VMware

Exam Questions 2V0-41.23

VMware NSX 4.x Professional

NEW QUESTION 1
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-88

NEW QUESTION 2
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: CD

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/

NEW QUESTION 3
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 Tler-0 and Tler-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 4
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. Tier-1 gateway in distributed only mode
C. An Interface Group for the NSX Edge uplinks

D. A Punting Traffic Group for the NSX Edge uplinks

Answer: C

Explanation:
To enable stateful active-active SNAT on a Tier-0 or Tier-1 gateway, you must configure an Interface Group for the NSX Edge uplinks. An Interface Group is a logical grouping of NSX Edge interfaces that belong to the same failure domain. A failure domain is a set of NSX Edge nodes that share the same physical network infrastructure and are subject to the same network failures. By configuring an Interface Group, you can ensure that the stateful services are distributed across different failure domains and can recover from network failures1

NEW QUESTION 5
Which CLI command on NSX Manager and NSX Edge is used to change NTP settings?

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

Answer: D

Explanation:
The CLI command on NSX Manager and NSX Edge that is used to change NTP settings is set ntp-server. Th command allows the user to configure one or more NTP servers for time synchronization12. The other options are incorrect because they are not valid CLI commands for changing NTP settings. The get timezone and timezone commands are used to display and configure the timezone of the system1. The get
time-server command is used to display the current time server configuration1. There are no CLI commands for using RADIUS or BootP for NTP settings.
References: NSX-T Command-Line Interface
Reference, vSphere ESXi 7.0 U3 and later versions NTP configuration steps

NEW QUESTION 6
Which TraceFlow traffic type should an NSX administrator use tor 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 7
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

NEW QUESTION 8
In an NSX environment, an administrator is observing low throughput and congestion between the Tier-O Gateway and the upstream physical routers.
Which two actions could address low throughput and congestion? (Choose two.)

A. Configure NAT on the Tier-0 gateway.
B. Configure ECMP on the Tier-0 gateway.

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

Answer: BC

Explanation:
ECMP (Equal Cost Multi-Path) is a routing protocol that increases the north and south communication bandwidth by adding an uplink to the tier-0 logical router and configure it for each Edge node in an NSX Edge cluster2. The ECMP routing paths are used to load balance traffic and provide fault tolerance for failed paths2. The tier-0 logical router must be in active-active mode for ECMP to be available2. A maximum of eight ECMP paths are supported2. Configuring ECMP on the tier-0 gateway can address low throughput and congestion by distributing the traffic among multiple paths and avoiding bottlenecks.
Deploying Large size Edge node/s can also address low throughput and congestion by providing more resources (memory, CPU, disk) for the Edge node to handle the network traffic. The NSX Edge VM system requirements vary depending on the appliance size, which affects the bandwidth, NAT/firewall, load balancer, and VPN capabilities of the Edge node1. A Large size Edge node has 32 GB memory, 8 vCPU, 200 GB disk space, and can support 2-10 Gbps bandwidth, L2-L4 features, and L7 load balancer1. An Extra Large size Edge node has 64 GB memory, 16 vCPU, 200 GB disk space, and can support more than 10 Gbps bandwidth, L2-L4 features, L7 load balancer, and VPN1. Deploying a larger size Edge node can improve the performance and capacity of the tier-0 gateway.
References: 2: Understanding ECMP Routing - VMware Docs(https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/administration/GUID-443B6B0D-F179-42 NSX Edge VM System Requirements - VMware Docs(https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.2/installation/GUID-22F87CA8-01A9-4F2E

NEW QUESTION 9
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

A. Mastered
B. Not Mastered

Answer: A

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:

NEW QUESTION 10
In which VPN type are the Virtual Tunnel interfaces (VTI) used?

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

Answer: B

Explanation:
Route-based VPN is a VPN type that uses Virtual Tunnel interfaces (VTI) to establish IPSec tunnels between an NSX Edge node and remote sites2. A VTI is a logical interface that is assigned an IP address and is associated with a physical or virtual interface. The VTI acts as an end point of the IPSec tunnel and routes traffic between the NSX Edge node and the remote site2. Route & SSL based VPNs, Policy & Route based VPNs, and SSL-based VPN are not VPN types that use VTI. References: Virtual Private Network (VPN)

NEW QUESTION 10
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.

NEW QUESTION 11
When configuring OSPF on a Tler-0 Gateway, which three of the following must match in order to establish a neighbor relationship with an upstream router? (Choose three.)

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

Answer: BCF

Explanation:
ccording to the VMware NSX Documentation, these are the three parameters that must match in order to establish an OSPF neighbor relationship with an upstream router on a tier-0 gateway:

- MTU of the Uplink: The maximum transmission unit (MTU) of the uplink interface must match the MTU of the upstream router interface. Otherwise, OSPF packets may be fragmented or dropped, causing neighbor adjacency issues.
- Subnet mask: The subnet mask of the uplink interface must match the subnet mask of the upstream router interface. Otherwise, OSPF packets may not reach the correct destination or be rejected by the upstream router.
- Area ID: The area ID of the uplink interface must match the area ID of the upstream router interface.

Otherwise, OSPF packets may be ignored or discarded by the upstream router.

NEW QUESTION 16
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.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
The correct answer is to enable the OSPF toggle and to add an Area Definition for the Tier-0 gateway in 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. The you can double-click on the Set button next to Area Definition to add an area definition. A pop-up windo 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-B

NEW QUESTION 19
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 24
Which command is used to display the network configuration of the Tunnel Endpoint (TEP) IP on a bare metal transport node?

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

Answer: B

Explanation:
The command ifconfig is used to display the network configuration of the Tunnel Endpoint (TEP) IP on a ba metal transport node2. The TEP IP is assigned to a network interface on the bare metal server that is used for overlay traffic. The ifconfig command can show the IP address, netmask, broadcast address, and other information of the network interface. For example, the following command shows the network configuration
of the TEP IP on a bare metal transport node with interface name ens192:
ifconfig ens192
The output of the command would look something like this:
ens192: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 10.10.10.10 netmask 255.255.255.1 broadcast 10.10.10.255 inet6
fe80::250:56ff:fe9a:1b8c prefixlen 64 scopeid 0x20<link> ether 00:50:56:9a:1b:8c txqueuelen 1000 (Ethernet) RX packets 123456 bytes 123456789 (123.4 MB)
RX errors 0
dropped 0 overruns 0 frame 0 TX packets 234567 bytes 234567890 (234.5 MB) TX errors 0 dropped 0
overruns 0 carrier 0 collisions 0
The TEP IP in this example is 10.10.10.10. References:
- IBM Cloud Docs

NEW QUESTION 25
Which table on an ESXi host is used to determine the location of a particular workload for a frame-forwarding decision?

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

Answer: B

Explanation:
The MAC table on an ESXi host is used to determine the location of a particular workload for a
frame-forwarding decision. The MAC table maps the MAC addresses of the workloads to their corresponding tunnel endpoint (TEP) IP addresses. The TEP IP address identifies the ESXi host where the workload resides. The MAC table is populated by learning the source MAC addresses of the incoming frames from the workloads. The MAC table is also synchronized with other ESXi hosts in the same transport zone by using the NSX Controller.
https://nsx.techzone.vmware.com/resource/nsx-reference-design-guide

NEW QUESTION 29
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 U
D. It must be done via command line interface.
E. 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-A016-4A081AC8

NEW QUESTION 32
Where does an administrator configure the VLANs used In VRF Lite? (Choose two.)

A. segment connected to the Tler-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 37
Sort the rule processing steps of the Distributed Firewall. Order responses from left to right.

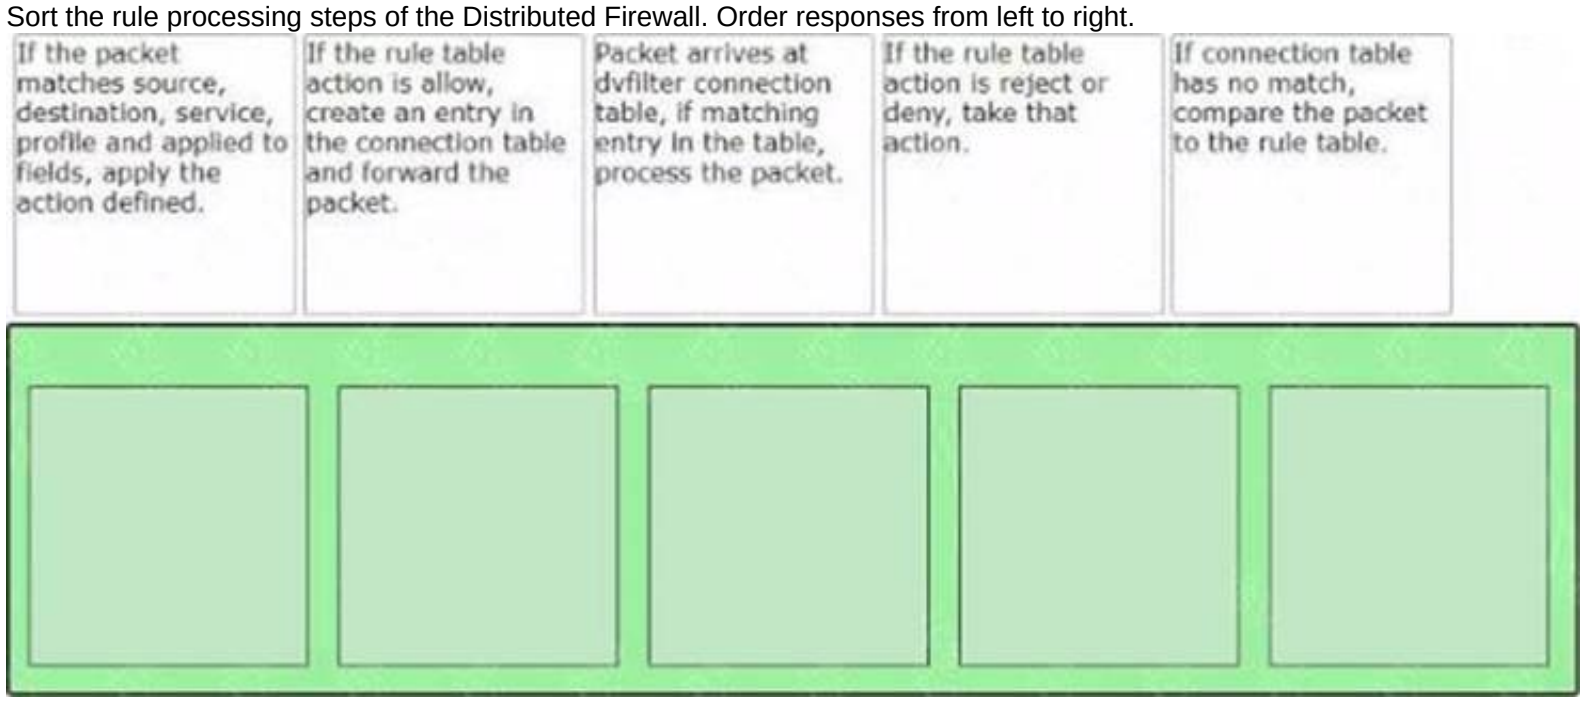

A. Mastered
B. Not Mastered

Answer: A

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 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.

NEW QUESTION 39
Which CLI command is used for packet capture on the ESXi Node?

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

Answer: C

Explanation:
According to the VMware Knowledge Base, this CLI command is used for packet capture on the ESXi node. pktcap-uw stands for Packet Capture User World and is a tool that allows you to capture packets from various points in the network stack of an ESXi host. You can use this tool to troubleshoot network issues or analyze traffic flows.
The other options are either incorrect or not available for this task. tcpdump is not a valid CLI command for packet capture on the ESXi node, as it is a tool that runs on Linux systems, not on ESXi hosts. debug is not a valid CLI command for packet capture on the ESXi node, as it is a generic term that describes the process of finding and fixing errors, not a specific tool or command. set capture is not a valid CLI command for packet capture on the ESXi node, as it does not exist in the ESXi CLI.

NEW QUESTION 44
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

```
esxcli network ip connection list | grep 1234
esxcli network connection list | grep 1234
esxcli network ip connection list | grep 1235
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 47
Which choice is a valid insertion point for North-South network introspection?

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

Answer: C

Explanation:
A valid insertion point for North-South network introspection is Tier-0 gateway. North-South network introspection is a service insertion feature that allows third-party network services to be integrated with
NSX. North-South network introspection enables traffic redirection from the uplink of an NSX Edge node to a service chain that consists of one or more service profiles1. The Tier-0 gateway is the logical router that connects the NSX Edge node to the physical network and provides North-South routing and network services2.
https://docs.vmware.com/en/VMware-NSX/4.1/administration/GUID-D5933474-34A2-4DCE-AE9B-A82FF33

NEW QUESTION 52
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 53
Match the NSX Intelligence recommendations with their correct purpose.

Recommendations:		Purposes:
security policy recommendations		Are service objects that were used by applications in the VMs or physical servers that an administrator had specified, but the services are not yet defined in the NSX inventory.
security group recommendations		Are VMs or physical servers whose traffic flows were analyzed for the time period and the boundary an administrator had specified.
service recommendations		Are East-West distributed firewall (DFW) security policies in the application category.

A. Mastered
B. Not Mastered

Answer: A

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.

NEW QUESTION 54
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-E606AB89

NEW QUESTION 59
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. esxcfg-nics -1l
B. esxcli network ip interface ipv4 get
C. esxcli network nic list
D. esxcfg-vmknic -1l
E. net-dvs

Answer: BD

Explanation:
To check the IP address of the VMkernel port for the Geneve protocol on the ESXi transport node, an NSX administrator can use the following commands:
- esxcli network ip interface ipv4 get: This command displays the IPv4 configuration of all VMkernel interfaces on the host, including their IP addresses, netmasks, and gateways. The Geneve protocol uses a VMkernel interface named geneve0 by default1
- esxcfg-vmknic -l: This command lists all VMkernel interfaces on the host, along with their MAC addresses, MTU, and netstack. The Geneve protocol uses a netstack named nsx-overlay by default

NEW QUESTION 62
Which CLI command shows syslog on NSX Manager?

A. get log-file auth.lag
B. /var/log/syslog/syslog.log
C. show log manager follow
D. get log-file syslog

Answer: D

Explanation:
According to the VMware NSX CLI Reference Guide, this CLI command shows the syslog messages on the NSX Manager node. You can use this command to view the system logs for troubleshooting or monitoring purposes.
The other options are either incorrect or not available for this task. get log-file auth.log is a CLI command that shows the authentication logs on the NSX Manager node, not the syslog messages. /var/log/syslog/syslog.log is not a CLI command, but a file path that may contain syslog messages on some Linux systems, but not on the NSX Manager node. show log manager follow is not a valid CLI command, as there is no show log command or manager option in the NSX CLI.
NSX Cli command get log-file <fiilename>
get log-file <filename> follow
Below are commonly used log files, there are many more log files
get log-file <auth.log | controller | controller-error | http.log | kern.log | manager.log | node-mgmt.log | policy.log | syslog> [follow]
use [follow] to continuing monitor Example: get log-file syslog follow get log-file syslog

NEW QUESTION 64
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 67
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 tailed 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 70
What are tour 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-A2FD426

NEW QUESTION 73
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 74
Which VPN type must be configured before enabling a L2VPN?

A. Route-based IPSec VPN
B. Policy based IPSec VPN
C. SSL-bosed 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 77
What are three NSX Manager roles? (Choose three.)

A. master
B. cloud
C. zookeepet
D. manager
E. policy
F. controller

Answer: DEF

Explanation:
According to the VMware NSX 4.x Professional documents and tutorials, an NSX Manager is a standalone appliance that hosts the API services, the management plane, control plane, and policy management. The NSX Manager has three built-in roles: policy, manager, and controller2. The policy role handles the declarative configuration of the system and translates it into desired state for the manager role. The manager role receives and validates the configuration from the policy role and stores it in a distributed persistent database. The manager role also publishes the configuration to the central control plane. The controller role implements the central control plane that computes the network state based on the configuration and topology information3. The other roles (master, cloud, and zookeeper) are not valid NSX Manager roles.

NEW QUESTION 81
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 85
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? O

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

Answer: B

Explanation:
To configure NSX Manager for two-factor authentication (2FA), an NSX administrator must have VMware
Identity Manager (vIDM) with an OAuth Client added. vIDM provides identity management services and supports various 2FA methods, such as VMware Verify, RSA SecurID, and RADIUS. An OAuth Client is a configuration entity in vIDM that represents an application that can use vIDM for authentication and authorization. NSX Manager must be registered as an OAuth Client in vIDM before it can use
2FA. References: : VMware NSX-T Data Center Installation Guide, page 19. : VMware NSX-T Data Center Administration Guide, page 102. : VMware Blogs: Two-Factor Authentication with VMware NSX-T

NEW QUESTION 86
Which NSX CLI command is used to change the authentication policy for local users?

A. Set cli-timeout
B. Get auth-policy minimum-password-length
C. Set hardening- policy
D. Set auth-policy

Answer: D

Explanation:
According to the VMware NSX Documentation4, the set auth-policy command is used to change the authentication policy settings for local users, such as password length, lockout period, and maximum authentication failures. The other commands are either used to view the authentication policy settings (B), change the CLI session timeout (A), or change the hardening policy settings ©.

NEW QUESTION 90
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 VPN
B. Layer 2 bridge
C. Layer 2 broadcast domain
D. Layer 3 route

Answer: C

Explanation:
An overlay segment is a logical construct that provides Layer 2 connectivity between virtual machines that are attached to it. An overlay segment can span multiple hosts and can be extended across different subnets or locations using Geneve encapsulation3. Therefore, two virtual machines on the same overlay segment belong to the same Layer 2 broadcast domain, which means they can communicate with each other using their MAC addresses without requiring any routing. The other options are incorrect because they involve Layer 3 or higher network boundaries, which require routing or tunneling to connect different segments. References: VMware NSX Documentation

NEW QUESTION 95
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 96
Which two statements describe the characteristics of an Edge Cluster in NSX? (Choose two.)

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

Answer: AE

Explanation:
Two statements that describe the characteristics of an Edge Cluster in NSX are:
- An Edge Cluster can have a maximum of 8 edge nodes2. This is the upper limit for scaling out the Edge Cluster and providing high availability and load balancing for network services.
- An Edge Cluster must contain only one type of edge nodes (VM or bare metal)3. This is because different types of edge nodes have different performance and resource requirements, and mixing them in the same cluster can cause inconsistency and instability. The other options are incorrect because they do not describe the characteristics of an Edge Cluster in NSX. An Edge Cluster can have either
active-active or active-standby edge nodes, depending on the configuration and services4. An Edge Cluster cannot contain multiple types of edge nodes, as explained above. References: Enhanced NSX Edge and Networking Services in NSX 4.0.1.1, NSX Edge Installation Requirements, NSX-T Edge Node Cluster

NEW QUESTION 99
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 104
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-CD

NEW QUESTION 109
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 .

NEW QUESTION 114
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/v1/cluster/api-certificate? action=set_cluster_certificate&certificate_id=<certificate_id>
B. Send an API call to https://<nsx-mgr>/api/v1/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 vip install<certificate_id>
D. SSH as admin into the NSX manager with the cluster VIP IP and run nsxcli cluster certificate node install <certificate_id>

Answer: A

Explanation:
https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/G

NEW QUESTION 116
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

2V0-41.23 Practice Exam Features:

* 2V0-41.23 Questions and Answers Updated Frequently

* 2V0-41.23 Practice Questions Verified by Expert Senior Certified Staff

* 2V0-41.23 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* 2V0-41.23 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year