

2V0-41.23 Dumps

VMware NSX 4.x Professional

<https://www.certleader.com/2V0-41.23-dumps.html>



NEW QUESTION 1

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

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

Answer: D

NEW QUESTION 2

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

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

Answer: BC

Explanation:

According to the VMware NSX 4.x Professional documents and tutorials, IPsec VPN secures traffic flowing between two networks connected over a public network through IPsec gateways called endpoints. NSX Edge supports a policy-based or a route-based IPsec VPN. Beginning with NSX-T Data Center 2.5, IPsec VPN services are supported on both Tier-0 and Tier-1 gateways¹. NSX Edge also leverages the DPDK accelerated performance library to optimize the performance of IPsec VPN².

NEW QUESTION 3

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

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

Answer: B

Explanation:

Unicast is the traffic type that an NSX administrator should use for validating connectivity between App and DB virtual machines that reside on different segments. According to the VMware documentation¹, unicast traffic is the traffic type that is used to send a packet from one source to one destination. Unicast traffic is the most common type of traffic in a network, and it is used for applications such as web browsing, email, file transfer, and so on². To perform a traceflow with unicast traffic, the NSX administrator needs to specify the source and destination IP addresses, and optionally the protocol and related parameters¹. The traceflow will show the path of the packet across the network and any observations or errors along the way³. The other options are incorrect because they are not suitable for validating connectivity between two specific virtual machines. Multicast traffic is the traffic type that is used to send a packet from one source to multiple destinations simultaneously². Multicast traffic is used for applications such as video streaming, online gaming and group communication⁴. To perform a traceflow with multicast traffic, the NSX administrator needs to specify the source IP address and the destination multicast IP address¹. Broadcast traffic is the traffic type that is used to send a packet from one source to all devices on the same subnet². Broadcast traffic is used for applications such as ARP, DHCP, and network discovery. To perform a traceflow with broadcast traffic, the NSX administrator needs to specify the source IP address and the destination MAC address as FF:FF:FF:FF:FF:FF¹. Anycast traffic is not a valid option, as it is not supported by NSX Traceflow. Anycast traffic is a traffic type that is used to send a packet from one source to the nearest or best destination among a group of devices that share the same IP address. Anycast traffic is used for applications such as DNS, CDN, and load balancing.

NEW QUESTION 4

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.

NEW QUESTION 5

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.

NEW QUESTION 6

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

NEW QUESTION 7

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.

References:

- Filtering Specific Domains (FQDN/URLs)
- FQDN Filtering

NEW QUESTION 8

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

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

Answer: C

Explanation:

The insertion point for East-West network introspection is the Guest VM vNIC. Network introspection is a service insertion feature that allows third-party network services to be integrated with NSX. Network introspection enables traffic redirection from the Guest VM vNIC to a service virtual machine (SVM) that runs the partner service. The SVM can then inspect, monitor, or modify the traffic before sending it back to the original destination¹. The other options are incorrect because they are not the insertion points for East-West network introspection. The Tier-0 router is used for North-South routing and network services. The partner SVM is the service virtual machine that runs the partner service, not the insertion point. The host physical NIC is not involved in network introspection. References: Network Introspection Settings

NEW QUESTION 9

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

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

Answer: BD

Explanation:

According to the VMware NSX Documentation, these are the two places where you need to configure the VLANs used in VRF Lite:

- Uplink trunk segment: This is a segment that connects a tier-0 gateway to a physical network using multiple VLAN tags. You need to configure the VLAN IDs for each VRF on this segment.
- Uplink interface of the VRF gateway: This is an interface that connects a VRF gateway to an uplink trunk segment using a specific VLAN tag. You need to configure the VLAN ID for each VRF on this interface.

NEW QUESTION 10

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-9175-5514750E9>

NEW QUESTION 10

Which command is used to test management connectivity from a transport node to NSX Manager?

- A. `esxcli network ip connection list | grep 1234`
- B. `esxcli network connection list | grep 1235`
- C. `esxcli network ip connection list | grep 1235`
- D. `esxcli network connection list | grep 1234`

Answer: A

Explanation:

The NSX Manager management plane communicates with the transport nodes by using APH Server over NSX-RPC/TCP through port 1234. CCP communicates with the transport nodes by using APH Server over NSX-RPC/TCP through port 1235.

NEW QUESTION 14

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-nxedge-01> get gateways

Logical Router

UUID                               VRF    GW-ID    Name                Type
-----
736a80e3-23f6-5a2d-81d6-bbefb2786666 0       0        SR-T1-LR-01         SERVICE_ROUTER_TIER1
B10ef54e-d5f3-49e5-99b7-8a51366d0592 1       1025     SR-T0-LR-01         SERVICE_ROUTER_TIER0
5a5ddd63-3764-4d28-b82e-ee4c964a0d7d 3       2049     DR-T0-LR-01         DISTRIBUTED_ROUTER_TIER0
0E0784db-511f-fa72-ae0b-1ccaa0262ad2 4       7        DR-T0-LR-01         DISTRIBUTED_ROUTER_TIER0
```

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

- A. vrf 1
- B. vrf 4
- C. sa-nxedge-01(tier1_sr> get bgp neighbor
- D. sa-nxedge-01(tier0_sr> get bgp neighbor
- E. sa-nxedge-01(tier1_dr)> get bgp neighbor
- F. vrf 3

Answer: DF

Explanation:

BGP will be configured on the T0 SR. Connect to the VRF for the T0 SR and run get bgp neighbor once connected to it.

<https://docs.vmware.com/en/VMware-Validated-Design/5.1/sddc-deployment-of-vmware-nsx-t-workload-doma>

For the BGP configuration on NSX-T, the Tier-0 Service Router (SR) is typically where BGP is configured. To check the BGP neighbor status:

Connect to the VRF for the T0 SR, which is VRF 3 based on the provided output. Run the command to get BGP neighbor status once connected to it.

NEW QUESTION 16

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

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

Answer: BE

Explanation:

According to the VMware NSX Documentation¹, Packet Capture and Traceflow are two built-in VMware tools that can help identify the cause of packet loss on VLAN segments.

Packet Capture allows you to capture packets on a specific interface or segment and analyze them using tools such as Wireshark or tcpdump. Packet Capture can help you diagnose network issues such as misconfigured MTU, incorrect VLAN tags, or firewall drops.

Traceflow allows you to inject synthetic packets into the network and trace their path from source to destination. Traceflow can help you verify connectivity, routing, and firewall rules between virtual machines or segments. Traceflow can also show you where packets are dropped or modified along the way.

NEW QUESTION 18

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: CE

Explanation:

The Service Engines in NSX Advanced Load Balancer are VM-based applications that handle all data plane operations by receiving and executing instructions from the Controller. The Service Engines perform the following functions:

➤ They perform application load-balancing operations for all client- and server-facing network interactions. They support various load-balancing algorithms, health monitors, SSL termination, and persistence profiles.

➤ They provide a user interface to perform configuration and management tasks. The user interface is accessible through a web browser or a REST API. The user interface allows the user to create and modify virtual services, pools, health monitors, policies, analytics, and other load-balancing settings

<https://docs.vmware.com/en/VMware-Telco-Cloud-Platform/3.0/vmware-telco-cloud-reference-architecture-gui>

NEW QUESTION 23

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.

References:

- Overview of NSX IDS/IPS and NSX Malware Prevention

NEW QUESTION 26

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

- A. set vrf <ID>show logical-routers show <LR-D> bgp
- B. show logical-routers get vrfshow ip route bgp
- C. get gateways vrf <number>get bgp neighbor
- D. enable <LR-D> get vrf <ID>show bgp neighbor

Answer: C

Explanation:

The sequence of commands that could be used to check the BGP connection status between the Tier-O Gateway and the upstream physical router on NSX Edge node is get gateways, vrf <number>, get bgp neighbor. These commands can be executed on the NSX Edge node CLI after logging in as admin6. The first command, get gateways, displays the list of logical routers (gateways) configured on the Edge node, along with their IDs and VRF numbers⁷. The second command, vrf <number>, switches to the VRF context of the desired Tier-O Gateway, where <number> is the VRF number obtained from the previous command⁷. The third command, get bgp neighbor, displays the BGP neighbor summary for the selected VRF, including the neighbor IP address, AS number, state, uptime, and prefixes received⁸. The other options are incorrect because they either use invalid or incomplete commands or do not switch to the correct VRF context. References: NSX-T Command-Line Interface Reference, NSX Edge Node CLI Commands, Troubleshooting BGP on NSX-T Edge Nodes

NEW QUESTION 27

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.

NEW QUESTION 28

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

- A. VXIAN
- B. UDP
- C. STT
- D. TEP

Answer: D

Explanation:

According to the VMware NSX Documentation, TEP stands for Tunnel End Point and is a logical interface that must be configured on transport nodes for encapsulation and decapsulation of Geneve protocol. Geneve is a tunneling protocol that encapsulates the original packet with an outer header that contains metadata such as the virtual network identifier (VNI) and the transport node IP address. TEPs are responsible for adding and removing the Geneve header as the packet traverses the overlay network.

NEW QUESTION 33

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

Answer: A

Explanation:

According to the VMware NSX Documentation, a non-preemptive failover policy means that the original failed node will not become the active node upon recovery, unless the current active node fails again. This policy can help avoid unnecessary failovers and ensure stability.

The other options are either incorrect or not available for this configuration. Preemptive is the opposite of non-preemptive, meaning that the original failed node will become the active node upon recovery, if it has a higher priority than the current active node. Enable Preemptive and Disable Preemptive are not valid options for the failover policy, as the failover policy is a drop-down menu that only has two choices: Preemptive and Non-Preemptive.

NEW QUESTION 35

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

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

Answer: C

Explanation:

The log for a firewall rule on an ESXi transport node is stored in the /var/log/dfwpklogs.log file. This file contains information about the packets that match or do not match the firewall rules, such as the source and destination IP addresses, ports, protocols, actions, and rule IDs. The log file can be viewed using the esxcli network firewall get command or the vSphere Client.

<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/administration/GUID-D57429A1-A0A9-42BE-A>

NEW QUESTION 36

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-B98F8D1F>

NEW QUESTION 38

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

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

Answer: BD

Explanation:

According to the VMware NSX Documentation¹, you can configure BGP neighbors for VRF-Lite by specifying the neighbor IP address, remote AS number, source IP address, and route filter. You can also configure route distribution for VRF-Lite by selecting the route redistribution sources and the route map to apply.

<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/administration/GUID-4CB5796A-1CED-4F0E-A>

NEW QUESTION 39

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

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

Answer: B

Explanation:

Federation is a feature of NSX that allows the administrator to manage multiple NSX instances with a single pane of glass view, create gateways and segments that span one or more locations, and configure and enforce firewall rules consistently across locations¹. Federation provides centralized policy management for security and networking services for all locations and pushes it down to NSX Local Managers at the respective sites for enforcement¹. Federation also enables disaster recovery and workload mobility scenarios by providing consistent network and security policies across different sites¹. References: 1: NSX Federation - VMware Docs(<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/administration/GUID-D5B6DC79-6733-44>)

NEW QUESTION 41

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. Downlink Interface

- B. VLAN Uplink
- C. Loopback Router Port
- D. Service Interface

Answer: B

Explanation:

To make NSX Edge Services available to a VM on a VLAN-backed logical switch, you need to configure a VLAN Uplink on the Tier-0 Gateway. A VLAN Uplink is a logical interface that connects the Tier-0 Gateway to the physical network and provides external connectivity for the NSX Edge Services1. A VLAN Uplink can be configured on the NSX Manager UI by selecting Networking > Tier-0 Gateways > Interfaces > Set > Add Interface1.

<https://docs.vmware.com/en/VMware-NSX/4.1/administration/GUID-D641380B-4C8E-4C8A-AF64-4261A266>

NEW QUESTION 42

Which VPN type must be configured before enabling a L2VPN?

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

Answer: A

Explanation:

According to the VMware NSX Documentation, this VPN type must be configured before enabling a L2VPN. L2VPN stands for Layer 2 VPN and is a feature that allows you to extend your layer 2 network across different sites using an IPsec tunnel. Route-based IPsec VPN is a VPN type that uses logical router ports to establish IPsec tunnels between sites.

<https://docs.vmware.com/en/VMware-NSX/4.1/administration/GUID-86C8D6BB-F185-46DC-828C-1E1876B8>

NEW QUESTION 43

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

Answer: DE

Explanation:

According to the VMware NSX Documentation, these are the prerequisites for adding nodes to an NSX Management Cluster using the NSX UI:

- All nodes must be in the same subnet and have IP connectivity with each other.
- A compute manager must be configured and associated with the NSX Manager node.
- The NSX Manager node must have a valid license.
- The NSX Manager node must have a valid certificate.

NEW QUESTION 46

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

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

Answer: B

Explanation:

According to the VMware NSX Documentation, this is where you can download the support bundle for NSX Manager from the NSX UI:

- System > Support Bundle: This option allows you to download a support bundle that contains logs, configuration files, and diagnostic information from your NSX Manager node and cluster. You can use this option to troubleshoot issues or provide information to VMware support.

<https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-794C691E-B950-4838-9> <https://docs.vmware.com/en/VMware-NSX/4.1/administration/GUID-73D9AF0D-4000-4EF2-AC66-6572AD1>

NEW QUESTION 50

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

- A. SFRVICE_ROUTER_TJERO
- B. TIERO_DISTRI BUTE D_ ROUTER
- C. DISTRIBUTED_ROUTER_TIER1
- D. DISTRIBUTED_ROUTER_TIER0
- E. SERVICE_ROUTER_TIER1

Answer: CD

Explanation:

<https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-design.doc/GUID-74>

NEW QUESTION 51

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

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

Answer: A

Explanation:

The RouterLink port is automatically created when a Tier-1 Gateway is connected with a Tier-0 Gateway from the NSX UI1. The RouterLink port is a logical interface that is assigned an IP address and is associated with a physical or virtual interface. The RouterLink port acts as an end point of the IPsec tunnel and routes traffic between the Tier-1 Gateway and the Tier-0 Gateway2. The other options are incorrect because they involve manual creation of logical switches or segments, which are not required for RouterLink port creation. References: Configure NSX for Virtual Networking from vSphere Client, Virtual Private Network (VPN)

NEW QUESTION 53

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. References: Activate NSX Malware Prevention

NEW QUESTION 57

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

- A. An alarm can be suppressed for a specific duration in seconds.
- B. An alarm can be suppressed for a specific duration in days.
- C. An alarm can be suppressed for a specific duration in minutes.
- D. An alarm can be suppressed for a specific duration in hours.

Answer: D

Explanation:

The answer is D. An alarm can be suppressed for a specific duration in hours.

According to the VMware NSX documentation, an alarm can be in one of the following states: Open, Acknowledged, Suppressed, or Resolved12

An alarm in a Suppressed state means that the status reporting for this alarm has been disabled by the user for a user-specified duration12

When a user moves an alarm into a Suppressed state, they are prompted to specify the duration in hours. After the specified duration passes, the alarm state reverts to Open. However, if the system determines the condition has been corrected, the alarm state changes to Resolved13

To learn more about how to manage alarm states in NSX, you can refer to the following resources:

- VMware NSX Documentation: Managing Alarm States 1
- VMware NSX Documentation: View Alarm Information 2
- VMware NSX Intelligence Documentation: Manage NSX Intelligence Alarm States 3 <https://docs.vmware.com/en/VMware-NSX-Intelligence/1.2/user-guide/GUID-EBD3C5A8-F9AB-4A22-BA40->

NEW QUESTION 60

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 2V0-41.23 Exam with Our Prep Materials Via below:

<https://www.certleader.com/2V0-41.23-dumps.html>