choose two.)

A. Sysloq Server
B. VMware Aria Automation
C. VMware Aria Operations for Logs
D. VMware Aria Operations for Networks
E. VMware Aria Operations

Answer: A, C

Explanation:

Syslog Server: NSX supports forwarding logs to a centralized syslog server, which is a standard tool for centralized logging in network environments.

VMware Aria Operations for Logs (formerly known as vRealize Log Insight): This tool provides centralized log management and analytics, specifically designed to integrate with VMware environments, including NSX, for enhanced log collection, analysis, and troubleshooting.

Question: 50

An administrator is configuring service insertion for Network Introspection.
Which two places can the Network Introspection be configured? (Choose two.)

A. Edge Node
B. Host pNIC
C. Tier-0 gateway
D. Tier-1 gateway
E. Partner SVM

Answer: D, E

Explanation:

Tier-1 gateway: Network introspection services can be configured at the Tier-1 gateway level to inspect and control East-West traffic between workloads.

Partner SVM (Service Virtual Machine): Network introspection is often implemented through integration with a Partner SVM, which is a virtual machine provided by a third-party security partner to perform deep packet inspection and other security functions.

Question: 51

Where can an administrator see a visual overview of network connections between different VMs and different networks, within the NSX domain?

A. Network Intelligence
B. NSX Intelligence
C. VMware Aria Operations
D. VMware Aria Operations for Networks

Answer: B

Explanation:

NSX Intelligence provides a visual overview of network connections within the NSX domain, allowing administrators to see the traffic flows between different VMs and networks. It offers detailed visibility into network traffic patterns, application dependencies, and security posture, making it a valuable tool for monitoring and troubleshooting within NSX environments.

Question: 52

An NSX administrator is troubleshooting a connectivity issue with virtual machines running on an ESXi transport node. Which feature in the NSX UI shows the mapping between the virtual NIC and the host's physical adapter?

A. Port Mirroring
B. Activity Monitoring
C. IPF1X
D. Switch Visualization

Answer: D

Explanation:

Switch Visualization in the NSX UI provides a clear mapping between virtual NICs (vNICs) and the physical adapters on the host. This feature allows administrators to see how virtual network interfaces connect to the underlying physical network infrastructure, which is essential for troubleshooting connectivity issues on transport nodes.

Question: 53

What are the four types of role-based access control (RBAC) permissions? (Choose four.)

A. Auditor
B. Full access
C. Enterprise Admin
D. None
E. Execute
F. Read
G. Network Admin

Answer: A, B, D, F

Explanation:

Auditor: Allows users to view settings and logs without making changes.
Full access: Provides complete control over all NSX settings and configurations.
None: No permissions are granted, restricting access completely.
Read: Allows users to view configurations and settings without editing capabilities.

Question: 54

Which CLI command does an NSX administrator run on the NSX Manager to generate support bundle logs if the NSX UI is inaccessible?

A. esxcli system syslog config logger set --id=nsxmanager
B. get support-bundle file vcpnv.tgz
C. vm-support
D. set support-bundle file vcpnv.tgz

Answer: B

Explanation:

When the NSX UI is inaccessible, an NSX administrator can use the get support-bundle file vcpnv.tgz command in the CLI on the NSX Manager to generate a support bundle. This command creates a compressed file (vcpnv.tgz) containing logs and diagnostic information needed for troubleshooting.

Question: 55

Which NSX CLI command is used to change the authentication policy for local users?

A. set hardening-policy
B. get auth-policy minimum-password-length
C. set cli-timeout
D. set auth-policy

Answer: D

Explanation:

The set auth-policy command in the NSX CLI is used to configure the authentication policy for local users. This command allows administrators to adjust settings related to password policies, lockout policies, and other authentication-related parameters for local user accounts on NSX Manager.

Question: 56

When configuring OSPF on Tier-0 Gateway, which three of the following must match in order to establish a neighbor relationship with an upstream router? (Choose three.)

A. Area ID
B. MTU of the Uplink
C Naming convention
D. Address of the neighbor
E. Subnet mask
F. Protocol and Port

Answer: A, B, E

Explanation:

Area ID: Both routers must belong to the same OSPF area for a neighbor relationship to form.
MTU of the Uplink: Mismatched MTU settings can prevent the OSPF adjacency from forming, as OSPF packets may be dropped if they exceed the MTU size.
Subnet mask: Both routers must have the same subnet mask on the interface where OSPF is configured to establish a neighbor relationship.

Question: 57

Which command is used to set the NSX Manager's logging-level to debug mode for troubleshooting?

A. set service manager log-level debug
B. sec service nsx-manager logging-level debug
C. sec service nsx-manager log-level debug
D. sec service manager logging-level debug

Answer: D

Explanation:

The set service nsx-manager log-level debug command is used to set the NSX Manager's logging level to debug mode. Setting the log level to debug can provide more detailed logging information, which is useful for troubleshooting issues within the NSX Manager.

Question: 58

Which CLI command shows syslog on NSX Manager?

A. (show log manager follow
B. gee log-file syslog
C. [get log-file auch.log
D. /var/log/syslog/syslog.log

Answer: A

Explanation:

The show log manager follow command is used on the NSX Manager to view the syslog in real-time. This command displays

ongoing log entries, allowing administrators to monitor syslog messages as they are generated, which is helpful for troubleshooting and real-time analysis.

Question: 59

Which two BGP configuration parameters can be configured in the VRF Lite gateways? (Choose two.)

A. Route Aggregation
B. Route Distribution
C. BGP Neighbors
D. Graceful Restart
E. Local AS

Answer: C, E

Explanation:

BGP Neighbors: This parameter is essential for establishing BGP sessions with other routers.

Configuring BGP neighbors allows VRF Lite gateways to exchange routing information with adjacent BGP-enabled devices.

Local AS: The Local Autonomous System (AS) number can be set for the VRF Lite gateway, which is necessary for BGP operations within a specific routing domain.

Question: 60

Which CLI command would an administrator use to allow syslog on an ESXi transport node when using the esxcli utility?

A. esxcli network firewall ruleset set -a -e false
B. esxcli network firewall ruleset set -r syslog -e false
C. esxcli network firewall ruleset -e syslog
D. esxcli network firewall ruleset set -r syslog -e true

Answer: D

Explanation:

The esxcli network firewall ruleset set -r syslog -e true command is used to enable the firewall ruleset for syslog on an ESXi host. Setting the -e flag to true allows syslog traffic through the ESXi firewall, enabling remote logging for syslog messages from the transport node.

Question: 61

Which troubleshooting step will resolve an error with code 1001 during the configuration of a time based firewall rule?

A. Restarting the NTPservice on the ESXi host.
B. Reconfiguring the ESXi host with a local NTP server.
C. Re-installing the NSX VIBs on the ESXi host.

D. Changing the time zone on the ESXi host.

Answer: A

Explanation:

An error with code 1001 during the configuration of a time-based firewall rule often indicates a time synchronization issue. Restarting the NTP service on the ESXi host can resolve this issue by ensuring that the host's time is synchronized correctly, which is essential for time-based rules to function accurately.

Question: 62

Which statement is true about an alarm in a Suppressed state?

A. An alarmcan be suppressed for a specific durationin hours.
B. An alarmcan be suppressed for a specific durationinseconds.
C. An alarmcan be suppressed for a specific durationindays.
D. An alarmcan be suppressed for a specific durationinminutes

Answer: A

Explanation:

In NSX and VMware environments, an alarm in a suppressed state can typically be set to remain suppressed for a specific duration measured in hours. This allows administrators to temporarily ignore the alarm for a set period while working on a resolution without continuous alerts.

Question: 63

What are four NSX built-in role-based access control (RBAC) roles? (Choose four.)

A. None

B. Read
C. Auditor
D. Full Access
E. Network Admin
F. Enterprise Admin
G. Operator

Answer: A, B, C, D

Explanation:

None: No permissions are granted, restricting the user's access entirely.
Read: Grants read-only access, allowing the user to view configurations and settings without making changes.
Auditor: Similar to Read, but typically includes access to audit logs and more detailed viewing permissions for compliance purposes.

Full Access: Grants complete control over all NSX configurations and settings, allowing unrestricted access.

Question: 64

An administrator has deployed 10 Edge Transport Nodes in their NSX Environment, but has forgotten to specify an NTP server during the deployment.
What is the efficient way to add an NTP server to all 10 Edge Transport Nodes?

A. Use a Node Profile
B. Use Transport Node Profile
C. Use the CLI on each Edge Node
D. Use a PowerCLI script

Answer: B

Explanation:

Using a Transport Node Profile allows the administrator to apply configuration changes, such as specifying an NTP server, across multiple Edge Transport Nodes efficiently. This method is scalable and avoids the need for manual configuration on each individual node. Once the Transport Node Profile is updated, the configuration can be pushed to all associated nodes.

Question: 65

Which table on an ESXi host is used to determine the location of a particular workload for a frameforwarding decision?

A. Routing Table
B. ARP Table
C. TEP Table
D. MAC Table

Answer: D

Explanation:

The MAC Table on an ESXi host is used to determine the location of a particular workload for frameforwarding decisions. This table maps MAC addresses to specific interfaces, enabling the ESXi host to forward frames to the correct destination based on the MAC address of the workload. This is crucial for efficient Layer 2 forwarding decisions within the host.

Question: 66

What is the most restrictive NSX built-in role which will allow a user to apply configuration changes on an NSX Edge?

A. Network Engineer

B. Cloud Service Administrator
C. NSX Administrator
D. Network Operator

Answer: A

Explanation:

The Network Engineer role in NSX is a built-in role that provides permissions to apply configuration changes on NSX components, including NSX Edge. It is the most restrictive role that still allows users to make changes, whereas roles like Network Operator are typically limited to read-only access.

Question: 67

Refer to the exhibit.

An administrator would like to change the private IP address of the NAT VM 172.16.101.11 to a public address of 80.80.80.1 as the packets leave the NAT-Segment network.

Which type of NAT solution should be implemented to achieve this?

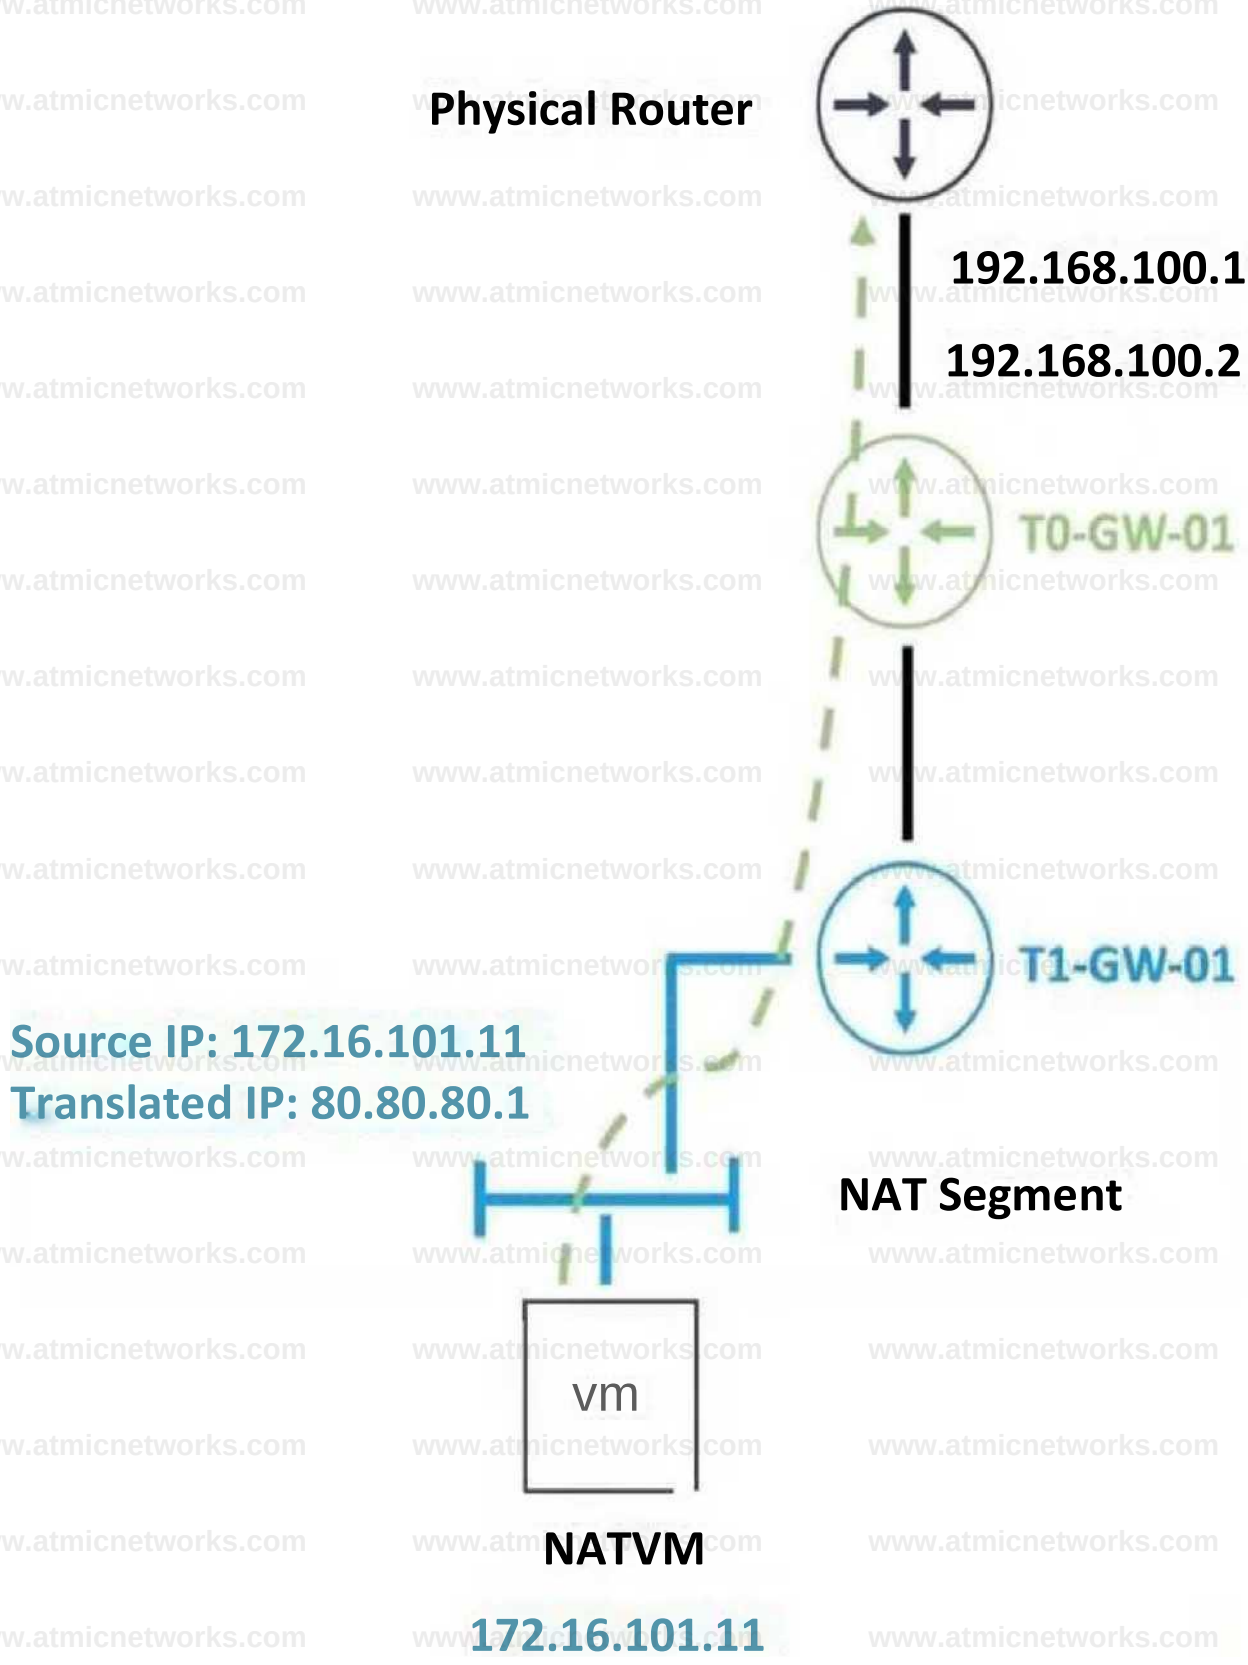

A. NAT64
B. Reflexive NAT
C. DNAT
D. SNAT

Answer: D

Explanation:

Source NAT (SNAT) is used to translate the private IP address (172.16.101.11) of the NAT VM to a public IP address (80.80.80.1) as the packets leave the NAT-Segment network. SNAT changes the source IP of outbound packets, allowing private IP addresses within the internal network to be mapped to public IP addresses for communication with external networks.

Question: 68

Which command is used to display the network configuration of the Tunnel Endpoint (TEP) IP on a bare metal transport node?

A. debug
B. tcpdump
C. tcpconfig
D. ifconfig

Answer: D

Explanation:

The ifconfig command is used to display the network configuration of interfaces, including the Tunnel Endpoint (TEP) IP on a bare metal transport node. This command provides details about IP addresses, subnet masks, and other network settings for each interface on the node.

Question: 69

Which VPN type must be configured before enabling an L2VPN?

A. Policy-based IPSec VPN
B. Port-based IPSec VPN
C. SSL-based IPSec VPN
D. Route-based IPSec VPN

Answer: D

Explanation:

Before enabling an L2VPN (Layer 2 VPN) in NSX, a Route-based IPSec VPN must be configured. Routebased VPNs create a secure tunnel over which Layer 2 traffic can be extended, allowing for the creation of L2VPN connections. This setup is required to establish the underlying secure connectivity that L2VPN relies on for traffic between sites.

Question: 70

DRAG DROP

Refer to the exhibits.

Drag and drop the NSX graphic element icons on the left found in an NSX Intelligence visualization graph to Its correct description on the right.

Answer Area

This icon represents a physical server that is part of your NSX environment A physical server can belong to more than one group

This icon represents a group on which security policies, Including East West firewall rules, can be applied A group can be a collection of VMs, physical servers, or sets of IP addresses

This is tire Icon lot the public IP addresses on the Internet. If at least one compute entity In your NSX environment communicated with a public IP address during the selected time period, that traffic flow Is included In the current visualisation.

This is the Icon used for a virtual machine (VM) that Is part of your NSX environment. A VM can belong to more than one group

Answer:

Explanation:

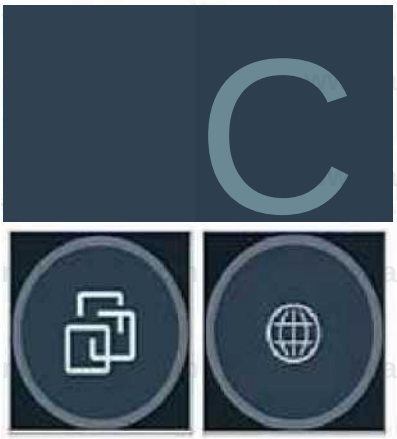

Answer Area

This icon represents a physical server that is part of your NSX environment. A physical server can belong to more than one group

This icon represents a group on which security policies, including East-West firewall rules, can be applied. A group can be a collection of VMs, physical servers, or sets of IP addresses

This is the icon for the public IP addresses on the Internet. If at least one compute entity in your NSX environment communicated with a public IP address during the selected time period, that traffic flow is included in the current visualization.

Hits is the Icon used for a virtual machine (VM) that is part of your NSX environment. A VM can belong to more than one group

https://docs.vmware.com/en/VMware-NSX-Intelligence/4.0/user-guide/GUID-DC78552B-2CC4-410D-A6C9-3FE0DCEE545B.html

Question: 71

Which three DHCP Services are supported by NSX? (Choose three.)

A. Gateway DHCP
B. Segment DHCP
C. DHCP Relay
D. Port DHCP per VNF
E. VRF DHCP Server

Answer: A, B, C

Explanation:

Gateway DHCP: NSX supports DHCP services configured on the gateway, allowing it to provide IP addresses to clients within the network.

Segment DHCP: NSX can provide DHCP services at the segment level, where DHCP is configured directly on a network segment to assign IP addresses to connected clients.

DHCP Relay: NSX supports DHCP Relay, which allows forwarding of DHCP requests to an external DHCP server for IP address assignment.

Question: 72

Which choice is a valid insertion point for North-South network introspection?

A. Host Physical NIC
B. Tier-0 gateway
C. Guest VM vNIC
D. Partner SVM

Answer: B

Explanation:

The Tier-0 gateway is the valid insertion point for North-South network introspection in an NSX environment. North-South traffic refers to traffic entering or leaving the data center, and the Tier-0 gateway is responsible for routing this traffic between the NSX environment and external networks. Placing network introspection at the Tier-0 gateway allows for inspection and security policies to be applied to traffic as it moves in and out of the data center.

Question: 73

Which of the following settings must be configured in an NSX environment before enabling stateful active-active SNAT?

A. Tier-1 gateway in active-standby mode
B. A Punting Traffic Group for the NSX Edge uplinks
C. An Interface Group for the NSX Edge uplinks
D. Tier-1 gateway in distributed only mode

Answer: B

Explanation:

In an NSX environment, a Punting Traffic Group for the NSX Edge uplinks must be configured before enabling stateful active-active Source NAT (SNAT). This configuration ensures that traffic is appropriately handled and forwarded between the NSX Edge nodes in an active-active setup, allowing stateful connections to be maintained across multiple Edge nodes.

Question: 74

An NSX administrator is reviewing syslog and notices that Distributed Firewall Rules hit counts are not being logged. What could cause this issue?

A. Zero Trust Security is not enabled.
B. Syslog is not configured on the NSX Manager.
C. Syslog is not configured on the ESXi transport node.
D. Distributed Firewall Rule logging is not enabled.

Answer: D

Explanation:

If Distributed Firewall Rule hit counts are not being logged, it is likely because Distributed Firewall Rule logging is not enabled. For hit counts to appear in the logs, logging must be explicitly enabled on each firewall rule where tracking is required. Without enabling logging at the rule level, no hit count information will be recorded in syslog.

Question: 75

How is the RouterLink port created between a Tier-1 Gateway and Tier-0 Gateway?

A. Automatically created when Tier-1 is created.
B. Manually create a Logical Switch and connect to bother Tier-1 and Tier-0 Gateways.
C. Manually create a Segment and connect to both Tier-1 and Tier-0 Gateways.
D. Automatically created when Tier-1 is connected with Tier-0 from NSX Ul.

Answer: D

Explanation:

The RouterLink port between a Tier-1 Gateway and a Tier-0 Gateway is automatically created when the Tier-1 Gateway is connected to the Tier-0 Gateway through the NSX UI. This link enables routing between the Tier-1 and Tier-0 gateways without the need for manual configuration of segments or logical switches.

Question: 76

Which two of the following parameters are required for deploying the NSX Application Platform? (Choose two.)

A. Interface Name
B. Upload XML File
C. Cluster Format Type
D. Interface Service Name
E. Upload Kubernetes Configuration File

Answer: B, E

Explanation:

Cluster Format Type: This parameter specifies the type of cluster format that will be used for the NSX Application Platform deployment.

Upload Kubernetes Configuration File: NSX Application Platform requires a Kubernetes environment, and the configuration file for Kubernetes must be uploaded to facilitate the deployment.

Question: 77

Which three of the following describe the Border Gateway Routing Protocol (BGP) configuration on a Tier-0 Gateway?

(Choose three.)

A. It supports a 4-byte autonomous system number.
B. Can be used as an Exterior Gateway Protocol.
C. The network is divided into areas that are logical groups.
D. EIGRP is disabled by default.
E. BGP is enabled by default.

Answer: A, B, E

Explanation:

It supports a 4-byte autonomous system number: BGP on a Tier-0 Gateway supports 4-byte AS (Autonomous System) numbers, which are necessary for larger routing domains.

Can be used as an Exterior Gateway Protocol: BGP is commonly used as an Exterior Gateway Protocol to establish routing between different autonomous systems (AS).

BGP is enabled by default: On a Tier-0 Gateway, BGP is typically enabled by default, allowing administrators to configure it for external routing.

Question: 78

Which CLI command on NSX Manager and NSX Edge is used to change NTP settings?

A. set timezone
B. set ntp-server
C. get timezone
D. get time-server

Answer: B

Explanation:

The set ntp-server command is used on NSX Manager and NSX Edge to configure the NTP (Network Time Protocol) settings. This command allows administrators to specify the NTP server, ensuring that the NSX components synchronize their time accurately with the designated time server.

Question: 79

What needs to be configured on a Tier-0 Gateway to make NSX Edge Services available to a VM on a VLAN-backed logical switch?

A. VLAN Uplink
B. Downlink interface
C. Loopback Router Port
D. Service interface

Answer: D

Explanation:

A Service interface on the Tier-0 Gateway is required to make NSX Edge Services, such as NAT or load balancing, available to a VM on a VLAN-backed logical switch. The Service interface allows the Tier-0 Gateway to connect directly to the VLAN-backed network, enabling Edge Services to interact with VMs on that network.

Question: 80

The security administrator turns on logging for a firewall rule.
Where is the log stored on an ESXi transport node?

A. /var/log/messages.log
B. /var/log/vmware/nsx/firewall.log
C. /var/log/fw.log
D. /var/log/dfwpktlogs.log

Answer: D

Explanation:

When logging is enabled for a firewall rule in NSX, the logs are stored on the ESXi transport node in the /var/log/vmware/nsx/firewall.log file. This file contains information about firewall rule hits and is useful for monitoring and troubleshooting firewall activity on the transport node.

Question: 81

An administrator needs to download the support bundle for NSX Manager.
Where does the administrator download the log bundle from?

A. System > Support Bundle
B. System > Settings
C. System > Utilities > Tools
D. System > Settings > Support Bundle

Answer: A

Explanation:

To download the support bundle for NSX Manager, an administrator navigates to System > Support

Bundle in the NSX Manager UI. This section provides options to generate and download the log bundle, which contains diagnostic information useful for troubleshooting and support.

Question: 82

Which two statements are true for IPSec VPN? (Choose two.)

A. IPSec VPN services can be configured at Tier-0 and Tier-1 gateways.
B. Dynamic routing is supported for any IPSec mode in NSX.
C. IPSec VPNs use the DPDK accelerated performance library.
D. VPNs can be configured on the command line interface on the NSX manager.

Answer: A, C

Explanation:

IPSec VPN services can be configured at Tier-0 and Tier-1 gateways: In NSX, IPSec VPN services can be applied to both Tier-0 and Tier-1 gateways, allowing secure site-to-site connections from these gateway levels.
IPSec VPNs use the DPDK accelerated performance library: NSX leverages the Data Plane Development Kit (DPDK) for optimized performance, which accelerates packet processing for IPSec VPNs and improves throughput.

Question: 83

An NSX administrator noticed that the nsxcli command times out after 600 secs of idle time.
Which CLI command disables the nsxcli time out value on NSX Manager?

A. set cli-timeout 1
B. set cli-timeout enabled
C. set cli-timeout disabled
D. set cli-timeout 0

Answer: D

Explanation:

Setting the cli-timeout value to 0 disables the CLI timeout on NSX Manager, preventing the nsxcli session from timing out due to inactivity. This ensures that the session remains active indefinitely until manually closed.

Question: 84

Which tool could be used to configure BGP on a Tier-0 Gateway?

A. ESX CLI
B. NSX CLI
C. API
D. iPerf3

Answer: B, C

Explanation:

Question: 85

Which of the following statements is true regarding the use of a Dynamic Routing Protocol on a Tier- 1 Gateway?

A. Both BGP and OSPF can be used on a Tier-1 Gateway.
B. You can only use OSPF on the Tier-1 Gateway
C. A Dynamic Routing Protocol cannot be used on a Tier-1 Gateway.
D. You can only use BGP on the Tier-1 Gateway.

Answer: D

Explanation:

In NSX, BGP is the only supported dynamic routing protocol on a Tier-1 Gateway. OSPF is not supported at the Tier-1 level; it is only available on Tier-0 Gateways. This limitation means that for dynamic routing on a Tier-1 Gateway, administrators can configure BGP to exchange routing information with connected Tier-0 Gateways.

Question: 86

Which two choices are solutions offered by the VMware NSX portfolio? (Choose two.)

A. VMware Tanzu Kubernetes Grid
B. VMware Tanzu Kubernetes Cluster
C. VMware NSX Advanced Load Balancer
D. VMware NSX Distributed IDS/IPS
E. VMware Aria Automation

Answer: C, D

Explanation:

VMware NSX is a portfolio of networking and security solutions that enables consistent policy, operations, and automation across multiple cloud environments1

The VMware NSX portfolio includes the following solutions:

VMware NSX Data Center: A platform for data center network virtualization and security that delivers a complete L2-L7 networking stack and overlay services for any workload1

VMware NSX Cloud: A service that extends consistent networking and security to public clouds such as AWS and Azure1

VMware NSX Advanced Load Balancer: A solution that provides load balancing, web application firewall, analytics, and monitoring for applications across any cloud12

VMware NSX Distributed IDS/IPS: A feature that provides distributed intrusion detection and prevention for workloads across any cloud12

VMware NSX Intelligence: A service that provides planning, observability, and intelligence for network and micro-segmentation1

VMware NSX Federation: A capability that enables multi-site networking and security management with consistent policy and operational state synchronization1

VMware NSX Service Mesh: A service that connects, secures, and monitors microservices across multiple clusters and clouds1

VMware NSX for Horizon: A solution that delivers secure desktops and applications across any device, location, or

network1

VMware NSX for vSphere: A solution that provides network agility and security for vSphere environments with a built-in console in vCenter1

VMware NSX-T Data Center: A platform for cloud-native applications that supports containers, Kubernetes, bare metal hosts, and multi-hypervisor environments1

VMware Tanzu Kubernetes Grid and VMware Tanzu Kubernetes Cluster are not part of the VMware NSX portfolio. They are solutions for running Kubernetes clusters on any cloud3

VMware Aria Automation is not a real product name. It is a fictional name that does not exist in the VMware portfolio.

https://blogs.vmware.com/networkvirtualization/2020/01/nsx-hero.html/

Question: 87

A company Is deploying NSX micro-segmentation in their vSphere environment to secure a simple application composed of web. app, and database tiers.

The naming convention will be:

- WKS-WEB-SRV-XXX
- WKY-APP-SRR-XXX
- WKI-DB-SRR-XXX

What is the optimal way to group them to enforce security policies from NSX?

A. Use Edge as a firewall between tiers.
B. Do a service insertion to accomplish the task.
C. Group all by means of tags membership.
D. Create an Ethernet based security policy.

Answer: C

Explanation:

The answer is C. Group all by means of tags membership.

Tags are metadata that can be applied to physical servers, virtual machines, logical ports, and logical segments in NSX. Tags can be used for dynamic security group membership, which allows for granular and flexible enforcement of security policies based on various criteria1

In the scenario, the company is deploying NSX micro-segmentation to secure a simple application composed of web, app, and database tiers. The naming convention will be:

WKS-WEB-SRV-XXX

WKY-APP-SRR-XXX

WKI-DB-SRR-XXX

The optimal way to group them to enforce security policies from NSX is to use tags membership. For example, the company can create three tags: Web, App, and DB, and assign them to the corresponding VMs based on their names. Then, the company can create three security groups: Web- SG, App-SG, and DB-SG, and use the tags as the membership criteria. Finally, the company can create and apply security policies to the security groups based on the desired rules and actions2 Using tags membership has several advantages over the other options:

It is more scalable and dynamic than using Edge as a firewall between tiers. Edge firewall is a centralized solution that can create bottlenecks and performance issues when handling large amounts of traffic3

It is more simple and efficient than doing a service insertion to accomplish the task. Service insertion is a feature that allows for integrating third-party services with NSX, such as antivirus or intrusion prevention systems. Service insertion is not

necessary for basic micro-segmentation and can introduce additional complexity and overhead.

It is more flexible and granular than creating an Ethernet based security policy. Ethernet based security policy is a type of policy that uses MAC addresses as the source or destination criteria. Ethernet based security policy is limited by the scope of layer 2 domains and does not support logical constructs such as segments or groups.

To learn more about tags membership and how to use it for micro-segmentation in NSX, you can refer to the following resources:

VMware NSX Documentation: Security Tag 1

VMware NSX Micro-segmentation Day 1: Chapter 4 - Security Policy Design 2

VMware NSX 4.x Professional: Security Groups

VMware NSX 4.x Professional: Security Policies

Question: 88

Which three NSX Edge components are used for North-South Malware Prevention? (Choose three.)

A. Thin Agent
B. RAPID
C. Security Hub
D. IDS/IPS
E. Security Analyzer
F. Reputation Service

Answer: BCD

Explanation:

https://docs.vmware.com/en/VMware-NSX/4.1/administration/GUID-69DF70C2-1769-4858-97E7-B757CAED08F0.html#:~:text=On%20the%20north%2Dsouth%20traffic,Guest%20Introspection%20(GI)%20platform.

The main components on the edge node for north-south malware prevention perform the following functions:

- IDS/IPS engine: Extracts files and relays events and data to the security hub

North-south malware prevention uses the file extraction features of the IDS/IPS engine that runs on NSX Edge for north-south traffic.

- Security hub: Collects file events, obtains verdicts for known files, sends files for local and cloud-based analysis, and sends information to the security analyzer
- RAPID: Provides local analysis of the file
- ASDS Cache: Caches reputation and verdicts of known files

Question: 89

What are two valid options when configuring the scope of a distributed firewall rule? (Choose two.)

A. DFW
B. Tier-1 Gateway
C. Segment
D. Segment Port
E. Group

Answer: AE

Explanation:

A group is a logical construct that represents a collection of objects in NSX, such as segments, segment ports, virtual machines, IP addresses, MAC addresses, tags, or security policies. A group can be used to define dynamic membership criteria based on various attributes or filters. A group can also be used as the scope of a distributed firewall rule, which means that the rule will apply to all the traffic that matches the group membership criteria32

Reference: https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/administration/GUID- 41CC06DF-1CD4-4233-B43E-492A9A3AD5F6.htmlhttps://docs.vmware.com/en/VMware-NSX-Data- Center-for-vSphere/6.4/com.vmware.nsx.admin.doc/GUID-D44C8923-992F-4695-B9C0- 5CC271679D09.html

Question: 90

Which two statements are true about IDS Signatures? (Choose two.)

A. Users can upload their own IDS signature definitions.
B. An IDS signature contains data used to identify known exploits and vulnerabilities.
C. An IDS signature contains data used to identify the creator of known exploits and vulnerabilities.
D. IDS signatures can be High Risk, Suspicious, Low Risk and Trustworthy.
E. An IDS signature contains a set of instructions that determine which traffic is analyzed.

Answer: BE

Explanation:

According to the Network Bachelor article1, an IDS signature contains data used to identify an attacker's attempt to exploit a known vulnerability in both the operating system and applications. This implies that statement B is true. According to the VMware NSX Documentation2, IDS/IPS Profiles are used to group signatures, which can then be applied to select applications and traffic. This

implies that statement E is true. Statement A is false because users cannot upload their own IDS signature definitions, they have to use the ones provided by VMware or Trustwave3. Statement C is false because an IDS signature does not contain data used to identify the creator of known exploits and vulnerabilities, only the exploits and vulnerabilities themselves. Statement D is false because IDS signatures are classified into one of the following severity categories: Critical, High, Medium, Low, or Informational1.
Reference: 3: Distributed IDS/IPS Settings and Signatures - VMware Docs 2: Distributed IDS/IPS - VMware Docs 1: NSX-T: Exploring Distributed IDS - Network Bachelor https://docs.vmware.com/en/VMware-SD-WAN/5.4/VMware-SD-WAN-Administration-Guide/GUID- 0BB81F8D-70EB-42D4-ABAF-F80C8F77A4CB.html

Question: 91

HOTSPOT
Refer to the exhibit.
An administrator configured NSX Advanced Load Balancer to load balance the production web server traffic, but the end users are unable to access the production website by using the VIP address. Which of the following Tier-1 gateway route advertisement settings needs to be enabled to resolve the problem? Mark the correct answer by clicking on the image.

Answer Area

Answer:

Explanation:

The correct answer is to enable the option All LB VIP Routes on the Tier-1 gateway route advertisement settings. This option allows the Tier-1 gateway to advertise the NSX Advanced Load Balancer LB VIP routes to the Tier-0 gateway and other peer routers, so that the end users can reach the production website by using the VIP address1. The other options are not relevant for this scenario.
To mark the correct answer by clicking on the image, you can click on the toggle switch next to All LB VIP Routes to turn it on. The switch should change from gray to blue, indicating that the option is enabled. See the image below for reference:

Question: 92

HOTSPOT
Refer to the exhibit.
Which two items must be configured to enable OSPF for the Tler-0 Gateway in the Image? Mark your answers

by clicking twice on the image.

Answer:

Explanation:

The correct answer is to enable the OSPF toggle and to add an Area Definition for the Tier-0 gateway in the image. These two items are required to configure OSPF on the Tier-0 gateway, as explained in the web search results123.

To mark your answers by clicking twice on the image, you can double-click on the toggle switch next to OSPF to turn it on. The switch should change from gray to blue, indicating that the option is enabled. Then, you can double-click on the Set button next to Area Definition to add an area definition. A pop-up window should appear where you can specify the area ID and type.

1. Click the OSPF toggle to enable OSPF 2. In the Area Definition field, click Set to add an area definition https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/administration/GUID-5BEC626C-5312-467D-B873-8E117349E9FC.html

Question: 93

Where in the NSX UI would an administrator set the time attribute for a time-based Gateway Firewall rule?

A. The option to set time-based rule is a clock Icon in the rule.
B. The option to set time based rule is a field in the rule Itself.
C. There Is no option in the NSX UI. It must be done via command line interface.
D. The option to set time-based rule is a clock Icon in the policy.

Answer: D

Explanation:

According to the VMware documentation1, the clock icon appears on the firewall policy section that you want to have a time window. By clicking the clock icon, you can create or select a time window that applies to all the rules in that policy section. The other options are incorrect because they either do not exist or are not related to the time-based rule feature. There is no option to set a time-based rule in the rule itself, as it is a policy-level setting. There is also an option to set a time-based rule in the NSX UI, so it does not require

using the command line interface.
https://docs.vmware.com/en/VMware-NSX/4.1/administration/GUID-8572496E-A60E-48C3-A016-4A081AC80BE7.html

Question: 94

Which three data collection sources are used by NSX Network Detection and Response to create correlations/Intrusion campaigns? (Choose three.)

A. Files and anti-malware (lie events from the NSX Edge nodes and the Security Analyzer
B. East-West anti-malware events from the ESXi hosts
C. Distributed Firewall flow data from the ESXi hosts
D. IDS/IPS events from the ESXi hosts and NSX Edge nodes
E. Suspicious Traffic Detection events from NSX Intelligence

Answer: ADE

Explanation:

The correct answers are A. Files and anti-malware (file) events from the NSX Edge nodes and the Security Analyzer, D. IDS/IPS events from the ESXi hosts and NSX Edge nodes, and E. Suspicious Traffic Detection events from NSX Intelligence. According to the VMware NSX Documentation3, these are the three data collection sources that are used by NSX Network Detection and Response to create correlations/intrusion campaigns.
The other options are incorrect or not supported by NSX Network Detection and Response. East-West anti-malware events from the ESXi hosts are not collected by NSX Network Detection and Response3. Distributed Firewall flow data from the ESXi hosts are not used for correlation/intrusion campaigns by NSX Network Detection and Response3.
https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/administration/GUID-14BBE50D-9931-4719-8FA7-884539C0D277.html

Question: 95

DRAG DROP

Sort the rule processing steps of the Distributed Firewall. Order responses from left to right.

If the packet matches source, destination, service, profile and applied to fields, apply the action defined	If the rule table action is allow, create an entry in the connection table and forward the packet	Packet arrives at dvfliter connection table, if matching entry In the table, process the packet.	|l the rule table action is reject or deny, take that action	If connection table has no match, compare the packet to the rule table

Answer:

Explanation:

The correct order of the rule processing steps of the Distributed Firewall is as follows: Packet arrives at vfilter connection table. If matching entry in the table, process the packet. If connection table has no match, compare the packet to the rule table.

If the packet matches source, destination, service, profile and applied to fields, apply the action defined.

If the rule table action is allow, create an entry in the connection table and forward the packet.

If the rule table action is reject or deny, take that action.

This order is based on the description of how the Distributed Firewall works in the web search results1. The first step is to check if there is an existing connection entry for the packet in the vfilter connection table, which is a cache of flow entries for rules with an allow action. If there is a match, the packet is processed according to the connection entry. If there is no match, the packet is compared to the rule table, which contains all the security policy rules. The rules are evaluated from top to bottom until a match is found. The match criteria include source, destination, service, profile and applied to fields. The action defined by the matching rule is applied to the packet. The action can be allow, reject or deny. If the action is allow, a new connection entry is created for the packet and the packet is forwarded to its destination. If the action is reject or deny, the packet is dropped and an ICMP message or a TCP reset message is sent back to the source.

Question: 96

Which two choices are use cases for Distributed Intrusion Detection? (Choose two.)

A. Use agentless antivirus with Guest Introspection.
B. Quarantine workloads based on vulnerabilities.
C. Identify risk and reputation of accessed websites.
D. Gain Insight about micro-segmentation traffic flows.
E. Identify security vulnerabilities in the workloads.

Answer: BE

Explanation:

According to the VMware NSX Documentation, these are two of the use cases for Distributed Intrusion Detection, which is a feature of NSX Network Detection and Response:

Quarantine workloads based on vulnerabilities: You can use Distributed Intrusion Detection to detect vulnerabilities in your workloads and apply quarantine actions to isolate them from the network until they are remediated.

Identify security vulnerabilities in the workloads: You can use Distributed Intrusion Detection to scan your workloads for known vulnerabilities and generate reports that show the severity, impact, and remediation steps for each vulnerability.

Question: 97

Which two of the following features are supported for the Standard NSX Application Platform Deployment? (Choose two.)

A. NSX Intrusion Detection and Prevention
B. NSX Intelligence
C. NSX Network Detection and Response
D. NSX Malware Prevention Metrics
E. NSX Intrinsic Security

Answer: CD

Explanation:

The NSX Application Platform Deployment features are divided into three form factors: Evaluation, Standard, and Advanced. Each form factor determines which NSX features can be activated or installed on the platform1. The Evaluation form factor supports only NSX Intelligence, which provides network visibility and analytics for NSX-T environments2. The Standard form factor supports both NSX Intelligence and NSX Network Detection and Response, which provides network threat detection and response capabilities for NSX-T environments3. The Advanced form factor supports all four features: NSX Intelligence, NSX Network Detection and Response, NSX Malware Prevention, and NSX Metrics1.

https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/nsx-application-platform/GUID- 85CD2728-8081-45CE-9A4A-D72F49779D6A.html

Question: 98

NSX improves the security of today's modern workloads by preventing lateral movement, which feature of NSX can be used to achieve this?

A. Network Segmentation
B. Virtual Security Zones
C. Edge Firewalling
D. Dynamic Routing

Answer: A

Explanation:

According to the web search results, network segmentation is a feature of NSX that improves the security of today’s modern workloads by preventing lateral movement. Lateral movement is a technique used by attackers to move from one compromised system to another within a network, exploiting vulnerabilities or credentials .

Network segmentation prevents lateral movement by dividing a network into smaller segments or zones, each with its own security policies and controls. This way, if one segment is compromised, the attacker cannot access other segments or resources . NSX enables network segmentation by using micro-segmentation, which applies granular firewall rules at the virtual machine level, regardless of the physical network topology .

Question: 99

Which of the following exist only on Tler-1 Gateway firewall configurations and not on Tier-0?

A. Applied To
B. Actions
C. Profiles
D. Sources

Answer: C

Explanation:

Question: 100

What are four NSX built-in rote-based access control (RBAC) roles? (Choose four.)

A. Network Admin
B. Enterprise Admin
C. Full Access
D. Read
E. LB Operator
F. None
G. Auditor

Answer: ABEG

Explanation:
https://docs.vmware.com/en/VMware-NSX/4.1/administration/GUID-26C44DE8-1854-4B06-B6DA-A2FD426CDF44.html

Question: 101

Which two are requirements for FQDN Analysis? (Choose two.)

A. The NSX Edge nodes require access to the Internet to download category and reputation definitions.
B. ESXi control panel requires access to the Internet to download category and reputation definitions.
C. The NSX Manager requires access to the Internet to download category and reputation definitions.
D. A layer 7 gateway firewall rule must be configured on the Tier-1 gateway uplink.
E. A layer 7 gateway firewall rule must be configured on the Tier-0 gateway uplink.

Answer: AD

Explanation:
https://docs.vmware.com/en/VMware-NSX/4.1/administration/GUID-C5CD87FD-8095-49F3-97CE-E606AB89162E.html?hWord=N4IghgNiBcIGYEcAmA7ABGFkCeBnAlriAL5A

Question: 102

Which of the two following characteristics about NAT64 are true? (Choose two.)

A. NAT64 is stateless and requires gateways to be deployed in active-standby mode.
B. NAT64 is supported on Tier-1 gateways only.
C. NAT64 is supported on Tier-0 and Tier-1 gateways.
D. NAT64 requires the Tier-1 gateway to be configured in active-standby mode.
E. NAT64 requires the Tier-1 gateway to be configured in active-active mode.

Answer: CD

Explanation:
https://docs.vmware.com/en/VMware-NSX/4.1/administration/GUID-69604E49-BC8B-4777-BFD8-B98F8D1FF064.html

Question: 103

What are two valid BGP Attributes that can be used to influence the route path traffic will take? (Choose two.)

A. AS-Path Prepend
B. BFD
C. Cost
D. MED

Answer: AD

Explanation:
AS-Path Prepend: This attribute allows you to prepend one or more AS numbers to the AS path of a route, making it appear longer and less preferable to other BGP routers. You can use this attribute to manipulate the inbound traffic from your BGP peers by advertising a longer AS path for some routes and a shorter AS path for others.
MED: This attribute stands for Multi-Exit Discriminator and allows you to specify a preference value for a route among multiple exit points from an AS. You can use this attribute to manipulate the outbound traffic to your BGP peers by advertising a lower MED value for some routes and a higher MED value for others.

Question: 104

HOTSPOT
Refer to the exhibit.

An administrator configured NSX Advanced Load Balancer to redistribute the traffic between the web servers.

However, requests are sent to only one server

Which of the following pool configuration settings needs to be adjusted to resolve the problem?

Mark the correct answer by clicking on the image.

Answer:

Explanation:

Load Balancing Algorithm

You specify the following parameters during the creation of a server pool:

- Name: A unique name for the server pool.
- Cloud: The cloud connector details for the NSX environment.
- VRF Context: Virtual Routing Framework (VRF) is a method to isolate traffic in a system. VRF is also called a route domain in the load balancer community. A global VRF context is created by default. Network administrators might create custom VRF contexts to isolate traffic between different tenants or subsets.
- Default Server Port: New connections to servers will use this destination service port. The default port is 80.
- Load-balancing algorithm: The selected load-balancing algorithm controls how the incoming connections are distributed among the servers in the pool.
- Tier-1 gateway (logical router): Specify the Tier-1 gateway that you want to attach the server pool to. This value matches the Tier-1 gateway specified for the virtual service and VIP.

Question: 105

Which two of the following are used to configure Distributed Firewall on VDS? (Choose two.)

A. vSphere API
B. NSX API
C. NSX CU
D. vCenter API
E. NSX UI

Answer: BE

Explanation:
According to the VMware NSX Documentation, these are two of the ways that you can use to configure Distributed Firewall on VDS:
NSX API: This is a RESTful API that allows you to programmatically configure and manage Distributed Firewall on VDS using HTTP methods and JSON payloads. You can use tools such as Postman or curl to send API requests to the NSX Manager node.
NSX UI: This is a graphical user interface that allows you to configure and manage Distributed Firewall on VDS using menus, tabs, buttons, and forms. You can access the NSX UI by logging in to the NSX Manager node using a web browser.
https://docs.vmware.com/en/VMware-NSX/4.1/administration/GUID-0DEF9F18-608D-4B5C-9175-5514750E901B.html

Question: 106

Which three security features are dependent on the NSX Application Platform? (Choose three.)

A. NSX Intelligence
B. NSX Firewall
C. NSX Network Detection and Response
D. NSX TLS Inspection
E. NSX Distributed IDS/IPS
F. NSX Malware Prevention

Answer: ACF

Explanation:
https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/nsx-application-platform/GUID-42EDE0AD-CD65-41AC-9694-AD0CCEC35969.html

Question: 107

Which NSX feature can be leveraged to achieve consistent policy configuration and simplicity across sites?

A. VRF Lite
B. Ethernet VPN
C. NSX MTML5 UI
D. NSX Federation

Answer: D

Explanation:

According to the VMware NSX Documentation, this is the NSX feature that can be leveraged to achieve consistent policy configuration and simplicity across sites:

NSX Federation: This feature allows you to create and manage a global network infrastructure that spans across multiple sites using a single pane of glass. You can use this feature to synchronize policies, segments, gateways, firewalls, VPNs, load balancers, and other network services across sites.

Question: 108

Which steps are required to activate Malware Prevention on the NSX Application Platform?

A. Select Cloud Region and Deploy Network Detection and Response.
B. Activate NSX Network Detection and Response and run Pre-checks.
C. Activate NSX Network Detection and Response and Deploy Malware Prevention.
D. Select Cloud Region and run Pre-checks.

Answer: D

Explanation:

To activate Malware Prevention on the NSX Application Platform, the steps are:

In the NSX Manager UI, select System and in the Configuration section, select NSX Application Platform.

Navigate to the Features section, locate the NSX Malware Prevention feature card, and click Activate or anywhere in the card.

In the NSX Malware Prevention activation window, select one of the available cloud regions from which you can access the NSX Advanced Threat Prevention cloud service.

Click Run Prechecks. This precheck process can take some time as the system validates that the minimum license requirement is met and that it is eligible for use with the NSX Advanced Threat Prevention cloud service. The system also validates that the selected cloud region is reachable. Click Activate. This step can take some time1. Therefore, the correct answer is D. The other options are incorrect because they involve activating or deploying NSX Network Detection and Response, which is a different feature from Malware Prevention.

Reference: Activate NSX Malware Prevention

Question: 109

Which field in a Tier-1 Gateway Firewall would be used to allow access for a collection of trustworthy web sites?

A. Source
B. Profiles -> Context Profiles
C. Destination
D. Profiles -> L7 Access Profile

Answer: D

Explanation:

The field in a Tier-1 Gateway Firewall that would be used to allow access for a collection of

trustworthy web sites is Profiles -> L7 Access Profile. This field allows the user to create a Layer 7 access profile that defines a list of allowed or blocked URLs based on categories, reputation, or custom entries1. The user can then apply the L7 access profile to a firewall rule to control the traffic based on the URL filtering criteria1. The other options are incorrect because they are not related to URL filtering. The Source field specifies the source IP address or group of the firewall rule1. The Destination field specifies the destination IP address or group of the firewall rule1. The Profiles -> Context Profiles field allows the user to create a context profile that defines a list of application signatures or attributes that can be used to identify and classify network traffic1. Reference: Gateway Firewall

Question: 110

When running nsxcli on an ESXi host, which command will show the Replication mode?

A. get logical-switch <Local-Switch-UUID> status
B. get logical-switch <Logical-Switch-UUID>
C. get logical-switches
D. get logical-switch status

Answer: C

Explanation:
https://vdc-download.vmware.com/vmwb-repository/dcr-public/c3fd9cef-6b2b-4772-93be-3fe60ce064a1/1f67b9e1-b111-4de7-9ea1-39931d28f560/NSX-T%20Command-Line%20Interface%20Reference.html#get%20logical-switch%20%3Clogical-switch-id%3E

Question: 111

Which two statements are correct about East-West Malware Prevention? (Choose two.)

A. A SVM is deployed on every ESXi host.
B. NSX Application Platform must have Internet access.
C. An agent must be installed on every ESXi host.
D. An agent must be installed on every NSX Edge node.
E. NSX Edge nodes must have Internet access.

Answer: AB

Explanation:

Reference: https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/administration/GUID- 0A8BF7D8-9C2E-48A5-8219-17C00F1EC13A.htmlhttps://www.wwt.com/blog/primer-series-napp- malware-prevention

Question: 112

A security administrator needs to configure a firewall rule based on the domain name of a specific application.
Which field in a distributed firewall rule does the administrator configure?

A. Profile
B. Service
C. Policy
D. Source

Answer: A

Explanation:
To configure a firewall rule based on the domain name of a specific application, the administrator needs to use the Profile field in a distributed firewall rule. The Profile field allows the administrator to select a context profile that contains one or more attributes for filtering traffic. One of the attributes that can be used is Domain (FQDN) Name, which specifies the fully qualified domain name of the application. For example, if the administrator wants to filter traffic to *.office365.com, they can create a context profile with the Domain (FQDN) Name attribute set to *.office365.com and use it in the Profile field of the firewall rule.
Reference:
Filtering Specific Domains (FQDN/URLs)
FQDN Filtering

Question: 113

What are two functions of the Service Engines in NSX Advanced Load Balancer? (Choose two.)

A. It collects real-time analytics from application traffic flows.
B. It stores the configuration and policies related to load-balancing services.
C. It performs application load-balancing operations.
D. It deploys web servers to perform load-balancing operations.
E. It provides a user interface to perform configuration and management tasks.

Answer: AC

Explanation:
Reference: https://docs.vmware.com/en/VMware-NSX-Advanced-Load-Balancer/22.1/Administration_Guide/GUID-84139C37-0129-40A7-A7AB-5A93E1F65B6D.html

Question: 114

An architect receives a request to apply distributed firewall in a customer environment without making changes to the network and vSphere environment. The architect decides to use Distributed Firewall on VDS. Which two of the following requirements must be met in the environment? (Choose two.)

A. vCenter 8.0 and later
B. NSX version must be 3.2 and later
C. NSX version must be 3.0 and later
D. VDS version 6.6.0 and later

Answer: BD

Explanation:

Distributed Firewall on VDS is a feature of NSX-T Data Center that allows users to install Distributed Security for vSphere Distributed Switch (VDS) without the need to deploy an NSX Virtual Distributed Switch (N-VDS). This feature provides NSX security capabilities such as Distributed Firewall (DFW), Distributed IDS/IPS, Identity Firewall, L7 App ID, FQDN Filtering, NSX Intelligence, and NSX Malware Prevention. To enable this feature, the following requirements must be met in the environment: The NSX version must be 3.2 and later1. This is the minimum version that supports Distributed Security for VDS.
The VDS version must be 6.6.0 and later1. This is the minimum version that supports the NSX host preparation operation that activates the DFW with the default rule set to allow.

Reference:
Overview of NSX IDS/IPS and NSX Malware Prevention

Question: 115

An NSX administrator would like to create an L2 segment with the following requirements:

- L2 domain should not exist on the physical switches.
- East/West communication must be maximized as much as possible.

Which type of segment must the administrator choose?

A. VLAN
B. Overlay
C. Bridge
D. Hybrid

Answer: B

Explanation:

An overlay segment is a layer 2 broadcast domain that is implemented as a logical construct in the NSX-T Data Center software. Overlay segments do not require any configuration on the physical switches, and they allow for optimal east/west communication between workloads on different ESXi hosts. Overlay segments use the Geneve protocol to encapsulate and decapsulate traffic between the hosts. Overlay segments are created and managed by the NSX Manager.
https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/administration/GUID-316E5027- E588-455C-88AD-A7DA930A4F0B.html