Cisco

(300-715)

Implementing and Configuring Cisco Identity Services Engine (300-715 SISE)

Total: **311 Questions**

Link:

Question: 1

Which personas can a Cisco ISE node assume?

A. policy service, gatekeeping, and monitoring
B. administration, monitoring, and gatekeeping
C. administration, policy service, and monitoring
D. administration, policy service, gatekeeping

Answer: C

Explanation:

The correct answer is C: administration, policy service, and monitoring. Cisco Identity Services Engine (ISE) nodes are designed to fulfill specific roles within the overall ISE architecture. An ISE node can be configured to operate as an Administration node, which is responsible for the centralized management and configuration of the ISE system. This includes tasks like defining policies, managing users and devices, and handling system-level settings. The Policy Service persona is essential for enforcement. These nodes evaluate access requests, apply defined security policies, and authorize user and device access to network resources based on pre-configured rules. Lastly, Monitoring nodes provide critical visibility into network activity, capturing authentication, authorization, and accounting (AAA) logs, and providing operational insights. The gatekeeping function, implied by options A, B, and D, is essentially incorporated within the policy service role of ISE. There isn't a separate 'gatekeeping' node persona. ISE nodes must assume at least one of these roles to function within the network. Therefore, the personas that a Cisco ISE node can assume are administration, policy service, and monitoring.

Authoritative link for further research:

Cisco Identity Services Engine (ISE) User Guide: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-user-guide-list.html(Refer to documentation on node roles and deployments)

Question: 2

What occurs when a Cisco ISE distributed deployment has two nodes and the secondary node is deregistered?

A. The secondary node restarts.
B. The primary node restarts.
C. Both nodes restart.
D. The primary node becomes standalone.

Answer: C

Explanation:

Here's a detailed justification for why the correct answer is C, "Both nodes restart," when a Cisco ISE distributed deployment with two nodes has its secondary node deregistered:

In a Cisco ISE deployment, especially with a primary and secondary node, the nodes rely on each other for high availability and service continuity. When the secondary node is deregistered, it fundamentally alters the established distributed system configuration. The system must then reorganize to operate with a single node instead of a redundant setup.

Deregistering the secondary node isn't a simple removal of a component; it's an operational change that impacts the primary node's view of the system's health and configuration. As a result, a synchronization and

reconfiguration event is triggered across all active nodes to ensure that the system is consistent and operational. This synchronization process typically includes restarts of services and components to reflect the change. Therefore, deregistering the secondary node forces both the primary node to acknowledge the loss of its failover partner and reconfigure accordingly while the deregistered node reboots to a standalone state.

The system restarts both the primary and the affected former secondary nodes during this process.

The primary node, even though it remains the primary, needs to update its operating configuration to reflect the absence of its secondary peer, and this update involves restarting services to ensure the system is stable.

The deregistered node needs to update its configuration and start as a standalone node, which requires a restart. The answer that specifies the primary node becoming standalone is incorrect because the primary node continues to be a primary (with no peer), not a single, standalone node until it is reconfigured as a standalone system.

In essence, the deregistration disrupts the distributed setup, requiring a controlled restart across both nodes to maintain system integrity. This process is similar to how distributed systems handle node removals or failures, where updates or restarts across participating systems are required to reflect the new configuration.

Authoritative Links for Further Research:

Cisco Identity Services Engine (ISE) Administration Guide:
https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html - Refer to the section covering distributed deployment, node administration and troubleshooting. Search for specific sections related to "deregistering a node" and "node behavior upon deregistration."
Cisco Community:https://community.cisco.com/t5/security/bd-p/5671j-disc-sec - Search the Cisco community forums for related discussions. Look for questions and answers related to ISE node deregistration and how the system behaves.

Question: 3

DRAG DROP -
Drag the steps to configure a Cisco ISE node as a primary administration node from the left into the correct order on the right.
Select and Place:

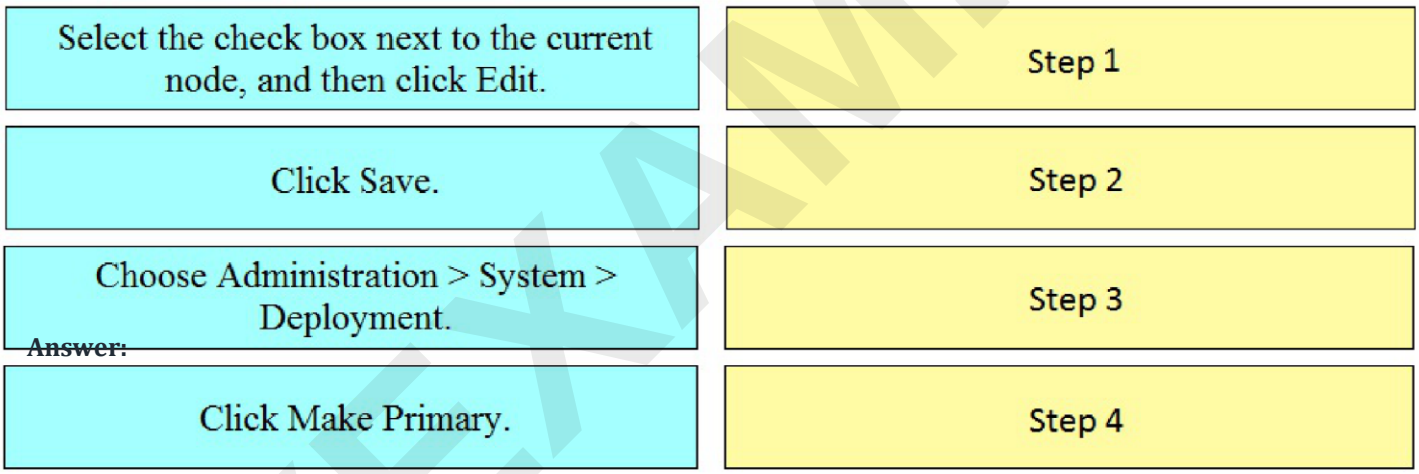

Answer:

Select the check box next to the current node, and then click Edit.	Choose Administration > System > Deployment.
Click Save.	Select the check box next to the current node, and then click Edit.
Choose Administration > System > Deployment.	Click Make Primary.
Click Make Primary.	Click Save.

Explanation:

Reference:
https://www.cisco.com/en/US/docs/security/ise/1.0/user_guide/ise10_dis_deploy.html

Question: 4

Which two features are available when the primary admin node is down and the secondary admin node has not been promoted? (Choose two.)

A. new AD user 802.1X authentication
B. hotspot
C. posture
D. guest AUP
E. BYOD

Answer: BD

Explanation:

Here's a detailed justification for why the correct answer is B and D (Hotspot and Guest AUP) when the primary admin node is down and the secondary admin node has not been promoted in Cisco ISE:

In a Cisco ISE deployment with a primary and secondary admin node, the admin nodes are primarily responsible for configuration and policy management. When the primary admin node fails and the secondary has not taken over, core functions requiring admin-level access are unavailable.

Specifically:

Authentication services (A) such as 802.1X with new Active Directory (AD) user authentication, rely on the live configuration and policy data which is primarily handled by the admin nodes, are unavailable. These are dynamic processes that are not cached or replicated for use in this scenario. The failed primary admin node and non-promoted secondary prevents this.

Posture assessment (C), another vital process that relies on the ISE policy engine which resides on the admin node, would also fail. Posture checks verify device compliance, which requires ongoing communication with the policy management engine and active sessions, and is not a static service.

However, certain basic functions continue to operate that do not require live administrative interaction. These include:

Hotspot (B): This typically uses pre-configured settings, often using a sponsored guest setup, and will continue to offer a basic login page. The authentication might still work through a pre-existing RADIUS server.

Guest AUP (D): The Acceptable Use Policy page is usually cached or is a static page. Therefore, the guest

user can still acknowledge it when they are accessing the network via the hotspot.

Therefore, even if the primary admin node is down and the secondary has not been promoted, hotspot services and Guest AUP acknowledgment can still function because they don't require active administrative management or policies to work.

Authoritative Links for Further Research:

Cisco ISE High Availability Guide:https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/admin_guide/b_ise_admin_2_7/b_ise_admin_27_chapter_01011.html
Cisco ISE Failover Scenarios:https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/admin_guide/b_ise_admin_2_7/b_ise_admin_27_chapter_01011.html#id_55132

Question: 5

Which supplicant(s) and server(s) are capable of supporting EAP-CHAINING?

A. Cisco Secure Services Client and Cisco Access Control Server
B. Cisco AnyConnect NAM and Cisco Identity Service Engine
C. Cisco AnyConnect NAM and Cisco Access Control Server
D. Windows Native Supplicant and Cisco Identity Service Engine

Answer: B

Explanation:

EAP-CHAINING, a feature enabling multiple EAP authentication methods sequentially, requires specific supplicant and server support. Cisco AnyConnect Network Access Manager (NAM) is designed to handle complex authentication scenarios, including EAP-CHAINING, acting as a versatile supplicant. Cisco Identity Services Engine (ISE) serves as the policy enforcement and authentication server, and is equipped to process multi-layered EAP authentication sequences, making it compatible with EAP-CHAINING. Neither Cisco Access Control Server (ACS) nor the Windows Native Supplicant are specifically designed to handle EAP-CHAINING, they typically only handle a single EAP authentication at a time. ACS is an older product, and Windows Native Supplicant has limited EAP capabilities compared to NAM. Therefore, only the pairing of Cisco AnyConnect NAM and Cisco ISE fulfill the requirement for initiating and processing EAP-CHAINING. The other options, which pair ACS with either NAM or CSC and Windows Native Supplicant with ISE lack the necessary EAP-CHAINING capability on either the supplicant or server side, or both.

Further Research:

Cisco ISE documentation: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-technical-reference-list.html
Cisco AnyConnect NAM documentation: https://www.cisco.com/c/en/us/support/security/anyconnect-secure-mobility-client/products-technical-reference-list.html

Question: 6

What is a requirement for Feed Service to work?

A. TCP port 8080 must be opened between Cisco ISE and the feed server.
B. Cisco ISE has access to an internal server to download feed update.
C. Cisco ISE has a base license.
D. Cisco ISE has Internet access to download feed update.

Answer: B

Explanation:

The correct answer is **D. Cisco ISE has Internet access to download feed update.**

Here's the justification:

Cisco Identity Services Engine (ISE) relies on the Feed Service to obtain the latest updates for various security intelligence feeds, such as threat intelligence data, vulnerability information, and location-based services.

These feeds are essential for maintaining the efficacy of ISE's security policies. For the Feed Service to function correctly, ISE must have the ability to connect to external servers hosted by Cisco (or other providers). This connection happens over the internet.

Option D directly states this requirement: Internet access to download feed updates. This is crucial because Cisco doesn't typically host these updates on internal networks; they are publicly available through the internet via specific URLs. The Feed Service downloads these feeds at scheduled intervals to ensure ISE has the most current data.

Option A, while mentioning a port, is incorrect. The communication between ISE and the feed servers does not specifically rely on port 8080 for all feed downloads. The port used for the feed download is usually port 443 (HTTPS).

Option B is incorrect because while ISE can download feeds from an internal server, this is not a typical or a requirement to get updates from the internet. This would require manual configuration and is often used in air-gapped environments. In most cases, ISE will download directly from the internet.

Option C is incorrect as the Feed Service is a core function of ISE and not dependent on a specific license type beyond a base license that is already required to operate ISE.

Therefore, having internet access for downloading feed updates from external servers is a fundamental prerequisite for the Feed Service to operate in a typical deployment of Cisco ISE.

Authoritative Links:

Cisco Identity Services Engine (ISE) Administrator Guide:
https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/admin_guide/b_ise_admin_guide_3_2/b_ise_admin_guide_32_chapter_0100.html#reference_g24_wzn_cqb (Search for "Feed Service")
Cisco ISE Configuration:https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/admin_guide/b_ise_admin_guide_3_2/b_ise_admin_guide_32_chapter_0100.html (This page contains various information on ISE including feed services)

Question: 7

What is a method for transporting security group tags throughout the network?

A. by embedding the security group tag in the 802.1Q header
B. by the Security Group Tag Exchange Protocol
C. by enabling 802.1AE on every network device
D. by embedding the security group tag in the IP header

Answer: B

Explanation:

The correct method for transporting Security Group Tags (SGTs) throughout a network is using the Security Group Tag Exchange Protocol (SGT Exchange Protocol or SXP). SXP is a crucial component of Cisco's TrustSec solution. It facilitates the propagation of SGT information between network devices that may not have direct line-of-sight access to the Identity Services Engine (ISE). SXP operates independently of the data plane; it's a control plane mechanism. This separation allows for scalability and reduces the overhead on data packets. Options A and D, embedding SGTs in the 802.1Q or IP header, are not standard or supported methods for SGT transport. While 802.1Q headers are used for VLAN tagging, and IP headers are used for network addressing, they are not designed to carry SGT information in a scalable, standard way across the entire network. 802.1AE, while related to security, focuses on MACsec encryption and doesn't directly facilitate SGT propagation. Therefore, option B, using the SXP protocol, is the only option that aligns with how Cisco TrustSec networks are designed to operate, allowing for consistent enforcement of security policies across a network based on user identity, device, and application. SXP essentially acts as a translator, allowing devices that lack direct access to ISE to understand and enforce the security policies associated with different SGTs.

Cisco TrustSec Solution GuideCisco Identity Services Engine (ISE) Design Guide

Question: 8

An engineer is configuring a virtual Cisco ISE deployment and needs each persona to be on a different node. Which persona should be configured with the largest amount of storage in this environment?

A. Monitoring and Troubleshooting
B. Policy Services
C. Primary Administration
D. Platform Exchange Grid

Answer: A

Explanation:

The correct answer is **A. Monitoring and Troubleshooting**. In a distributed Cisco ISE deployment where each persona resides on a separate node, the Monitoring and Troubleshooting (MnT) persona typically requires the largest storage allocation. This is because the MnT node is responsible for collecting, storing, and processing vast amounts of operational data, including logs, session details, endpoint information, and historical records. This data volume is significantly higher than that of other personas. The Policy Service Node (PSN) primarily focuses on policy enforcement and doesn't retain extensive historical data. The Primary Administration Node (PAN) manages the ISE deployment configuration and doesn't necessitate large storage. The Platform Exchange Grid (PxGrid) persona facilitates information exchange with other security systems and doesn't demand significant storage for data retention. Effective monitoring and troubleshooting relies heavily on the ability to retain and analyze large volumes of log data. Therefore, the MnT persona requires a substantial storage capacity to function efficiently and to maintain the necessary historical context for issue diagnosis. Insufficient storage on the MnT node could impede visibility into network activity, hinder troubleshooting, and potentially affect overall network security. For deeper understanding, consider exploring the Cisco ISE deployment best practices and sizing guidelines documentation.

Relevant links:

Cisco Identity Services Engine (ISE) Design and Implementation Guide:
https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/ise_design/b_ise_design_guide_27.html (Look for sections on deployment models and sizing)
Cisco ISE Installation Guide:https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/installation_guide/b_ise_installation_guide_27.html (Details storage requirements for different personas)

Question: 9

In a standalone Cisco ISE deployment, which two personas are configured on a node? (Choose two.)

A. subscriber
B. primary
C. administration
D. publisher
E. policy service

Answer: CE

Explanation:

In a standalone Cisco ISE deployment, the two essential personas configured on a single node are **Administration** and **Policy Service**. The Administration persona is crucial for managing the ISE system itself.

This includes tasks like configuration, monitoring, and maintenance. It provides the graphical interface and underlying processes required for an administrator to interact with the system. The Policy Service persona handles all aspects related to network access policy. This includes authentication, authorization, and accounting (AAA) functions. It's the engine that enforces the policies defined by the administrator, enabling secure access control to the network. In a standalone setup, both of these functionalities reside on the same physical node. The absence of other nodes eliminates the need for secondary roles like publisher, subscriber, or primary, as these are pertinent to distributed deployments. A single node performs all system and policy operations, making administration and policy service functionalities essential for core function.

Therefore, the correct answer is **C. Administration** and **E. Policy Service**.

Authoritative links for further research:

Cisco ISE documentation: https://www.cisco.com/c/en/us/support/security/identity-services-engine/tsd-products-support-series-home.html
Cisco ISE Deployment Models: https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/install_guide/b_ise_installation_guide_27/b_ise_installation_guide_27_chapter_010.html

Question: 10

A network engineer must enforce access control using special tags, without re-engineering the network design. Which feature should be configured to achieve this in a scalable manner?

A. RBAC
B. dACL
C. SGT
D. VLAN

Answer: C

Explanation:

The correct answer is C, Security Group Tags (SGTs). SGTs provide a scalable method for enforcing access control based on user, device, or application classifications, rather than relying on network topology like VLANs or IP addresses. They allow for policy definition decoupled from IP addressing schemes, simplifying network management and policy changes. RBAC (Role-Based Access Control), option A, focuses on user roles and their associated permissions but doesn't directly tag or categorize traffic on the network level. Option B,

dACL (Dynamic Access Control Lists), although flexible, is primarily configured at the interface level, thus not easily scalable for user or application groupings and might be heavily dependent on the network topology. Option D, VLANs (Virtual LANs), segments the network based on broadcast domains, which is fundamentally a network topology division and does not provide a tag for application or user-based access control. SGTs, part of Cisco's TrustSec framework, are embedded within network packets and utilized by network devices, particularly the Cisco ISE, to enforce consistent access policies. This "tag-based" approach is significantly more scalable because policies can be applied based on the SGT value, irrespective of IP address or location in the network. When new resources or users are added, they simply receive the appropriate SGT and automatically inherit the policy. This mechanism also enhances micro-segmentation and zero-trust security principles.

For further reading, consult Cisco documentation on TrustSec and Security Group Tags:

Cisco TrustSec:https://www.cisco.com/c/en/us/solutions/enterprise/trustsec/index.html
Cisco ISE:https://www.cisco.com/c/en/us/products/security/identity-services-engine/index.html
Implementing Cisco TrustSec:https://www.ciscopress.com/articles/article.asp?p=2989079&seqNum=2

Question: 11

A network engineer is configuring a network device that needs to filter traffic based on security group tags using a security policy on a routed interface.
Which command should be used to accomplish this task?

A.cts role-based policy priority-static
B.cts cache enable
C.cts authorization list
D.cts role-based enforcement

Answer: D

Explanation:

The correct answer is **D. cts role-based enforcement**. This command is used to activate Security Group Tag (SGT)-based enforcement on a Cisco device's interface. SGTs are crucial for implementing role-based access control, enabling network devices to apply security policies based on user or device classifications rather than just IP addresses. The cts role-based enforcement command instructs the device to look for SGTs within traffic frames and match them against defined security policies, which are configured based on source and destination SGTs. Option A, cts role-based policy priority-static, is used to set a priority for static policies related to SGTs but doesn't enable enforcement on the interface. Option B, cts cache enable, is used to enable SGT caching, which improves performance but isn't directly involved in enforcing policies. Option C, cts
authorization list, relates to authentication and authorization using SGTs but not actual traffic filtering on an interface. To apply a security policy based on SGTs to filter traffic on a routed interface, you must first enable SGT enforcement on the interface using cts role-based enforcement. Only then will policies based on SGTs be considered for traffic filtering. In summary, while other commands relate to SGT configurations, only cts role-based enforcement directly allows enforcement of SGT based policies on an interface.

For further research, refer to Cisco's official documentation on TrustSec and SGTs:

Cisco TrustSec Configuration Guide: https://www.cisco.com/c/en/us/td/docs/solutions/CVD/Campus/cisco-campus-network-design-1-0-external.html (Search for "Security Group Tags (SGTs)")
Cisco Identity Services Engine (ISE) Configuration Guides: Search for relevant guides for your ISE version. These documents will detail SGT based policy configurations.

Question: 12

In a Cisco ISE split deployment model, which load is split between the nodes?

A.log collection
B.device admission
C.AAA
D.network admission

Answer: C

Explanation:

The correct answer is C, **AAA (Authentication, Authorization, and Accounting)**. In a Cisco ISE split deployment, the primary function of distributing the workload is to handle AAA requests. This means that the processing of authentication, authorization, and accounting functions is divided among the nodes in the deployment. Specifically, a primary Policy Administration Node (PAN) is present to manage configurations while multiple Secondary Monitoring Nodes (MN) and Secondary Policy Service Nodes (PSN) handle these AAA processes. Log collection (A) and device admission (B) are primarily handled by the Monitoring Node and the primary PAN. Network admission (D) is related to the overall ISE functionality, but the split workload in a distributed ISE environment specifically concerns how AAA requests are managed across the nodes. By distributing the AAA load, ISE achieves high availability and scalability. This prevents a single point of failure and allows the system to handle a larger number of simultaneous authentication and authorization requests.

Essentially, client authentication attempts are directed to these PSNs. This design allows a system to scale and support large enterprise networks. Distributing the AAA load ensures that the system remains responsive even with a large number of devices connecting to the network.Further information can be found in the Cisco documentation:

Cisco Identity Services Engine Design Guide
Cisco ISE Deployment Models

Question: 13

What is the deployment mode when two Cisco ISE nodes are configured in an environment?

A.standalone
B.distributed
C.standard
D.active

Answer: B

Explanation:

When two Cisco Identity Services Engine (ISE) nodes are deployed, they operate in a distributed deployment mode. This mode is characterized by the distribution of functions and data across multiple ISE nodes. Unlike a standalone deployment, where a single node handles all functions, a distributed setup provides redundancy and scalability. One node is designated as the primary Policy Administration Node (PAN), responsible for configuration and policy management. The other node typically acts as a secondary PAN, serving as a backup and capable of taking over primary duties in case of failure. The remaining functions such as Monitoring, Logging, and Policy Service Nodes (PSN), are distributed between the two nodes, ensuring higher availability.

An active-active or active-standby design can be employed. The active-active design allows both nodes to process requests simultaneously to distribute the load. In the active-standby design, the secondary node is

idle until the primary fails. This configuration provides high availability, fault tolerance, and improved performance compared to a single node setup. A standard term or single term does not accurately describe such an setup. Therefore, the distributed mode is the correct choice when multiple ISE nodes are present.

Authoritative Links:

Cisco Identity Services Engine Administrator Guide:https://www.cisco.com/c/en/us/td/docs/security/ise/3-1/admin_guide/b_ise_admin_3_1/b_ise_admin_31_chapter_01010.html (Refer to sections on Deployment Models)
Cisco ISE Deployment Models:https://www.cisco.com/c/en/us/products/collateral/security/identity-services-engine/guide-c07-733026.html (Explains different ISE deployment options)

Question: 14

An engineer is testing Cisco ISE policies in a lab environment with no support for a deployment server. In order to push supplicant profiles to the workstations for testing, firewall ports will need to be opened.
From which Cisco ISE persona should this traffic be originating?

A.administration
B.authentication
C.policy service
D.monitoring

Answer: C

Explanation:

The correct answer is C, Policy Service Persona. Cisco ISE's architecture relies on distinct personas, each with specific functions. The Policy Service persona is responsible for handling authentication, authorization, and policy enforcement. This includes pushing supplicant profiles to endpoints. In the context of testing without a deployment server, the policy service node would be the one initiating the communication to push these profiles. The other personas have different roles; the Administration persona handles configuration and management tasks, the Authentication persona deals with identity verification, and the Monitoring persona gathers logs and reports. Since pushing supplicant profiles falls directly under policy enforcement, the Policy Service persona is the logical origin of the traffic. This process uses protocols like RADIUS or potentially proprietary methods specific to ISE and its endpoint posture assessment capabilities. Essentially, the policy service node acts as the central point for distributing and enforcing access control policies including those related to supplicant configuration, to network endpoints.

Relevant documentation that explains ISE personas and their functions can be found in the Cisco ISE documentation:

Cisco Identity Services Engine Administrator Guide: This is the primary resource for understanding ISE architecture and features: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html
Cisco ISE deployment models: Understand the different personas role on different deployment models here:https://www.cisco.com/c/en/us/td/docs/security/ise/3-0/admin_guide/b_ise_admin_3_0/b_ise_admin_30_chapter_010.html

Question: 15

What does a fully distributed Cisco ISE deployment include?

A.PAN and MnT on the same node while PSNs are on their own dedicated nodes.

B.All Cisco ISE personas are sharing the same node.

C.All Cisco ISE personas on their own dedicated nodes.

D.PAN and PSN on the same node while MnTs are on their own dedicated nodes.

Answer: C

Explanation:

The correct answer is **C. All Cisco ISE personas on their own dedicated nodes.** This defines a fully distributed Cisco ISE deployment, which aims for high availability and scalability by separating the functions of each persona. A fully distributed deployment means each of the ISE personas – Policy Administration Node (PAN), Monitoring and Troubleshooting Node (MnT), and Policy Service Node (PSN) – operates on its own dedicated server or virtual machine. This separation of duties allows for independent scaling and resilience. If a PSN fails, the others continue to function, reducing downtime. The PAN, responsible for administration and configuration, benefits from dedicated resources, ensuring a responsive management interface. Similarly, having an MnT on a separate node ensures effective logging and troubleshooting capabilities, as it won't be impacted by other ISE functions. Option A is incorrect because it groups PAN and MnT, defeating the purpose of a fully distributed model. Option B suggests a single-node deployment which isn't distributed. Option D incorrectly groups PAN and PSN, and separates MnT, not providing optimal performance or resilience. In summary, separating all personas onto individual nodes offers the highest level of scalability, redundancy, and operational efficiency for a Cisco ISE deployment.

Authoritative Links for Further Research:

Cisco Identity Services Engine (ISE) Deployment Models:
https://www.cisco.com/c/en/us/td/docs/security/ise/3-1/install_guide/b_ise_InstallationGuide31/b_ise_InstallationGuide_chapter_01010.html#con_1114887 (Refer to the "Deployment Models" section for a breakdown of distributed deployments)
Cisco ISE Admin Guide: Search within the guide for "distributed deployment" for more detailed information.
https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html

Question: 16

An engineer is configuring 802.1X and wants it to be transparent from the users' point of view. The implementation should provide open authentication on the switch ports while providing strong levels of security for non-authenticated devices. Which deployment mode should be used to achieve this?

A. closed

B. high-impact

C. low-impact

D. open

Answer: C

Explanation:

The correct deployment mode is **C. low-impact**. Low-impact mode in 802.1X deployments, specifically within Cisco Identity Services Engine (ISE), is designed to minimize user disruption while still enforcing security policies. In this mode, devices initially connect to the network with open access (like open authentication, as described in the question). Then, in the background, ISE performs authentication checks. If authentication is successful, the device gains the appropriate network access. Conversely, non-authenticated devices are

subjected to defined access limitations, ensuring that security policies are enforced without impeding user experience. This achieves the desired transparency by not immediately blocking network access during initial connection and provides robust protection through conditional authentication. High-impact mode, conversely, would immediately block access until successful authentication, violating the transparency requirement, and closed mode would only allow authorized devices, thus impeding the initial open access needed. Open mode is a misnomer as it usually implies no authentication. Therefore, low-impact provides the ideal balance of seamless connectivity and security, aligning directly with the engineer's objective.

Further research:

Cisco ISE Deployment Modes:https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/admin_guide/b_ise_admin_guide_2_7/b_ise_admin_guide_27_chapter_01111.html (Specifically, search for "Low-Impact Mode")
Understanding 802.1X:https://networkdirection.net/articles/wireless/understanding-802-1x/

Question: 17

A network administrator changed a Cisco ISE deployment from pilot to production and noticed that the JVM memory utilization increased significantly. The administrator suspects this is due to replication between the nodes. What must be configured to minimize performance degradation?

A.Enable the endpoint attribute filter.
B.Review the profiling policies for any misconfiguration.
C.Ensure that Cisco ISE is updated with the latest profiler feed update.
D.Change the reauthentication interval.

Answer: A

Explanation:

The correct answer is **A. Enable the endpoint attribute filter.** Here's why:

When Cisco ISE is deployed, it constantly collects endpoint data (attributes) like device type, operating system, and more. This data is crucial for profiling and policy enforcement. However, in a large production environment, the sheer volume of attributes can overwhelm the system, especially during replication between ISE nodes. This excessive data transfer leads to increased JVM memory utilization, causing performance degradation.

Enabling the endpoint attribute filter allows the administrator to specify which attributes are replicated between nodes. By filtering out unnecessary attributes, the replication process becomes more efficient, reducing the data volume being transferred and processed. This directly translates to lower JVM memory consumption and improved performance. Think of it like only sending essential information across a network instead of everything.

Options B, C, and D are less relevant to the specific problem of replication-induced JVM memory issues.
Reviewing profiling policies (B) and ensuring the latest profiler feed (C) are good practices for endpoint visibility and identification, but do not directly address the replication bottleneck. Changing the reauthentication interval (D) affects the frequency of authentication requests, not the data volume replicated.

In essence, endpoint attribute filtering is a targeted approach to minimizing replication overhead in Cisco ISE, thereby mitigating performance degradation linked to high JVM memory utilization. It aligns with the concept of optimizing data transfer for efficiency in a distributed system.

Authoritative Links for Further Research:

Cisco Identity Services Engine (ISE) Administrator Guide: https://www.cisco.com/c/en/us/td/docs/security/ise/3-1/admin_guide/b_ise_admin_3_1.html (Search for sections on endpoint filtering and replication)
Cisco ISE Design Guide: https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/design_guide/b_ise_dg_2_7.html (Look for information regarding deployment best practices and performance tuning)

Question: 18

An administrator is attempting to replace the built-in self-signed certificates on a Cisco ISE appliance. The CA is requesting some information about the appliance in order to sign the new certificate.
What must be done in order to provide the CA this information?

A.Install the Root CA and intermediate CA.
B.Generate the CSR.
C.Download the CA server certificate.
D.Download the intermediate server certificate.

Answer: B

Explanation:

The correct answer is B. **Generate the CSR (Certificate Signing Request).**

Here's why: When replacing self-signed certificates on a Cisco ISE appliance with CA-signed certificates, you must first create a CSR. This CSR contains vital information about the appliance that the Certificate Authority (CA) needs to create a signed certificate. This information includes the appliance's fully qualified domain name (FQDN), organization information, and the public key generated by the appliance. The CA will use the public key, organization info, and FQDN from the CSR to create a certificate. The private key of the new key pair remains with the ISE appliance.

Options A, C, and D are incorrect because they are relevant only after the CSR is generated and the CA returns the signed certificate. Installing the root and intermediate CA certificates is something that you do on clients, so that they know to trust the server certificates. The CA server certificate and the intermediate server certificate would come from the CA, and would be imported by the ISE appliance after receiving the signed certificates. Generating the CSR is the first step in the process of replacing the self-signed certificates with CA-signed certificates.

In summary: A CSR is necessary for requesting a signed certificate from a CA. The CSR provides the CA with the required details to generate a certificate that is specific to the Cisco ISE appliance.

Authoritative Links for further research:

Cisco Identity Services Engine (ISE) Administrator Guide, Release 3.2: https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/admin_guide/b_ise_admin_3_2/b_ise_admin_32_chapter_010010.html (Search for "Generate Certificate Signing Request (CSR)")
Cisco ISE Certificate Management:https://www.cisco.com/c/en/us/support/docs/security/identity-services-engine/214842-ise-certificate-management.html
What is a CSR? How to Generate a CSR:https://www.globalsign.com/en/blog/how-to-generate-a-csr

Question: 19

An administrator is adding network devices for a new medical building into Cisco ISE. These devices must be in a network device group that is identifying them as
`Medical Switch` so that the policies can be made separately for the endpoints connecting through them. Which configuration item must be changed in the network device within Cisco ISE to accomplish this goal?

A. Change the device profile to Medical Switch.
B. Change the device type to Medical Switch.
C. Change the device location to Medical Switch.
D. Change the model name to Medical Switch.

Answer: B

Explanation:

The correct answer is **B. Change the device type to Medical Switch.**

Cisco ISE uses Network Device Groups (NDGs) to logically categorize network devices for easier policy management. These groups are defined by attributes of the network devices, and Device Type is a key attribute used for NDG membership. When an administrator wants to group devices from a new medical building under the Medical Switch NDG, they need to configure the Device Type of each device to reflect this categorization. Setting the device type to "Medical Switch" tells ISE that these devices belong to the "Medical Switch" group, allowing for specific authorization and authentication policies to be applied to endpoints connecting through those devices.

Options A, C, and D are incorrect because they deal with different device attributes. Device Profile helps identify the operating system and capabilities of a device, which doesn't directly influence NDG membership. Device Location helps track physical location for monitoring purposes, not for policy enforcement. Model Name provides more specific identification of the hardware, and is unrelated to network grouping.

Therefore, changing the Device Type to Medical Switch is the direct method to achieve the desired device grouping for policy enforcement within Cisco ISE. By using this attribute, the administrator can effectively apply policies to the specific medical building devices.

Authoritative Links:

Cisco Identity Services Engine (ISE) Administrator Guide: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html (Refer to the section on Network Device Administration and Grouping) **Cisco ISE Network Device Groups:** [Search on Cisco website using this keyword for specific documentation] **Cisco ISE Configuration Guides:** [Search on Cisco website for ISE Configuration examples with NDG configuration]

Question: 20

An organization wants to split their Cisco ISE deployment to separate the device administration functionalities from the main deployment. For this to work, the administrator must deregister any nodes that will become a part of the new deployment, but the button for this option is grayed out.
Which configuration is causing this behavior?

A.All of the nodes are actively being synched.
B.All of the nodes participate in the PAN auto failover.
C.One of the nodes is an active PSN.
D.One of the nodes is the Primary PAN.

Answer: D

Explanation:

Here's a detailed justification for why the correct answer is D, "One of the nodes is the Primary PAN," and why the other options are incorrect:

Justification for Option D:

In a Cisco ISE deployment, the Primary Policy Administration Node (PAN) holds the central configuration and management role. It's responsible for synchronizing configurations to secondary PANs and Policy Service Nodes (PSNs). Deregistering a node from an ISE deployment requires careful consideration, especially for the Primary PAN. The primary PAN cannot be deregistered directly without first promoting a secondary PAN to take over its primary role. This is a critical safety mechanism to prevent disruption to the ISE infrastructure and the services it provides. The deregister option for a node will be greyed out if it is the primary PAN to prevent accidental disruption. The system requires a controlled failover process before that node can be removed from the deployment.

Why other options are incorrect:

A. All of the nodes are actively being synched: Active synchronization is a normal and expected operation within an ISE deployment. While synchronization is essential for consistency, it doesn't prevent a node from being deregistered in theory; however, there are some sync states which block node deregistration.

B. All of the nodes participate in the PAN auto failover: Auto failover is a configuration designed for high availability. While it indicates a redundant setup, the presence of failover doesn't dictate if a node can or cannot be deregistered, unless, it is a primary PAN.

C. One of the nodes is an active PSN: Policy Service Nodes (PSNs) handle runtime authentication and authorization policies. While they are crucial to the deployment, deregistering a PSN does not involve the same process as the PAN and, therefore, would not be a reason the deregister button is greyed out. PSNs can be easily deregistered from the deployment without the complexity that is involved with the PAN.

Key Concepts:

Cisco ISE Architecture: Understanding the roles of PANs (Primary and Secondary) and PSNs is critical for administering an ISE deployment. The PAN is the core component responsible for the overall configuration and management, while PSNs enforce policies.

Primary PAN Functionality: The Primary PAN is the single source of truth for configuration data. Its role is paramount for a stable and operational ISE environment.

Controlled Failover: Deregistering the Primary PAN requires careful planning and a controlled failover mechanism to ensure continuous operations, therefore a secondary PAN is required before a primary can be deregistered.

Authoritative Links for Further Research:

Cisco Identity Services Engine Administrator Guide, Release 3.2:
https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/admin_guide/b_ise_admin_guide_3_2/m_ise_arch.html (This link provides general information about ISE architecture and roles.)
Cisco ISE Operations Guide: (Search on cisco.com for specific ISE version operations guide for detailed information on node registration, deregistration, and failover. The latest version should be readily available.)

In summary, the fact that the "deregister" button is greyed out for a node is a direct indication that the specific node is the Primary PAN, and therefore, a controlled failover is needed prior to any deregistration process.

Question: 21

A network administrator must configure Cisco ISE Personas in the company to share session information via syslog. Which Cisco ISE personas must be added to syslog receivers to accomplish this goal?

A.admin
B.policy services
C.monitor
D.pxGrid

Answer: C

Explanation:

The correct answer is **C. monitor**.

Cisco ISE uses different personas, each with specific roles, to manage the system effectively. The "monitor" persona is specifically responsible for collecting and reporting logs, including session information. To send session data via syslog, the syslog receiver must be configured to receive logs from the ISE node running the "monitor" persona.

The "admin" persona focuses on management tasks such as configuration and user access control. The "policy service" persona handles policy evaluation and enforcement, but does not directly generate or forward syslog messages containing session details. While "pxGrid" is used for sharing context and security information with other platforms, it does not directly handle the syslog forwarding of session data. Therefore, only the monitor persona is relevant for syslog forwarding of session information within Cisco ISE. Configuring the syslog receiver to accept logs from the monitor persona ensures that the system session information required by the administrator is received for auditing and analysis.

Further Research:

1. **Cisco Identity Services Engine (ISE) Deployment Architecture Guide:** https://www.cisco.com/c/en/us/td/docs/security/ise/3-3/ise_deployment/b_ise_deployment_3_3_guide/m_ise_persona.html This official Cisco documentation provides detailed information on ISE personas and their roles, including the monitoring persona and its logging capabilities.
2. **Cisco ISE Syslog Configuration:** Search on Cisco's website for "Cisco ISE Syslog Configuration" to find configuration guides, best practices, and detailed documentation.

Question: 22

What is the maximum number of PSN nodes supported in a medium-sized deployment?

A.two
B.three
C.five
D.eight

Answer: C

Explanation:

The correct answer is C, five. Cisco Identity Services Engine (ISE) deployments are categorized by size (small,

medium, and large), each with specific scalability limits. The number of Policy Service Nodes (PSNs), which handle authentication, authorization, and accounting requests, directly impacts the system's capacity. A medium-sized ISE deployment is designed to handle a moderate number of endpoints and concurrent users. Cisco documentation specifies that a medium-sized deployment supports a maximum of five PSNs. This limit is chosen to balance performance, redundancy, and management complexity. Deploying more than five PSNs in a medium-sized setup may lead to performance issues and is generally not recommended by Cisco. The limit is enforced to ensure predictable and reliable operation within the resources allocated for a medium-sized environment. This configuration ensures optimal resource utilization and avoids over-provisioning, aligning with cloud computing best practices of right-sizing resources. Exceeding the supported number of PSNs can lead to inefficiencies and potential stability problems. These limitations are crucial for maintaining a robust and manageable ISE infrastructure. The selection of a medium-sized deployment implies a specific scale and budget, with these limits accordingly scaled. The number of PSNs isn't just about load capacity but is also about management of resources and complexity. For a system to grow beyond the five PSN limit, upgrading to a large deployment architecture is necessary.**Authoritative Link:**

Cisco Identity Services Engine Hardware and Virtual Machine Installation Guides - While this is a link to the main documentation page, explore the installation and deployment guides within to find the recommended deployment sizes and PSN limitations.

Question: 23

How is policy services node redundancy achieved in a deployment?

A.by creating a node group
B.by deploying both primary and secondary node
C.by enabling VIP
D.by utilizing RADIUS server list on the NAD

Answer: A

Explanation:

The correct answer is **A. by creating a node group.** Policy Service Node (PSN) redundancy in Cisco ISE is primarily achieved by grouping multiple PSNs into a node group. This allows for load balancing and high availability. When a Network Access Device (NAD) authenticates a user or device, it sends the request to a specific PSN within the group. If that PSN becomes unavailable, the NAD will be directed to another active PSN in the group, ensuring continuous service. This is a fundamental concept in load balancing, distributing traffic across multiple resources. A node group is essentially a logical entity that treats multiple ISE instances as a single unit for policy processing. Option B is incorrect; while a primary and secondary admin node exist, they manage the ISE deployment, not the PSN redundancy directly. Option C, VIP (Virtual IP), is relevant for load balancing but is not how redundancy is achieved in the ISE PSN setup; node groups facilitate the process. Option D, the RADIUS server list on NAD, is how NADs locate and use ISE instances; the redundancy itself is within ISE, through node groups. The node group structure provides failover capabilities, meaning requests are automatically redirected, and is a vital aspect of ensuring network reliability and availability.

For further reading, refer to Cisco's official documentation:

Cisco Identity Services Engine Administrator Guide, Release 3.1:
https://www.cisco.com/c/en/us/td/docs/security/ise/3-1/admin_guide/b_ise_admin_3_1.html
Cisco Identity Services Engine Hardware Installation Guide, Release 3.1:
https://www.cisco.com/c/en/us/td/docs/security/ise/3-1/installation_guide/b_ise_hig_3-1.html (Look for sections on PSN deployment and high availability)

Question: 24

Which two fields are available when creating an endpoint on the context visibility page of Cisco ISE? (Choose two.)

A.Security Group Tag
B.Endpoint Family
C.Policy Assignment
D.Identity Group Assignment
E.IP Address

Answer: CD

Explanation:

Here's a detailed justification for why options C (Policy Assignment) and D (Identity Group Assignment) are the correct choices when creating an endpoint on the context visibility page of Cisco ISE, while A (Security Group Tag), B (Endpoint Family), and E (IP Address) are not directly available at this stage:

The Cisco ISE Context Visibility page provides a centralized view of network endpoints and their associated attributes. When manually creating an endpoint on this page, the primary focus is on defining its identity within the ISE system for policy application. Policy Assignment (C) allows you to immediately assign the endpoint to a specific policy set. This directly links the endpoint to a set of authentication, authorization, and accounting rules, defining its network access behavior. This is a key aspect of endpoint management.

Simultaneously, Identity Group Assignment (D) enables you to place the endpoint into a logical group for easier management and policy targeting. This facilitates grouping devices based on function, department, or other criteria, simplifying policy application on a larger scale. These groupings are integral for effective access control.

Options A (Security Group Tag) and B (Endpoint Family) are not directly available during creation of an endpoint on the Context Visibility page. Security Group Tags are applied as part of the authorization profile derived from policy application, and Endpoint Family is an attribute automatically discovered or added through profiling activities. Both are associated with an endpoint after initial identification, not during manual creation. Option E (IP Address), while fundamental to endpoint connectivity, isn't a direct field you set when manually creating an endpoint entry; instead, the system uses this information from profiling and network observations.

Essentially, manually creating an endpoint via the Context Visibility page primarily concerns its identity and how it's handled within ISE policies; it provides a mechanism for pre-defining endpoint attributes before the endpoint actively requests network access. Therefore, policy assignments and identity group placements are the most relevant.

Authoritative links:

Cisco ISE Admin Guide: You can often find documentation on endpoint creation within the Cisco Identity Services Engine (ISE) Administration guides. Here's a generic link, but ensure to select the guide version corresponding to your ISE version: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html Search the document for terms like "context visibility," "endpoint creation," and "manual endpoint."
Cisco Learning Network Community:https://community.cisco.com/t5/security/bd-p/5681j-disc-security -Search this community forum for ISE topics related to endpoint management and context visibility for discussions and more details.

Question: 25

In which two ways can users and endpoints be classified for TrustSec? (Choose two.)

A.VLAN
B.dynamic
C.QoS
D.SGACL
E.SXP

Answer: AB

Explanation:

The correct answer is A and B: VLAN and dynamic classification. TrustSec, Cisco's security architecture, leverages Security Group Tags (SGTs) to enforce policy. Users and endpoints can be classified using various methods to assign these SGTs. VLAN is a foundational method; devices on the same VLAN often have similar security requirements, making it a practical classification point. For example, all devices in the 'Guest' VLAN might be assigned a specific SGT that restricts their network access. Dynamic classification is another key approach. It involves using the Identity Services Engine (ISE) to determine the SGT based on attributes like username, endpoint type, or device posture, as it joins the network, and can be enforced through various authentication methods such as 802.1X. This dynamic capability allows for more granular and context-aware policy enforcement than relying solely on static parameters like VLANs. Quality of Service (QoS) is about prioritizing traffic, not classifying endpoints for security purposes, and while it could be used alongside TrustSec, it is not a basis for classification for the sake of applying SGT. Similarly, SXP is a protocol for SGT propagation, not a means of initial endpoint classification itself. It ensures that SGT information is shared across network devices. You should not confuse an action or propagation method with the means of initial classification. Therefore, VLAN and dynamic classification through ISE are the primary ways TrustSec classifies endpoints and users for SGT assignment.

Supporting Links:

1. Cisco TrustSec: https://www.cisco.com/c/en/us/solutions/enterprise-networks/trustsec/index.html 2. Cisco Identity Services Engine (ISE): https://www.cisco.com/c/en/us/products/security/identity-services-engine/index.html
3. TrustSec SGTs: https://www.cisco.com/c/en/us/td/docs/switches/lan/trustsec/configuration/guide/b_tsec_cg/b_tsec_cg_c

Question: 26

When configuring Active Directory groups, what does the Cisco ISE use to resolve ambiguous group names?

A. MIB
B. SID
C. MAB
D. TGT

Answer: B

Explanation:

When integrating with Active Directory, Cisco ISE needs a reliable way to uniquely identify groups, especially when similarly named groups might exist in different organizational units or domains. This is where the Security Identifier (SID) comes into play. SIDs are unique, immutable identifiers assigned to every security principal (users, groups, computers) within an Active Directory environment. Unlike group names, which can

be duplicated, SIDs are guaranteed to be unique across the entire Active Directory forest. Cisco ISE utilizes the SID to definitively identify the correct group, eliminating any ambiguity caused by duplicate group names. It does not rely on Management Information Base (MIB), which is used for SNMP-based network management, or Media Access Control (MAC) Address Bypass (MAB) which is authentication based on MAC addresses. A Ticket-Granting Ticket (TGT), used in Kerberos authentication, is not directly involved in group name resolution. The SID offers a robust and dependable mechanism for consistent group identification and management within the Cisco ISE environment. It ensures that policies, permissions, and access controls are applied correctly based on the true identity of the Active Directory groups. The reliance on SIDs provides greater reliability in integrating user and group access management.

Authoritative links for further research:

1. Microsoft Documentation on SIDs: https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/manage/understand-security-identifiers
2. Cisco ISE Configuration Guides: Search for "Cisco ISE Active Directory Integration" on Cisco.com

Question: 27

Which permission is common to the Active Directory Join and Leave operations?

A.Remove the Cisco ISE machine account from the domain.

B.Search Active Directory to see if a Cisco ISE machine account already exists.

C.Set attributes on the Cisco ISE machine account.

D.Create a Cisco ISE machine account in the domain if the machine account does not already exist.

Answer: B

Explanation:

The correct answer is **B. Search Active Directory to see if a Cisco ISE machine account already exists.** This is because both joining and leaving an Active Directory domain with Cisco ISE require checking the existence of the ISE machine account within the domain. During the join process, ISE must determine if a machine account already exists to avoid conflicts. Similarly, when leaving, ISE needs to locate its existing machine account to properly remove it from the domain. Options A, C, and D relate to specific actions performed during the join or leave process but aren't common to both. The action of searching for an existing account is fundamental to both operations, ensuring a consistent and safe interaction with Active Directory. Before creating a new account during a join or attempting to remove an account when leaving, the system must first establish if the respective account exists. This avoids duplicate accounts upon join or failed remove operations due to an unknown account. This aligns with principles of identity management within an enterprise environment which prioritizes verification and avoids potential issues. The underlying concept involves directory services interactions, focusing on account presence and management as core functionalities.

For further research, consider the following links:

Cisco ISE Administration Guide:https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html (Search within the guide for Active Directory integration and Domain join/leave procedures.)
Microsoft Active Directory Documentation:https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/active-directory-domain-services (Explore the fundamentals of domain joins and machine account management.)

Question: 28

Which interface-level command is needed to turn on 802.1X authentication?

A. dot1x system-auth-control
B. dot1x pae authenticator
C. aaa server radius dynamic-author
D. authentication host-mode single-host

Answer: B

Explanation:

The correct answer is **B. dot1x pae authenticator**.

802.1X authentication is a port-based network access control mechanism, and the dot1x pae authenticator command is the specific interface-level configuration needed to enable it on a Cisco device. "PAE" stands for Port Access Entity, and specifying "authenticator" mode means the device will act as the authenticator in the 802.1X process. The authenticator initiates the authentication with supplicants (end devices) and forwards credentials to the authentication server (like Cisco ISE).

Option A, dot1x system-auth-control, is a global configuration command, which enables or disables 802.1X globally on the device. It doesn't specify which interfaces will participate in 802.1X authentication. Option C, aaa server radius dynamic-author, configures dynamic authorization and does not enable 802.1X authentication on an interface. Option D, authentication host-mode single-host, sets the host mode for MAC address
authentication and is not specifically for 802.1X.

To implement 802.1X, you need both global configuration (using dot1x system-auth-control, typically), and importantly, the interface-level configuration using dot1x pae authenticator to specify which ports will actively use the 802.1X authentication process. Without the dot1x pae authenticator command, the interface won't be participating in the 802.1X process.

Authoritative Links:

Cisco's Documentation on 802.1X:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_8021x/configuration/15-sy/sec-usr-8021x-15-sy-book/sec-ieee-8021x.html
Cisco's Command Reference for dot1x pae:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/command/sec-cr-all-book/sec-cr-d41.html#wp3834442548

Question: 29

Which RADIUS attribute is used to dynamically assign the Inactivity active timer for MAB users from the Cisco ISE node?

A.session-timeout
B.termination-action
C.radius-server timeout
D.idle-timeout

Answer: D

Explanation:

The correct answer is **D. idle-timeout**. Here's the justification:

Machine Authentication Bypass (MAB) often involves devices that may remain connected without consistent user interaction. To manage these inactive connections efficiently, Cisco ISE uses the idle-timeout RADIUS attribute. This attribute signals to the network device (like a switch or wireless controller) how long a MAB-authenticated session can remain idle before it's terminated. Idle-timeout specifically controls the inactivity period, ensuring resources are not held by dormant connections.

Option A, session-timeout, controls the overall duration of a session, regardless of activity. Option B, termination-action, dictates how a session should be terminated (e.g., logoff, reauthenticate), but not when due to inactivity. Option C, radius-server timeout, usually refers to the timeout setting on the network device for its communication with the RADIUS server (Cisco ISE) and not a session attribute.

Therefore, when dynamically assigning an inactivity timer for MAB users from Cisco ISE, the idle-timeout RADIUS attribute is the precise tool for the task. The ISE dynamically configures the device based on the defined authorization policies, sending the idle-timeout value as a part of the RADIUS response. This allows for fine-grained control over when a MAB-authenticated session, lacking activity, should be terminated, enhancing security and resource management. Using this attribute ensures that resources associated with an idle connection are released to be used by other users.For further research and understanding, you can consult:

1. Cisco ISE Configuration Guides: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html
2. Understanding RADIUS Attributes: https://tools.ietf.org/html/rfc2865 - (While not specific to ISE, it details standard RADIUS attributes).
3. Cisco ISE documentation on MAB and authorization policies.

Question: 30

What does the dot1x system-auth-control command do?

A. globally enables 802.1x
B. causes a network access switch not to track 802.1x sessions
C. enables 802.1x on a network access device interface
D. causes a network access switch to track 802.1x sessions

Answer: A

Explanation:

The dot1x system-auth-control command globally enables 802.1X authentication on a Cisco network device. It essentially activates the 802.1X functionality across the device, making it the default method for port authentication. Without this global command, 802.1X features configured on individual interfaces would not function. The command prepares the device to participate in the 802.1X authentication process, which includes communication with an authentication server, such as Cisco ISE. When a device connects to an interface with 802.1X enabled, the device's identity is verified against the configured authentication policies. The command doesn't cause the switch not to track sessions; in fact, it's the opposite – it enables the tracking necessary for 802.1X operations. It also does not enable 802.1x on a specific interface; that's done with different interface commands. This global command is crucial to initiate a secure and controlled access environment. The command acts as an on/off switch, enabling the foundation for all subsequent 802.1x configurations on that device. Therefore, option A correctly describes the impact of the command. Further details on the command and its function in a Cisco network can be found in Cisco's official documentation on 802.1X and related commands.

Here are authoritative links for further research:

Cisco's 802.1X Configuration Guide:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_8021x/configuration/15-mt/sec-usr-8021x-15-mt-book/sec-ieee-8021x.html
Cisco Command Reference for dot1x system-auth-control:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/command/sec-cr-book/sec-d1.html#wp1344113

Question: 31

What should be configured on the Cisco ISE authentication policy for unknown MAC addresses/identities for successful authentication?

A.continue
B.pass
C.drop
D.reject

Answer: A

Explanation:

The correct action for unknown MAC addresses/identities in Cisco ISE's authentication policy is "Continue".
This setting directs ISE to proceed to the next policy rule rather than halting the authentication process. In scenarios where a MAC address or user identity is not specifically defined in an identity source, the 'Continue' option avoids immediate rejection. This allows the ISE system to evaluate subsequent authentication rules, potentially matching the unknown endpoint based on other attributes, for instance, posture or profiling policies, before final authentication decisions. "Pass" isn't appropriate here as it would signal successful authentication without proper identification, creating security risks. "Drop" and "Reject" would both terminate the authentication process, which is undesirable since we want the system to explore other possible matching rules. Using "Continue" allows for a more flexible, multi-layered approach, where multiple policy rules are evaluated before an authentication decision is reached, enhancing security while accommodating diverse endpoint configurations. This approach ensures that legitimate, yet not previously known endpoints can still potentially be authorized based on other policies. Failure to use 'Continue' may unnecessarily block legitimate devices that have not been previously onboarded. Proper use of 'Continue' with other policy configurations allows a better balance of flexibility and security. For more information, see the Cisco documentation on ISE authentication policies here: https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/admin_guide/b_ise_admin_guide_27/b_ise_admin_guide_27_chapter_0100.html and https://www.cisco.com/c/en/us/td/docs/security/ise/2-6/admin_guide/b_ise_admin_guide_26/b_ise_admin_guide_26_chapter_0110.html

Question: 32

Which command displays all 802.1X/MAB sessions that are active on the switch ports of a Cisco Catalyst switch?

A.show authentication sessions interface Gi1/0/x output
B.show authentication sessions
C.show authentication sessions output
D.show authentication sessions interface Gi 1/0/x

Answer: B

Explanation:

The correct command to display all active 802.1X/MAB sessions on a Cisco Catalyst switch is show authentication sessions. This command provides a comprehensive overview of all active authentication sessions regardless of the interface or authentication method (802.1X or MAB). Option A, show authentication sessions interface Gi1/0/x output, would only display authentication sessions on a specific interface (Gi1/0/x). Option C, show authentication sessions output, is not a valid command, as there's no such option. Option D, show authentication sessions interface Gi 1/0/x, attempts to specify an interface but includes a space in the interface name which is incorrect syntax. The show authentication sessions command, without any interface specification, shows a summary of all active sessions, including their associated user, authentication method, and VLAN information. This is crucial for network administrators to monitor and troubleshoot network access control issues. This command is foundational for managing and securing network access leveraging 802.1X and MAB protocols. It gives an aggregate view of the device's current authentication status.

For further information, you can refer to the Cisco documentation:

Cisco Command Reference Guide:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_data_8021x/command/sec-8021x-cr-book/sec-8021x-cr_show.html
Cisco Identity Services Engine (ISE) Documentation:
https://www.cisco.com/c/en/us/support/security/identity-services-engine/tsd-products-support-series-home.html

Question: 33

What are two requirements of generating a single certificate in Cisco ISE by using a certificate provisioning portal, without generating a certificate signing request?
(Choose two.)

A.Enter the IP address of the device.
B.Enter the common name.
C.Choose the hashing method.
D.Locate the CSV file for the device MAC.
E.Select the certificate template.

Answer: BE

Explanation:

Let's break down why options B and E are the correct answers for generating a single certificate in Cisco ISE using the provisioning portal without a CSR.

When a certificate is generated without a Certificate Signing Request (CSR), the Cisco ISE server takes on the responsibility of creating the key pair. The common name (B) becomes crucial because it's a primary identifier within the certificate, typically reflecting the device's hostname or fully qualified domain name (FQDN). This common name is essential for secure communication, ensuring that the certificate is used by the intended device.

The certificate template (E) is also necessary because it dictates the purpose and parameters of the certificate. Cisco ISE uses templates to define the certificate's validity period, key usage extensions, and other essential characteristics. By choosing a template, you instruct ISE on how the generated certificate should be structured and what functions it is permitted to perform.

Option A (Enter the IP address of the device) is incorrect. While devices need IP addresses for network connectivity, the IP address is not a fundamental part of a certificate's identity. Certificates identify entities (like devices) using names, and the IP address can change dynamically. Option C (Choose the hashing method) is not required at this stage. Hashing algorithms are used in the CSR generation process, which is explicitly

excluded in this scenario. Option D (Locate the CSV file for the device MAC) is related to bulk certificate enrollment processes and not relevant for generating a single certificate without a CSR.

In summary, because we are generating the certificate directly, ISE needs a common name to identify the certificate target and a template to define how the certificate should be structured. These are the basic elements of certificate generation without requiring an intermediary CSR.

Supporting Links:

1. **Cisco ISE Certificate Management:**https://www.cisco.com/c/en/us/td/docs/security/ise/3-1/admin_guide/b_ise_admin_3_1/m_certs.html (This Cisco documentation provides information on the certificate management in ISE).
2. **Understanding Certificate Templates:**https://www.pkisolutions.com/certificate-templates-explained/ (This external site gives a broader understanding of certificate templates and their use cases).

Question: 34

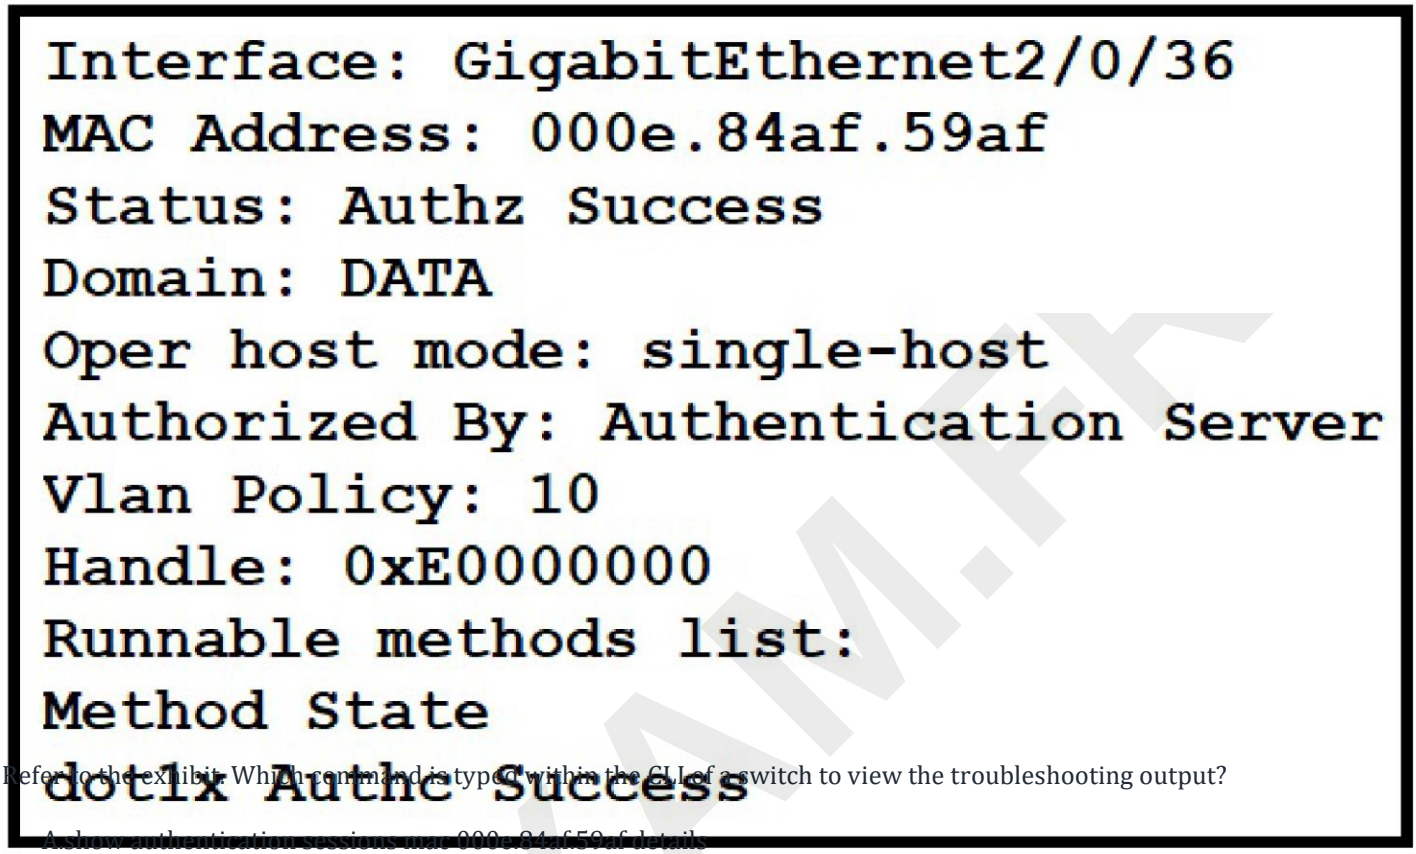

Refer to the exhibit. Which command is typed within the CLI of a switch to view the troubleshooting output?

A.show authentication sessions mac 000e.84af.59af details

B.show authentication registrations

C.show authentication interface gigabitethernet2/0/36

D.show authentication sessions method

Answer: A

Explanation:

A. This is the correct answer. Output looks the same on my 3650.B. show authentication registrations = Shows clients registered with Session ManagerC. show authentication interface ? = Invalid syntax ?D. show

authentication sessions method = Incomplete Command

Question: 35

What gives Cisco ISE an option to scan endpoints for vulnerabilities?

A.authentication policy
B.authorization profile
C.authentication profile
D.authorization policy

Answer: B

Explanation:

The correct answer is **B. authorization profile**. Cisco ISE uses authorization profiles to define the permissions and actions granted to an endpoint after successful authentication. Within an authorization profile, you can specify various attributes and settings that dictate how a device interacts with the network. One crucial feature is the ability to trigger posture assessments which can include vulnerability scans. While
authentication profiles (C) manage the actual authentication process, they don't directly dictate post-authentication actions like vulnerability checks. Similarly, authentication policies (A) define the conditions that must be met for authentication to succeed, and authorization policies (D) determine which authorization profile is applied based on endpoint characteristics and identity. Vulnerability scanning is a post-
authentication activity governed by the applied authorization profile. The profile can include settings that instruct ISE to initiate a vulnerability assessment using a third-party integration, or the ISE's built-in posture service, based on the endpoint type, user role, and other contextual data defined within that authorization profile. This posture assessment ensures that only secure and compliant devices can gain access to network resources, enhancing the organization's security posture.

For further research, consult the following Cisco documentation:

Cisco ISE Configuration Guides: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-configuration-examples-list.html
Cisco ISE Admin Guide: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html (Specifically, sections related to authorization policies, profiles and posture assessment)

Question: 36

Which two values are compared by the binary comparison function in authentication that is based on Active Directory?

A.user-presented certificate and a certificate stored in Active Directory
B.MS-CHAPv2 provided machine credentials and credentials stored in Active Directory
C.user-presented password hash and a hash stored in Active Directory
D.subject alternative name and the common name

Answer: A

Explanation:

The correct answer is **A: user-presented certificate and a certificate stored in Active Directory.**

Let's break down why. Cisco ISE, when authenticating against Active Directory using certificate-based authentication, primarily relies on a binary comparison. This means it directly compares two data structures to determine equality. In this context, the core comparison involves the digital certificate presented by the user or device seeking access against the corresponding certificate record held within Active Directory. This process checks that the presented certificate is identical to the one registered in AD for the authenticating entity.

Active Directory stores user and device certificates, making it the trusted repository for verifying identities. When a user presents a certificate, ISE retrieves the associated certificate from Active Directory based on the user's identity. The binary comparison function then performs a byte-by-byte comparison between the presented certificate and the stored certificate. If they match, the authentication succeeds, assuming other necessary checks (like revocation) pass.

Option B, while involving MS-CHAPv2 (which is a password-based protocol) and AD, isn't based on binary comparison. Passwords are hashed, not compared byte-by-byte. Option C is similar – passwords are not compared directly as a binary value, but rather as the result of a hashing function. Option D, involving subject alternative name and common name, is generally a string comparison, not a binary comparison.

The "binary" aspect emphasizes the exact matching required for cryptographic material like certificates, contrasting with the more abstract, one-way nature of hash comparisons or the flexible nature of string matches. Therefore, the direct comparison of certificate data in its binary form is the core of this
authentication method. This approach provides strong security by relying on the unforgeable nature of digital certificates.

Further reading:

Cisco ISE Identity Management: https://www.cisco.com/c/en/us/td/docs/security/ise/3-1/admin_guide/b_ise_admin_3_1/m_id_mgmt.html
Active Directory Certificate Services: https://learn.microsoft.com/en-us/windows-server/networking/core-network-guide/cncg/server-certs/ad-cs
Certificate-Based Authentication with ISE: https://www.cisco.com/c/en/us/support/docs/security/identity-services-engine/214979-ise-certificate-based-authentication-for.html

Question: 37

What happens when an internal user is configured with an external identity store for authentication, but an engineer uses the Cisco ISE admin portal to select an internal identity store as the identity source?

A. Authentication is redirected to the internal identity source.
B. Authentication is granted.
C. Authentication fails.
D. Authentication is redirected to the external identity source.

Answer: C

Explanation:

Okay, let's break down why the correct answer is C, "Authentication fails," in the context of Cisco ISE and identity store configurations.

Here's the justification:

The core of the issue lies in the mismatch between the user's configured identity source and the identity source actively selected during authentication. When an internal user is set up to authenticate against an external identity store (like Active Directory or LDAP), Cisco ISE expects credentials to be validated against

that specific external database. However, when an engineer, via the ISE admin portal, overrides this default behavior by selecting an internal identity store (which stores ISE-specific administrative users and sometimes a limited set of internal users) as the authentication source, a significant conflict arises.

ISE will attempt to locate the user's credentials within the internal database. Since the user's information is not present in this database (it is stored in the configured external identity store), the authentication process will fail. It doesn't redirect the authentication attempt to the external store automatically simply because it's configured; ISE follows the source specified during the authentication request. This failure ensures that ISE adheres to the authentication policies defined and prevents unauthorized access attempts via incorrect identity sources. A successful authentication always requires that the chosen identity source for
authentication is configured and that the authentication information matches. This is fundamental to access control and network security policies enforced by systems like Cisco ISE.

The intended behavior of ISE is not to fall back to the user's configured external store; instead, it will strictly evaluate using only the identity source that has been explicitly selected or specified for the current authentication attempt. Therefore, a mismatch between the user's intended authentication source and the configured source as per the user setup, results in authentication failure.

Further research on Cisco ISE Identity stores can be found at these official Cisco resources:

Cisco Identity Services Engine (ISE) Administrator Guide:
https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html (Search for "Identity Stores" within the relevant version's guide)
Cisco ISE Configuration Best Practices:https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-technical-reference-list.html (Look for configuration guides and best practices documents regarding identity management)

Question: 38

Which two actions occur when a Cisco ISE server device administrator logs in to a device? (Choose two.)

A.The Cisco ISE server queries the internal identity store.
B.The device queries the external identity store.
C.The device queries the Cisco ISE authorization server.
D.The device queries the internal identity store.
E.The Cisco ISE server queries the external identity store.

Answer: AE

Explanation:

The correct answer is AE. When a device administrator logs into a Cisco ISE (Identity Services Engine) server, the authentication process involves several steps. First, the administrator initiates the login attempt through the ISE interface. Cisco ISE, acting as an identity management system, needs to verify the administrator's credentials. To do this, it directly queries its internal identity store (Option A). This internal store houses the authorized administrator accounts within the ISE system. Simultaneously, Cisco ISE may need to check for group memberships or other attributes that could exist in external identity stores such as Active Directory or LDAP. This is where the external identity store is queried (Option E). The administrator's profile might include information stored outside of the ISE's internal database. Options B, C, and D are incorrect because the device, in this context, refers to the ISE server itself, and it does not query an external store or the authorization server for device administrator login; the ISE performs these actions. The authorization step within ISE follows successful authentication. In essence, ISE uses its internal store to confirm the username and password combination of an authorized administrator and uses external stores to retrieve associated

attributes. This dual-pronged approach ensures comprehensive verification and access control.

Here are some authoritative links for further research on Cisco ISE and identity management concepts:

1. **Cisco Identity Services Engine (ISE) - Cisco:** https://www.cisco.com/c/en/us/products/security/identity-services-engine/index.html (Official Cisco page on ISE.)
2. **Cisco ISE Deployment Models - Cisco Community:**https://community.cisco.com/t5/security-documents/cisco-ise-deployment-models/ta-p/3128429 (Explains different ISE deployment scenarios)
3. **Identity Management Systems - NIST:** https://csrc.nist.gov/glossary/term/identity_management_system (NIST definition of identity management systems)

Question: 39

An engineer is configuring a guest password policy and needs to ensure that the password complexity requirements are set to mitigate brute force attacks.
Which two requirements should be included in this policy? (Choose two.)

A. active username limit
B. password expiration period
C. access code control
D. username expiration date
E. minimum password length

Answer: BE

Explanation:

The correct answer is B and E. A password expiration period (B) and minimum password length (E) are crucial for mitigating brute-force attacks against guest accounts in Cisco Identity Services Engine (ISE). A password expiration policy forces users to change their passwords regularly, which reduces the window of vulnerability should a password be compromised via brute-force or other means. It ensures that even if a password becomes known, it won't be valid for an extended duration. A minimum password length (E) directly combats brute-force attempts by increasing the computational effort required to guess the password. Longer passwords have a much larger potential keyspace, making it exponentially more difficult for an attacker to try every possible combination. Active username limit (A) and username expiration date (D) are more related to account management than password complexity. Access code control (C), while relevant for specific authentication methods, is not a direct password complexity requirement. Password complexity requirements should enforce sufficient entropy, which directly relates to the randomness and length of a password. Both password expiration and length are standard security practices for any system and are vital to protecting accounts. These settings force attackers to put in more effort for brute forcing to work against them, leading to a better posture.

For further information on password complexity policies and best practices, refer to resources like the National Institute of Standards and Technology (NIST) Special Publication 800-63B, which details authentication and password guidelines: https://pages.nist.gov/800-63-3/sp800-63b.htmlAdditionally, Cisco documentation on ISE guest services can provide more detail about specific features and configurations: https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/admin_guide/b_ISE_admin_3_2/m_guest_services.html.

Question: 40

An engineer is using the low-impact mode for a phased deployment of Cisco ISE and is trying to connect to the network prior to authentication.
Which access will be denied in this deployment?

A.DNS
B.DHCP
C.EAP
D.HTTP

Answer: D

Explanation:

The correct answer is **D. HTTP**.

In a low-impact mode deployment of Cisco ISE, the primary goal is to gradually introduce the ISE system into an existing network without causing widespread disruptions. This phased approach typically starts by monitoring network activity without actively enforcing authentication or authorization initially. When a client connects to the network, it is allowed basic network access, meaning fundamental services essential for communication like DNS, DHCP, and EAP (for wired 802.1X or wireless) are permitted. However, access to higher-level services, especially applications running on the HTTP protocol, are generally blocked until the user is authenticated with ISE. This allows the engineer to see which devices would be authenticated and how they would connect, and also confirms that DNS and DHCP are working and do not need to be fixed before the policy is changed. This is crucial in identifying potential issues beforehand, allowing for smoother migration to full ISE functionality. EAP is passed to ISE for initial monitoring and reporting. By blocking HTTP traffic, the deployment does not impact the network initially, but will cause any web applications to fail without authentication. It also serves as an indicator for the user that the authentication process is required before access is fully granted.

Authoritative Links for further research:

Cisco Identity Services Engine (ISE) Documentation:
https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html (Specific installation guides and deployment examples will cover low-impact mode.)

Cisco ISE Deployment Best Practices:https://www.cisco.com/c/en/us/td/docs/security/ise/2-4/admin_guide/b_ise_admin_guide_24/b_ise_admin_guide_24_chapter_0101.html (Look for sections discussing phased deployments and monitoring.)

Question: 41

An administrator needs to connect ISE to Active Directory as an external authentication source and allow the proper ports through the firewall.
Which two ports should be opened to accomplish this task? (Choose two.)

A.TELNET: 23
B.HTTPS: 443
C.HTTP: 80
D.LDAP: 389
E.MSRPC:445

Answer: DE

Explanation:

Here's a detailed justification for why the correct answer to connecting Cisco ISE to Active Directory using specific ports is D (LDAP: 389) and E (MSRPC: 445):

Cisco ISE (Identity Services Engine) relies on directory services like Active Directory (AD) for user authentication and authorization. To interact with AD, ISE needs to communicate over specific network ports.

LDAP (Lightweight Directory Access Protocol) is the primary protocol for querying and retrieving user and group information from Active Directory. By default, LDAP operates on port 389. This allows ISE to look up user credentials and group memberships when a user attempts to authenticate.

MSRPC (Microsoft Remote Procedure Call) is also essential for ISE's integration with Active Directory. MSRPC uses port 445 (amongst others) and enables secure communication between ISE and Active Directory for tasks like user/computer authentication and retrieving security tokens (Kerberos and NTLM). Although LDAP facilitates authentication, often, an underlying service like MSRPC is used to facilitate the actual authentication process.

Therefore, opening ports 389 for LDAP and 445 for MSRPC on the firewall allows for a successful and secure integration between Cisco ISE and Active Directory. Ports 23 (Telnet), 443 (HTTPS), and 80 (HTTP) are not directly used for the primary purpose of integrating ISE with AD for user authentication. Telnet is an unsecure protocol and not recommended while HTTPS is for web access and HTTP is for unsecured web access neither of which directly connect to Active Directory for user authentication.

Authoritative Links:

Cisco ISE Active Directory Integration Guide: Consult Cisco's official documentation for the latest specific port requirements. While these can vary slightly between ISE versions, 389 and 445 are consistent for core AD integration. (Search on Cisco website using these keywords)
Microsoft Documentation on Active Directory Ports: Microsoft provides detailed lists of ports used by Active Directory for various services (search for "Active Directory Port Requirements Microsoft" on your preferred search engine). This documentation supports the importance of LDAP and MSRPC ports for this type of integration.
LDAP Protocol Documentation: Research on official LDAP standards will confirm port 389 as the common default for LDAP communications. (e.g. RFC 4511)
Microsoft RPC Documentation: Research MSRPC and its port utilization, specifically regarding Active Directory functions. (search for "Microsoft RPC Documentation").

Question: 42

An engineer is implementing Cisco ISE and needs to configure 802.1X. The port settings are configured for port-based authentication.
Which command should be used to complete this configuration?

A.aaa authentication dot1x default group radius
B.dot1x system-auth-control
C.authentication port-control auto
D.dot1x pae authenticator

Answer: B

Explanation:

The correct command is **B. dot1x system-auth-control**. This command globally enables 802.1X authentication

on the switch. While options A, C, and D are related to 802.1X, they don't accomplish the fundamental step of turning on the 802.1X process on the device. Specifically:

Option A aaa authentication dot1x default group radius: This command specifies that the RADIUS server should be used for 802.1X authentication. It's a necessary step, but it's secondary to enabling 802.1X in the first place.

Option C authentication port-control auto: This command is used to set a port's authentication mode to "auto," allowing 802.1X, MAC authentication bypass (MAB), and web authentication. It configures a port-specific behavior, not global activation.

Option D dot1x pae authenticator: This command configures the port's Port Access Entity (PAE) role as an authenticator for 802.1X, meaning the port will request authentication. This step is done at a port-by-port level, not at the system wide level, and it is also a secondary step.

The dot1x system-auth-control command is the pivotal starting point, enabling the overall 802.1X framework on the Cisco switch. Once enabled, further configurations like RADIUS server details and port-specific settings can be implemented. Without dot1x system-auth-control, 802.1X functionalities would not operate regardless of other configuration.

For further research, consult the official Cisco documentation on 802.1X configuration:

Cisco Identity Services Engine (ISE) Configuration Guides

Cisco IOS Security Configuration Guide (Search for "802.1X Authentication")

Question: 43

DRAG DROP -
An organization wants to implement 802.1X and is debating whether to use PEAP-MSCHAPv2 or PEAP-EAP-TLS for authentication. Drag the characteristics on the left to the corresponding protocol on the right.
Select and Place:

uses username and password for authentication

uses certificates for authentication

changes credentials through the admin portal

supports fragmentation after the tunnel is established

uses the X.509 format

supports auto-enrollment for obtaining credentials

PEAP-MSCHAPv2

PEAP-EAP-TLS

Answer:

PEAP-MSCHAPv2

- uses username and password for authentication
- changes credentials through the admin portal
- supports fragmentation after the tunnel is established

PEAP-EAP-TLS

- uses certificates for authentication
- uses the X.509 format
- supports auto-enrollment for obtaining credentials

Question: 44

DRAG DROP -
Drag the descriptions on the left onto the components of 802.1X on the right.
Select and Place:

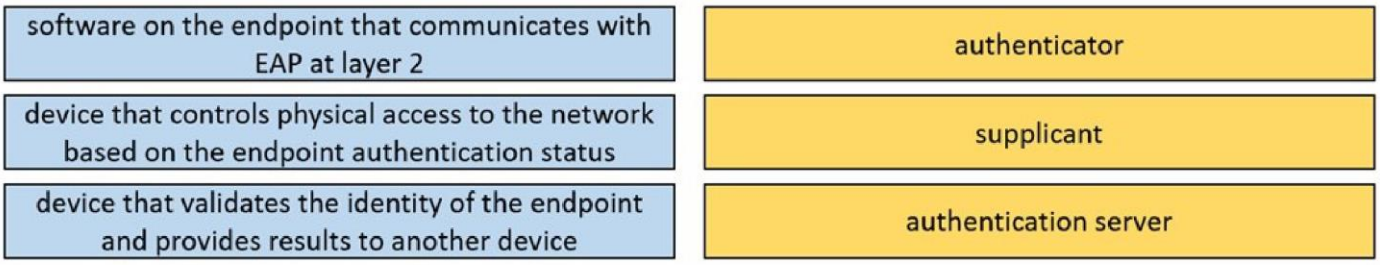

Answer:

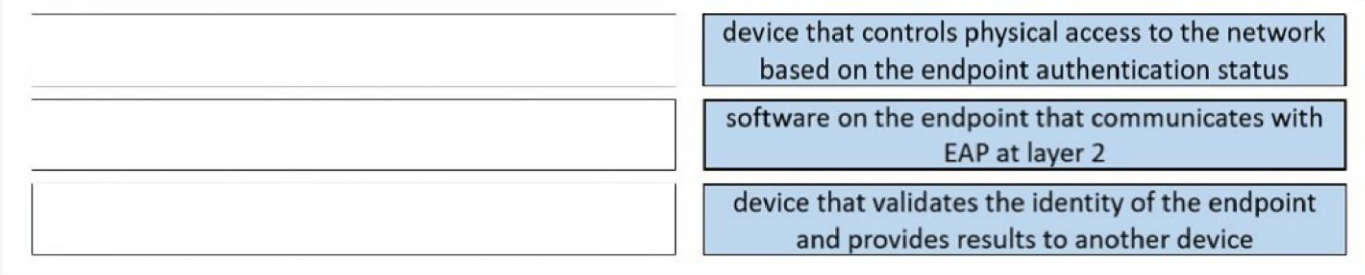

Explanation:
Authenticator " device that controls physical access to the network based on the authentication status Supplicant - software on the endpoint that communicates with EAP at layer 2
Authentication server " device that validates the identity of the endpoint and provides results to another device

Reference:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_8021x/configuration/xe-3se/3850/sec-user-8021 x-xe-3se-3850-book/config-ieee-802x- pba.html

Question: 45

A network engineer is configuring Cisco TrustSec and needs to ensure that the Security Group Tag is being transmitted between two devices.
Where in the Layer 2 frame should this be verified?

A.payload
B.802.1 AE header
C.CMD field
D.802.1Q field

Answer: C

Explanation:

The correct answer is **C. CMD field**.

Security Group Tags (SGTs) in Cisco TrustSec are embedded within the Cisco Meta Data (CMD) field of the Layer 2 Ethernet frame. This CMD field is not a standard part of the 802.1Q or 802.1AE headers, nor is it within the data payload. Cisco's proprietary CMD field resides within a specific vendor-defined Type-Length-Value (TLV) structure, also known as the SGT Exchange Protocol (SXP) TLV, which is carried as part of the CDPv2

(Cisco Discovery Protocol version 2) messaging between network devices or within a dedicated TrustSec-enabled Layer 2 frame.

The 802.1Q header, on the other hand, is used for VLAN tagging, and the 802.1AE header is used for MACsec encryption, neither of which are related to SGT transport. The payload carries the actual data being transmitted and will not contain the SGT metadata directly. Therefore, to verify that the SGT is being transmitted, the network engineer must inspect the CMD field within the Ethernet frame. This field contains the SGT as well as other control plane information related to TrustSec. Specifically, they should look for the SXP TLV within the CMD field of CDPv2 or TrustSec enabled frames.

Authoritative Links for Further Research:

Cisco TrustSec White Paper:
https://www.cisco.com/c/en/us/solutions/collateral/enterprise/trustsec/whitepaper_c11-537221.html **Cisco Security Group Tag Exchange Protocol (SXP) Explained:**https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_secure_group_tag/configuration/15-mt/sec-secure-group-tag-15-mt-book/sec-sxp.html **Cisco Implementing TrustSec at Layer 2:**
https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Security/TrustSec_L2.pdf

Question: 46

A network administrator must configure endpoints using an 802.1X authentication method with EAP identity certificates that are provided by the Cisco ISE. When the endpoint presents the identity certificate to Cisco ISE to validate the certificate, endpoints must be authorized to connect to the network.
Which EAP type must be configured by the network administrator to complete this task?

A. EAP-TTLS
B. EAP-TLS
C. EAP-FAST
D. EAP-PEAP-MSCHAPv2

Answer: B

Explanation:

The correct EAP type for authenticating endpoints using identity certificates provided by Cisco ISE is **EAP-TLS**. EAP-TLS (Extensible Authentication Protocol-Transport Layer Security) relies on digital certificates for mutual authentication between the client and the authentication server. In this scenario, the endpoints present their certificates, issued by Cisco ISE, to verify their identity. The ISE, configured as the
authentication server, then validates these certificates to authorize network access. This mechanism provides strong authentication as it leverages cryptographic keys within the certificates. EAP-TLS is highly secure because it does not rely on usernames and passwords and mitigates the risk of credential theft. Unlike EAP-TTLS and EAP-PEAP, which use a secured tunnel to transmit less secure authentication methods, EAP-TLS uses certificate-based mutual authentication for both the server and the client. EAP-FAST also relies on a secure tunnel but uses a Protected Access Credential (PAC) instead of certificates for initial authentication and subsequent key establishment. Therefore, given the requirement for certificate-based authentication, EAP-TLS is the appropriate solution.

For more information, you can refer to the following resources:

1. **Cisco ISE Configuration Guide**: Look for sections discussing EAP-TLS implementation for detailed configuration steps: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html (Select the relevant configuration guide based on your ISE version)
2. **Understanding EAP Types**: Search for articles or documentation explaining the different EAP

methods like EAP-TLS, EAP-TTLS, EAP-FAST, and EAP-PEAP. A good starting point is: https://en.wikipedia.org/wiki/Extensible_Authentication_Protocol

3. **Cisco Documentation on EAP-TLS**: Search specifically for EAP-TLS configuration examples and best practices using Cisco products. You can find relevant documentation in the Cisco's documentation portal for ISE.

Question: 47

An organization wants to standardize the 802.1X configuration on their switches and remove static ACLs on the switch ports while allowing Cisco ISE to communicate to the switch what access to provide.
What must be configured to accomplish this task?

A. dynamic access list within the authorization profile
B. extended access-list on the switch for the client
C. security group tag within the authorization policy
D. port security on the switch based on the client's information

Answer: A

Explanation:

The correct answer is **A. dynamic access list within the authorization profile.** Here's why:

The scenario describes a move towards centralized access control using Cisco ISE. Instead of static ACLs on switches, the goal is to have ISE dynamically dictate access policies. Dynamic Access Lists (dACLs) within ISE authorization profiles are the mechanism to achieve this. When a user authenticates via 802.1X, ISE, based on configured policies, can dynamically push down an ACL to the switch port, allowing granular control. This eliminates the need for pre-configured ACLs on every port, offering flexibility and centralized management.

Option B is incorrect because extended access lists on the switch contradict the requirement to remove static ACLs and delegate control to ISE. Option C, Security Group Tags (SGTs), are used for policy enforcement throughout the network fabric but not for port-level access control at the switch level in this specific scenario.

Option D, port security, although related to network access, focuses on limiting which MAC addresses can connect to a port, not on providing specific user-based access control like dACLs.

dACLs facilitate a scalable and manageable approach to network access, aligning with the core principles of identity and access management in modern networks. They enable a more granular, role-based access approach, enhancing security and operational efficiency compared to static ACLs. By using authorization profiles to control access, organizations can significantly reduce management overhead and streamline network policy changes.

Authoritative links for further research:

Cisco ISE Configuration Guide:https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/admin_guide/b_ise_admin_3_2/m_configure_authorization_policy.html
Dynamic Access Control Lists with Cisco ISE:https://www.cisco.com/c/en/us/support/docs/security/identity-services-engine/200292-Understanding-Dynamic-ACLs-with-ISE.html

Question: 48

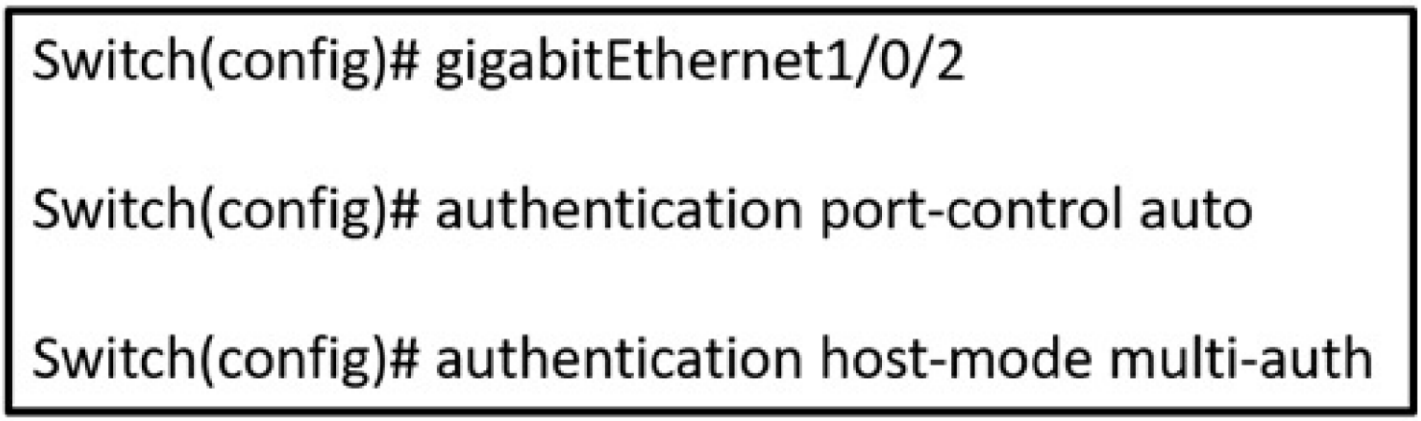

Refer to the exhibit.
In which scenario does this switch configuration apply?

A.when allowing a hub with multiple clients connected
B.when allowing multiple IP phones to be connected
C.when preventing users with hypervisor
D.when bypassing IP phone authentication

Answer: A

Explanation:
Reference;
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_8021x/configuration/xe-3se/5700/sec-user-8021x-xe-3se-5700-book/sec-ieee-802x-multi- auth.html

Question: 49

Refer to the exhibit.
Which switch configuration change will allow only one voice and one data endpoint on each port?

A. auto to manual
B. mab to dot1x
C. multi-auth to multi-domain
D. multi-auth to single-auth

Answer: C

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/16-6/configuration_guid

e/sec/b_166_sec_9300_cg/ configuring_ieee_802_1x_port_based_authentication.html

Question: 50

A network security engineer needs to configure 802.1X port authentication to allow a single host to be authenticated for data and another single host to be authenticated for voice.
Which command should the engineer run on the interface to accomplish this goal?

A. authentication host-mode multi-domain
B. authentication host-mode single-host
C. authentication host-mode multi-auth
D. authentication host-mode multi-host

Answer: A

Explanation:

The correct answer is **A. authentication host-mode multi-domain**. This mode allows for multiple devices to authenticate on a single port, specifically separating them into distinct domains like data and voice. The multi-domain host-mode enables the interface to support two separate authentication sessions, typically one for a data device (like a computer) and another for a voice device (like an IP phone). This functionality is crucial for network environments where distinct policies or VLANs are applied based on the device type. Single-host (B) only allows one device to authenticate, which contradicts the requirement. Multi-auth (C) while supporting multiple authentications, doesn't distinguish between domains for access. Multi-host (D) also allows multiple hosts but usually operates under the same policy, not dividing between distinct access types. Therefore, multi-domain is the only option that properly segregates devices based on their function within the authentication process. The Cisco documentation clarifies that multi-domain allows distinct authentication sessions for voice and data devices using a single port which aligns directly with the requirement stated. For further exploration of 802.1X and Cisco ISE configurations, please refer to the following links: Cisco Identity Services Engine (ISE) Configuration Guides and Cisco Catalyst Switch Security Configuration Guide. These provide authoritative resources for understanding the details of authentication methods and their implementations on Cisco networking devices.

Question: 51

An administrator connects an HP printer to a dot1x enable port, but the printer is nor accessible. Which feature must the administrator enable to access the printer?

A.change of authorization
B.MAC authentication bypass
C.TACACS authentication
D.RADIUS authentication

Answer: B

Explanation:

The correct answer is **B. MAC Authentication Bypass (MAB)**. Here's why:

Dot1x, a port-based network access control mechanism, typically authenticates users or devices using credentials. However, devices like printers often lack the ability to participate in 802.1X authentication. They lack the necessary supplicant software to handle the EAP (Extensible Authentication Protocol) exchanges. In

such cases, MAB acts as a fallback mechanism. MAB leverages the MAC address of the device as the identifier for authentication. When a device attempts to connect, if 802.1X fails or is absent, the network device (like a switch) can be configured to use the MAC address. The network switch sends a RADIUS request to the Identity Services Engine (ISE) with the device's MAC address. ISE then checks its database for a corresponding entry and if found, allows the device access. This enables non-802.1x devices like printers to obtain network access while maintaining network security using access lists or VLAN assignments via the ISE.

Option A, Change of Authorization (CoA), is a dynamic re-authentication method used to change a user or device's access level after the initial authentication, not the initial authentication method itself. TACACS (Option C) and RADIUS (Option D) are authentication protocols. RADIUS is used by Dot1x and MAB, but not as an alternative authentication method to MAB itself. MAB is the specifically designed feature to handle devices without 802.1x capabilities which makes it the right answer.References:

Cisco's documentation on MAC Authentication Bypass: https://www.cisco.com/c/en/us/td/docs/security/ise/2-6/admin_guide/b_ise_admin_2_6/b_ise_admin_26_chapter_0111.html
Understanding 802.1X: https://www.fortinet.com/resources/cyberglossary/802-1x

Question: 52

When configuring an authorization policy, an administrator cannot see specific Active Directory groups present in their domain to be used as a policy condition.
However, other groups that are in the same domain are seen. What
is causing this issue?

A.Cisco ISE's connection to the AD join point is failing.
B.Cisco ISE only sees the built-in groups, not user created ones.
C.The groups are not added to Cisco ISE under the AD join point.
D.The groups are present but need to be manually typed as conditions.

Answer: C

Explanation:

The correct answer is C: The groups are not added to Cisco ISE under the AD join point. Cisco ISE doesn't automatically discover all Active Directory (AD) groups within a domain. Instead, an administrator must explicitly choose which AD groups are imported into ISE's configuration via the AD join point. This process involves selecting specific groups for ISE to recognize and use in policy conditions. Therefore, if a specific AD group is missing from the list when configuring an authorization policy, it means that group hasn't been included in the groups selected for ISE to monitor during the AD join configuration. Options A, B, and D are not the primary causes. Although a faulty AD connection (A) could disrupt group visibility in general, other groups from the same domain would likely also be missing. ISE does see user-created groups (B); this isn't a limitation. And while you can type conditions manually in some scenarios, this isn't the resolution for missing groups that aren't pre-loaded in the available options (D). Hence, it's necessary to revisit the AD join point configuration in ISE and add the specific missing groups to resolve this visibility issue. This ensures that ISE has the necessary information to enforce access policies correctly based on user group membership. You must specify which groups you want ISE to be aware of.

Cisco ISE Active Directory IntegrationCisco ISE 3.1: Integrate with Microsoft Active Directory - Part 1

Question: 53

An engineer is implementing network access control using Cisco ISE and needs to separate the traffic based on

the network device ID and use the IOS device sensor capability.
Which probe must be used to accomplish this task?

A.NetFlow probe
B.HTTP probe
C.RADIUS probe
D.network scan probe

Answer: C

Explanation:

The correct answer is **C. RADIUS probe**. Here's why:

Cisco ISE relies heavily on the RADIUS protocol for authentication, authorization, and accounting (AAA). The RADIUS probe is the mechanism by which ISE interacts with network devices (like switches and routers) during the authentication process. When a user or device attempts to connect to the network, the network device sends a RADIUS request to ISE. This request includes information about the user/device and the network device itself.

The question specifies the need to separate traffic based on the "network device ID." This ID is typically included within the RADIUS attributes sent by the network device. The IOS device sensor capability, which can collect device-specific data, also integrates with RADIUS. The RADIUS probe intercepts these RADIUS messages allowing ISE to extract the network device's identity, compare it to configured policies, and make access control decisions.

While other probes have their specific functions (like NetFlow for traffic analysis and HTTP for web-based authentication), they don't directly provide the necessary device identification within the authentication workflow. The Network scan probe is used for device discovery but not for per-device authentication during normal access. The RADIUS probe, through the attributes exchanged during authentication, is specifically designed for this type of network access control. Therefore, for identifying and making decisions based on network device ID within the authentication context, the RADIUS probe is the essential component.

Authoritative links:

Cisco Identity Services Engine (ISE) Configuration Guide: https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/admin_guide/b_ise_admin_guide_2_7/b_ise_admin_guide_27_chapter_011.html (Specifically look for "Configure AAA Settings")
Cisco ISE RADIUS Services:https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/admin_guide/b_ise_admin_guide_2_7/b_ise_admin_guide_27_chapter_011.html
Understanding RADIUS Attributes:https://www.cisco.com/c/en/us/support/docs/security/identity-services-engine/118914-config-ise-00.html

Question: 54

What is an advantage of using EAP-TLS over EAP-MS-CHAPv2 for client authentication?

A.EAP-TLS uses a username and password for authentication to enhance security, while EAP-MS-CHAPv2 does not.
B.EAP-TLS uses multiple forms of authentication, while EAP-MS-CHAPv2 only uses one.
C.EAP-TLS uses a device certificate for authentication to enhance security, while EAP-MS-CHAPv2 does not.
D.EAP-TLS secures the exchange of credentials, while EAP-MS-CHAPv2 does not.

Answer: C

Explanation:

The correct answer is **C. EAP-TLS uses a device certificate for authentication to enhance security, while EAP-MS-CHAPv2 does not.** Here's why:

EAP-TLS (Extensible Authentication Protocol - Transport Layer Security) relies on digital certificates for mutual authentication. This means both the client device and the authentication server (like Cisco ISE) present certificates to verify their identities. The client's certificate, often stored on the device itself, provides strong assurance that the connecting device is legitimate. This certificate-based authentication eliminates the reliance on easily compromised usernames and passwords, significantly bolstering security. EAP-TLS establishes a secure, encrypted channel to exchange authentication data, further protecting the process.

On the other hand, EAP-MS-CHAPv2 (Extensible Authentication Protocol - Microsoft Challenge Handshake Authentication Protocol version 2) uses username and password credentials, which are susceptible to attacks like eavesdropping and password guessing. While EAP-MS-CHAPv2 does encrypt the authentication exchange, it doesn't offer the same level of security as certificate-based authentication. The inherent vulnerability of password-based authentication makes EAP-MS-CHAPv2 less secure compared to EAP-TLS.

It's worth noting that while EAP-TLS involves a single authentication factor (the certificate), the use of certificates themselves is a form of multi-factor, as the certificate file and/or private key represents something the client “has” and depending on if a certificate password/pin is used something the user "knows" . Option A is incorrect as EAP-TLS does not use usernames and passwords as primary auth factor, option B is misleading as EAP-TLS is generally considered a single factor authentication method as the client certificate is singular, option D is also incorrect, as both protocols encrypt their exchange of credentials though EAP-TLS is widely regarded as more secure.

Here are some authoritative links for further research:

Cisco: Extensible Authentication Protocol (EAP) Methods:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_usr_auth/configuration/15-mt/sec-usr-auth-15-mt-book/sec-cfg-eap.html
Wikipedia: Extensible Authentication Protocol:
https://en.wikipedia.org/wiki/Extensible_Authentication_Protocol
Microsoft: How does Extensible Authentication Protocol work?:https://learn.microsoft.com/en-us/windows-server/networking/technologies/extensible-authentication-protocol/how-eap-works

Question: 55

What must be configured on the WLC to configure Central Web Authentication using Cisco ISE and a WLC?

A.Use the ip access-group webauth in command.
B.Use the radius-server vsa send authentication command.
C.Set the NAC State option to SNMP NAC.
D.Set the NAC State option to RADIUS NAC.

Answer: D

Explanation:

The correct answer is **D. Set the NAC State option to RADIUS NAC.** This setting is essential for integrating a Wireless LAN Controller (WLC) with Cisco Identity Services Engine (ISE) for Central Web Authentication (CWA). CWA requires the WLC to redirect unauthenticated users to the ISE portal for authentication. The "NAC State" on the WLC determines how it interacts with a Network Access Control (NAC) system. When set to RADIUS NAC, the WLC leverages the RADIUS protocol, which ISE also uses, to send authentication

requests and receive authorization responses. This RADIUS communication allows ISE to enforce policies and redirect users to the CWA portal. Options A and B are not relevant for this specific scenario. The ip access-group webauth command is used for local web authentication, not ISE CWA. The radius-server vsa send authentication command pertains to vendor-specific attributes and doesn't control the overall NAC state.

Option C, SNMP NAC, is for simpler scenarios and doesn't work with ISE's advanced authentication capabilities. The RADIUS NAC setting establishes the foundational link between the WLC and ISE, enabling seamless redirection and policy enforcement during the CWA process. Without this correct configuration, the ISE-WLC integration for CWA will fail. In essence, RADIUS NAC signals the WLC to treat authentication and authorization with ISE, specifically for web authentication, as a RADIUS interaction.

Authoritative Links:

Cisco Identity Services Engine (ISE) Configuration Guides:
https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html (Refer to the WLC integration sections)
Cisco Wireless LAN Controller Configuration Guides:
https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/products-installation-and-configuration-guides-list.html (Search for RADIUS and NAC settings)

Question: 56

A network administrator is configuring authorization policies in Cisco ISE. There is a requirement to use AD group assignments to control access to network resources. After a recent power failure and Cisco ISE rebooting itself, the AD group assignments no longer work.
What is the cause of this issue?

A.The AD join point is no longer connected.
B.The certificate checks are not being conducted.
C.The network devices ports are shut down.
D.The AD DNS response time is slow.

Answer: A

Explanation:

The correct answer is **A. The AD join point is no longer connected.** Here's why:

Cisco ISE relies on a functional connection to Active Directory (AD) to retrieve group membership information for authorization. This connection is established through an "AD join point," which allows ISE to authenticate users against AD and query their group affiliations. When ISE reboots, it needs to re-establish this connection. A power failure and subsequent reboot can disrupt this connection, preventing ISE from communicating with the AD server. Without a valid connection, ISE cannot perform the necessary lookups to determine AD group assignments, rendering the authorization policies ineffective.

Option B, certificate checks, while important for secure communication, are not the primary cause of this issue after a reboot. Certificate validation problems would typically prevent communication entirely. Option C, network device port shutdowns, are irrelevant because the problem occurs within the ISE system and its communication with AD. Option D, slow AD DNS response, can cause performance issues but does not directly lead to a complete failure of AD lookups.

The core issue here is the loss of connectivity between ISE and the AD server, preventing the retrieval of group information necessary for the authorization policies to function correctly. The "AD join point" encapsulates this connection. Without a valid and active join, authorization based on AD groups fails.

Authoritative links:

Cisco ISE Identity Sources:https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/admin_guide/b_ise_admin_guide_2_7/b_ise_admin_guide_27_chapter_011.html - This Cisco document details how to configure and manage identity sources, including Active Directory, in ISE. It explains the importance of a functioning join point.

Cisco ISE Deployment Models:https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/install_guide/b_ise_install_guide_27/b_ise_install_guide_27_chapter_010.html - While this document is primarily about ISE deployment, it provides valuable context for understanding how ISE connects with external systems, including Active Directory.

Question: 57

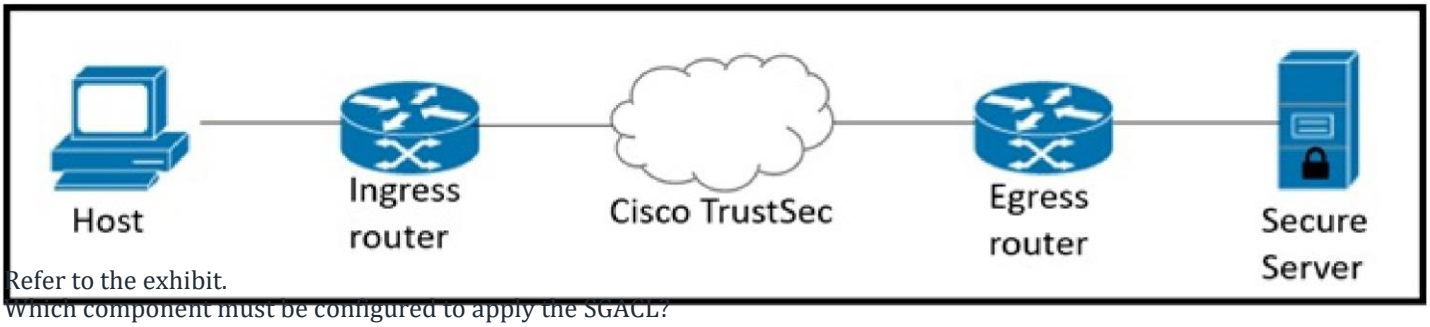

Refer to the exhibit.
Which component must be configured to apply the SGACL?

A. secure server
B. host
C. egress router
D. ingress router

Answer: C

Explanation:

Reference:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SY/configuration/guide/sy_swcg/trustsec.pdf

Question: 58

A laptop was stolen and a network engineer added it to the block list endpoint identity group. What must be done on a new Cisco ISE deployment to redirect the laptop and restrict access?

A. Select DROP under If Auth fail within the authentication policy.
B. Ensure that access to port 8444 is allowed within the ACL.
C. Ensure that access to port 8443 is allowed within the ACL.
D. Select DenyAccess within the authentication policy.

Answer: B

Explanation:

The correct answer is **B. Ensure that access to port 8444 is allowed within the ACL.**

Here's a detailed justification:

When a device is placed in the block list within Cisco ISE, the desired outcome is to restrict its access to the network. This involves a crucial redirection process. The device, upon authentication attempt, will be matched with the block list identity group and then, should trigger the authorization policy redirecting it to a restricted area. Crucially, this redirection often relies on communication with the Cisco ISE server on a specific port.

While port 8443 is primarily used for the ISE admin GUI and API, **port 8444** is the default port for the **ISE Portal**, including guest portals, posture assessment, and importantly, for redirection purposes. To successfully redirect the blocked laptop, the network infrastructure (specifically Access Control Lists or ACLs) must permit communication with the ISE server over port 8444. Without this access, the redirection will fail, and the blocked device might not be effectively contained or might cause issues. Option A, selecting "DROP" on authentication failure, would prevent any kind of action after authentication fail, even redirection. Option C focuses on the admin port (8443), not relevant to end-user redirection, while option D, "DenyAccess," is similar to A, without performing any redirection process. Therefore, permitting access to port 8444 is the sine qua non condition for effective redirection, making option B the only correct option in this scenario.

Authoritative Links for Further Research:

Cisco Identity Services Engine (ISE) Configuration Guides: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html (Refer to the specific guides on portal configuration and redirection policies) **Cisco ISE Ports:**https://www.cisco.com/c/en/us/td/docs/security/ise/3-3/admin_guide/b_ise_admin_3_3/m_ports.html (For detailed information on ports used by Cisco ISE)

Question: 59

When creating a policy within Cisco ISE for network access control, the administrator wants to allow different access restrictions based upon the wireless SSID to which the device is connecting.
Which policy condition must be used in order to accomplish this?

A. Network Access NetworkDeviceName CONTAINS <SSID Name>
B. DEVICE Device Type CONTAINS <SSID Name>
C. Airespace Airespace-Wlan-Id CONTAINS <SSID Name>
D. Radius Called-Station-ID CONTAINS <SSID Name>

Answer: D

Explanation:

The correct answer is **D. Radius Called-Station-ID CONTAINS** . Here's why:

Cisco ISE policies utilize conditions to match incoming authentication requests against configured criteria. These conditions help determine the appropriate authorization and access levels. When dealing with wireless network access, the SSID (Service Set Identifier) becomes a crucial factor for differentiating access needs.

The RADIUS protocol, widely used for network authentication and authorization, carries specific attributes during the authentication process. Specifically, the "Called-Station-ID" attribute in RADIUS encapsulates the SSID that the wireless client is trying to connect to. This attribute is sent from the wireless access point to the RADIUS server (in this case, Cisco ISE). Cisco ISE can then evaluate this "Called-Station-ID" within policy conditions to enforce specific rules.

Option A, Network Access NetworkDeviceName CONTAINS <SSID Name>, is incorrect. The NetworkDeviceName refers to the name of the network device (like the access point) itself and not the SSID. Option B, DEVICE Device Type CONTAINS <SSID Name>, is also incorrect as device type doesn't relate to the SSID. Option C, Airespace Airespace-Wlan-Id CONTAINS <SSID Name>, was used in older versions of Cisco wireless, and is

less relevant with modern configurations. The most standard and universal attribute to examine for SSID differentiation in RADIUS is Called-Station-ID. By configuring a policy using Radius Called-Station-ID CONTAINS <SSID Name>, Cisco ISE can effectively apply different access rules based on the connected wireless network.

For further research, refer to the Cisco ISE documentation on policy conditions, RADIUS attributes, and wireless network integration. A good starting point is: Cisco Identity Services Engine (ISE) Configuration Guides and RFC 2865, RADIUS.

Question: 60

A company manager is hosting a conference. Conference participants must connect to an open guest SSID and only use a preassigned code that they enter into the guest portal prior to gaining access to the network. How should the manager configure Cisco ISE to accomplish this goal?

A. Create logins for each participant to give them sponsored access.
B. Create entries in the guest identity group for all participants.
C. Create an access code to be entered in the AUP mode.
D. Create a registration code to be entered on the portal splash page.

Answer: C

Explanation:

Okay, let's break down why option C is the correct answer for configuring Cisco ISE to provide guest access with a pre-assigned code.

The scenario requires conference attendees to connect to a guest Wi-Fi and then authenticate using a code. Option C, "Create an access code to be entered in the AUP mode," directly addresses this. Here's why:

Cisco ISE (Identity Services Engine) offers various mechanisms for guest access. An "Acceptable Use Policy" (AUP) page is commonly presented to guests after they connect to the guest SSID, but before they are granted full network access. This page can be configured to require an access code for authentication, aligning perfectly with the requirement.

Here's why the other options are less suitable:

A. Create logins for each participant to give them sponsored access: This involves a more complex process where each user needs a dedicated username and password, often involving sponsorship workflows. This is cumbersome for a conference scenario where many guests require quick access.

B. Create entries in the guest identity group for all participants: Manually creating entries in an identity group is impractical for a large group of transient users like conference attendees. It requires advanced preparation and lacks the self-service element required in the scenario.

D. Create a registration code to be entered on the portal splash page: While registration can be used, the question focuses on using a "preassigned code" for access. The "registration" scenario usually requires the user to register their contact information and does not directly satisfy the requirement.

Therefore, leveraging the AUP mode with an access code on Cisco ISE allows the conference manager to generate a single or a set of pre-assigned access codes that can be distributed to participants. When a participant connects to the guest Wi-Fi, they are presented with the AUP page. They enter their code to gain network access after agreeing to the policy, achieving the requirements of the scenario.

Further Research:

Cisco ISE Guest Access Overview:https://www.cisco.com/c/en/us/td/docs/security/ise/2-4/admin_guide/b_ise_admin_2_4/b_ise_admin_24_chapter_01100.html (While this link refers to version 2.4, the core concepts regarding guest access and AUPs are largely consistent across versions.)
Cisco ISE Configuration Guides (Specific Version Needed for Detailed Steps): Search Cisco's documentation for the specific version of ISE in use.

In summary, using access codes within the AUP mode is the most efficient and straightforward method to implement the required functionality.

Question: 61

An organization has a fully distributed Cisco ISE deployment. When implementing probes, an administrator must scan for unknown endpoints to learn the IP-to-
MAC address bindings. The scan is complete on one PSN, but the information is not available on the others. What must be done to make the information available?

A.Cisco ISE must be configured to learn the IP-MAC binding of unknown endpoints via RADIUS authentication, not via scanning.

B.Cisco ISE must learn the IP-MAC binding of unknown endpoints via DHCP profiling, not via scanning.

C.Scanning must be initiated from the MnT node to centrally gather the information.

D.Scanning must be initiated from the PSN that last authenticated the endpoint.

Answer: D

Explanation:

The correct answer is D. In a distributed Cisco ISE deployment, each Policy Service Node (PSN) operates independently concerning active endpoint data. When a PSN performs an endpoint scan to discover IP-to-MAC address bindings, this information is initially stored locally on that specific PSN. It isn't automatically shared or synchronized across all PSNs within the deployment. Therefore, if you want to see the scanned IP-to-MAC binding information on another PSN, the scan needs to be initiated directly by the PSN that last authenticated the endpoint, not just any PSN, because that particular PSN would have relevant context about that endpoint from the authentication process. Scanning initiated by the MnT (Monitoring and
Troubleshooting) node (option C) wouldn't solve the issue because it does not facilitate the required node-specific IP-MAC learning. Similarly, relying only on RADIUS authentication (option A) or DHCP profiling (option B) does not involve the targeted scan function required for active IP-MAC detection across different PSNs.

These methods focus on endpoint categorization and authentication rather than a specific distributed scan across the PSNs. The process ensures that the information is available on the correct PSN, closest to the endpoint's network activity, and maintains scalability by distributing load.

Further research can be done on the following Cisco documentation pages:

Cisco ISE Architecture:https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/ise_deployment/b_ise_deployment_3_2/m_ise_deployment_models.html
Endpoint Profiling and Context Sharing:https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/admin_guide/b_ise_admin_3_2/m_prof_admin.html#Cisco_Concept.dita_2b5c5a6b-e183-4e7f-a796-a64569944409
Cisco ISE System Operations:https://www.cisco.com/c/en/us/td/docs/security/ise/3-2/admin_guide/b_ise_admin_3_2/m_sysops.html#Cisco_Concept.dita_84a88944-b54a-4a6c-9a81-f7e8058e5f25

Question: 62

An administrator is configuring a switch port for use with 802.1X.
What must be done so that the port will allow voice and multiple data endpoints?

A.Connect a hub to the switch port to allow multiple devices access after authentication.
B.Configure the port with the authentication host-mode multi-auth command.
C.Connect the data devices to the port, then attach the phone behind them.
D.Use the command authentication host-mode multi-domain on the port.

Answer: B

Explanation:

The correct answer is **B. Configure the port with the authentication host-mode multi-auth command.**

Here's why:

802.1X port authentication, by default, typically only allows a single authorized endpoint per port. However, real-world scenarios often require supporting multiple devices, particularly when combining voice (IP phones) and data endpoints on the same physical connection. Cisco switches offer different host modes to address this. The authentication host-mode multi-auth command specifically enables multi-authentication on a switch port. This allows multiple endpoints to authenticate independently using 802.1X, each subject to its own authorization policies, after the first device successfully authenticates.

Option A is incorrect because using a hub introduces a shared medium which bypasses the per-device authentication goals of 802.1X, thereby severely reducing security. Option C is incorrect because while some deployments may place phones behind data devices, this is not the solution for supporting multiple devices via 802.1X and requires specific configurations depending on the phone and device capabilities. Option D, authentication host-mode multi-domain, is used for voice and data separation where a single endpoint authenticates on each domain (voice or data) but does not enable multiple devices per domain. Therefore, it's not suitable for multiple data and voice endpoint on the same port.

The multi-auth mode provides the necessary granularity for an environment like this where different devices may have different access policies. It enables the switch to handle authentication requests from multiple devices plugged into the same port without compromising security.

Authoritative Links:

Cisco Documentation on 802.1X Port-Based Authentication:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3850/software/release/16-9/configuration/security/b_169_sec_3850_cg/b_169_sec_3850_cg_chapter_010100.html (Focus on the sections regarding host modes)
Cisco Configuration Guide - Port-Based Authentication:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/17-9/configuration_guide/sec/b_179_sec_9300_cg/port_based_authentication.html (Look for details on multi-auth mode)

Question: 63

An administrator is troubleshooting an endpoint that is supposed to bypass 802.1X and use MAB. The endpoint is bypassing 802.1X and successfully getting network access using MAB, however the endpoint cannot communicate because it cannot obtain an IP address.
What is the problem?

A.The endpoint is using the wrong protocol to authenticate with Cisco ISE.
B.The 802.1X timeout period is too long.

C.The DHCP probe for Cisco ISE is not working as expected.

D.An ACL on the port is blocking HTTP traffic.

Answer: B

Explanation:

The correct answer is **B. The 802.1X timeout period is too long.** Here's why:

When an endpoint attempts network access, it first tries 802.1X authentication. If this fails or isn't configured, the switch should then fall back to MAB (MAC Authentication Bypass). A long 802.1X timeout means the switch waits for a prolonged period for an 802.1X response that will never come (because the endpoint is configured to use MAB). During this waiting period, the switch does not initiate the MAB process. Only after the 802.1X timer expires does the switch proceed to the MAB authentication. Since the endpoint receives access but no IP address via MAB, the issue is not with MAB itself, but the delay caused by waiting on 802.1x.

This delay means the endpoint isn't granted access in time to perform DHCP, since it's waiting for authentication and won't get an IP until authentication is successful. Therefore, the device won't obtain an IP via DHCP until after the timeout, but by that point DHCP may have expired or the device may not request a new address. Reducing the 802.1X timeout allows the switch to quickly bypass 802.1X and trigger MAB, enabling the device to promptly initiate the DHCP process and get an IP address. This isn't an issue with the wrong authentication protocol (A), the DHCP Probe (C), or HTTP ACLs (D) since the endpoint bypasses 802.1x using MAB, which has successfully granted network access, but after the time the device can successfully request an IP address via DHCP. The timing is the critical factor causing the failure, not the access or configuration themselves.

Further Research:

Cisco ISE Configuration: Refer to Cisco's official documentation on configuring 802.1X and MAB for a comprehensive understanding. Specifically, search for timeout settings and their impact on authentication behavior.
https://www.cisco.com/c/en/us/support/security/identity-services-engine/tsd-products-support-series-home.html
802.1X and MAB Concepts: Review Cisco's documentation on the interaction of 802.1X and MAB, including how the switch decides which authentication method to use. Search specifically for the 802.1x timeout functionality and how this may impact the MAB process
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9300/software/release/16-12/configuration/security/b_1612_security_9300_cg/b_1612_security_9300_cg_chapter_01001.html

Question: 64

A Cisco ISE administrator must restrict specific endpoints from accessing the network while in closed mode. The requirement is to have Cisco ISE centrally store the endpoints to restrict access from.
What must be done to accomplish this task?

A.Create a profiling policy for each endpoint with the cdpCacheDeviceId attribute.

B.Create a logical profile for each device's profile policy and block that via authorization policies.

C.Add each MAC address manually to a blocklist identity group and create a policy denying access.

D.Add each IP address to a policy denying access.

Answer: C

Explanation:

The correct answer is **C. Add each MAC address manually to a blocklist identity group and create a policy denying access.** This approach leverages Cisco ISE's Identity Management capabilities for effective endpoint

restriction.

Here's why:

Identity Groups: Cisco ISE utilizes Identity Groups to categorize network devices and users. Creating a specific "blocklist" group is a logical way to organize endpoints that require restricted access.

MAC Address as Identity: MAC addresses are unique hardware identifiers associated with network interfaces, making them reliable identifiers for endpoints. Cisco ISE can use these to control access. **Authorization Policies:** Cisco ISE uses authorization policies to determine network access based on matching criteria. Once the devices are grouped into the blocklist, an authorization policy can be configured to deny access specifically for this group.

Centralized Management: Storing the MAC addresses within a blocklist identity group in Cisco ISE provides a centralized and efficient way to manage endpoint access. Instead of creating individual policies for each device (as proposed in options A and B), this method provides scalable management for a growing list of blocked endpoints.

Why other options are not optimal:

Option A (Profiling Policy based on CDP): While profiling helps identify devices, creating a profile based on cdpCacheDeviceId is not the correct approach for access restriction. The CDP is a discovery protocol that can help identify devices but is not suitable for managing an allow/deny list.

Option B (Logical Profile): Logical profiles do not handle specific endpoint identifiers like MAC addresses.

Using logical profiles to achieve this would make managing a list of blacklisted devices unnecessarily complex.

Option D (IP Address): Relying on IP addresses is unreliable because they can change. It's best to use MAC addresses as a unique identifier for network endpoints.

By using a blocklist identity group and associated authorization policies, the Cisco ISE administrator can centrally manage and restrict access for the specified endpoints.

Relevant Links for Further Research:

Cisco Identity Services Engine (ISE) Documentation: https://www.cisco.com/c/en/us/support/security/identity-services-engine/tsd-products-support-series-home.html

Cisco ISE Identity Management: Search for sections related to identity management, identity groups, and authorization policies within the Cisco ISE documentation for detailed information.

MAC Address Fundamentals:https://www.geeksforgeeks.org/what-is-a-mac-address/

Question: 65

An engineer is using profiling to determine what access an endpoint must receive. After configuring both Cisco ISE and the network devices for 802.1X and profiling, the endpoints do not profile prior to authentication.
What are two reasons this is happening? (Choose two.)

A.Closed mode is restricting the collection of the attributes prior to authentication.

B.The HTTP probe is malfunctioning due to closed mode being enabled.

C.The SNMP probe is not enabled.

D.NetFlow is not enable on the switch, so the attributes will not be collected.

E.The switch is collecting the attributes via RADIUS but the probes are not sending them.

Answer: AE

Explanation:

Okay, let's break down why options A and E are the correct reasons for profiling failures before authentication in Cisco ISE, focusing on the 802.1X context.

Option A: "Closed mode is restricting the collection of the attributes prior to authentication." This is correct. Cisco ISE's profiling feature relies on collecting network attributes from endpoints to accurately categorize them. When ISE is configured in "closed mode" for profiling, it **only** allows attribute collection after a successful authentication. This means if the endpoint hasn't authenticated, ISE won't actively probe or gather information about it to profile it. This hinders pre-authentication profiling.

Option E: "The switch is collecting the attributes via RADIUS but the probes are not sending them." This is also correct. While switches involved in 802.1X may send specific attributes (like MAC address) via RADIUS during authentication attempts, ISE's profiling usually relies on more detailed information. These probes include protocols like HTTP, DHCP, and SNMP. If these probes are not configured to send the detailed data, ISE won't receive the necessary information to profile correctly. Specifically, while the switch relays the RADIUS authentication request to ISE, the probes from the ISE are not sending back the profiling data based on observed attributes because the switch has not yet allowed traffic flow.

Why other options are incorrect:

Option B: While closed mode affects profiling, it doesn't specifically cause HTTP probes to malfunction. It prevents their activation entirely until authentication occurs. So the probe would not "malfunction", it is simply never initiated.

Option C: While SNMP is a useful profiling probe, its absence isn't a primary reason for the immediate failure of pre-authentication profiling. Other probes should still work if correctly configured outside of "closed mode".

Option D: NetFlow is a data source for network traffic analysis and is useful for post-authentication analysis but it is not needed for pre-authentication profiling. Profiling mainly depends on attribute collection from active probes.

In summary: The core issue here is that the system is likely in a closed-mode profiling configuration coupled with a misconfiguration or lack of activated attribute probes. These combined conditions prevent ISE from getting pre-authentication data needed for the endpoint to be accurately profiled.

Authoritative Links for Further Research:

Cisco Identity Services Engine (ISE) Administrator Guide: This is a primary resource for in-depth information on ISE configuration, including profiling options: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html
Cisco Validated Design for ISE: Provides recommended architectures and best practices, which often include detailed profiling guidance: https://www.cisco.com/c/en/us/solutions/enterprise/design-zone/index.html (Search for "ISE" within the design zone)
Cisco ISE Profiling: Cisco Documentation about Profilinghttps://www.cisco.com/c/en/us/td/docs/security/ise/3-3/admin_guide/b_ise_admin_3_3/m_ise_profiling_basics.html

Question: 66

Which two external identity stores support EAP-TLS and PEAP-TLS? (Choose two.)

A. RSA SecurID
B. RADIUS Token
C. Active Directory
D. Internal Database
E. LDAP

Answer: CE

Explanation:

The question asks which external identity stores support EAP-TLS and PEAP-TLS, authentication protocols that rely on digital certificates for secure communication. Active Directory (C) is a directory service widely used for managing users, computers, and other resources within a Windows domain environment. It inherently supports the use of user and computer certificates for authentication via EAP-TLS and PEAP-TLS. These protocols leverage the certificate infrastructure built into Active Directory to provide mutual authentication and secure session establishment. LDAP (E), while a directory access protocol, is often used to connect to underlying directory systems that may hold the necessary user and certificate information. Although LDAP itself does not contain user certificate information, it can access directory servers that store certificates, which means it can be indirectly used to authenticate using EAP-TLS and PEAP-TLS once a system (such as Active Directory) has the certificates. RSA SecurID (A) uses one-time passwords generated by hardware or software tokens, not certificates. RADIUS Token (B) similarly uses shared secrets or one-time passwords for authentication. The Internal Database (D) within Cisco ISE is not an external identity store. Thus, Active Directory (C) and LDAP (E) are the correct answers.

For further research, consider exploring these resources:

Microsoft's documentation on Active Directory and certificate services: https://learn.microsoft.com/en-us/windows-server/identity/
Cisco's documentation on Identity Services Engine (ISE) and external identity stores: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-technical-reference-list.html
RFC standards related to EAP and TLS protocols: https://datatracker.ietf.org/doc/html/rfc3748, https://datatracker.ietf.org/doc/html/rfc5280
Information on the PEAP protocol: https://www.cisco.com/c/en/us/support/docs/security-vpn/protected-extensible-authentication-protocol-peap/116158-technote-peap-00.html

Question: 67

An engineer deploys Cisco ISE and must configure Active Directory to then use information from Active Directory in an authorization policy.
Which two components must be configured, in addition to Active Directory groups, to achieve this goal? (Choose two.)

A.Identity Source Sequences
B.LDAP External Identity Sources
C.Active Directory External Identity Sources
D.Library Condition for Identity Group: User Identity Group
E.Library Condition for External Identity: External Groups

Answer: CE

Explanation:

The correct answer is **C. Active Directory External Identity Sources** and **E. Library Condition for External Identity: External Groups**.

Here's why:

To utilize Active Directory (AD) information within Cisco ISE authorization policies, you first need to integrate ISE with the AD domain. This is achieved by adding AD as an **Active Directory External Identity Source (C)**. This configuration allows ISE to query AD for user and group information. Without this, ISE cannot access or

validate users against AD.

Next, to use AD groups in authorization policies, you don't directly use the AD groups in their raw form. Instead, you use a **Library Condition for External Identity: External Groups (E)**. This condition acts as a filter and allows you to create policy rules based on group memberships retrieved from the configured external identity source (which is AD). This allows for dynamic authorization based on user group membership within Active Directory.

Identity Source Sequences (A) are used to order the sources that ISE queries for identity information, but is not essential to achieving authorization based on AD group membership. LDAP External Identity Sources (B) are for LDAP sources that are not Active Directory. Finally, Library Condition for Identity Group: User Identity Group (D) is used for internal ISE identity groups, not external groups from Active Directory.

Therefore, the combination of configuring an Active Directory External Identity Source and using the Library Condition for External Groups allows for the desired AD-based authorization policy creation.

Supporting Links:

Cisco ISE Admin Guide: https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/admin_guide/b_ise_admin_2_7/b_ise_admin_27_chapter_01001.html (Refer to the section on configuring external identity sources)
Cisco ISE Configuration Guides: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-configuration-examples-list.html (Search for AD integration and policy examples)

Question: 68

Which deployment mode allows for one or more policy service nodes to be used for session failover?

A. centralized
B. secondary
C. standalone
D. distributed

Answer: D

Explanation:

The correct answer is **D. distributed**. In a Cisco Identity Services Engine (ISE) distributed deployment, multiple Policy Service Nodes (PSNs) work together. This architecture enables load balancing and high availability through redundancy. If one PSN fails, another can seamlessly take over session processing, ensuring minimal disruption to network access and security. Centralized and standalone deployments, conversely, do not inherently offer such failover capabilities. A centralized deployment typically has a single Policy
Administration Node (PAN) and PSNs, but redundancy for policy enforcement relies on the distribution of PSNs. Standalone mode, while simplest to deploy, only has one node handling all functions, making it a single point of failure. A secondary node implies a specific role within a distributed setup, not a deployment mode itself. Therefore, the distributed model is the only one explicitly designed with the session failover
characteristic in mind, making it the ideal choice for resilient network access control. Each PSN replicates configuration data from the PAN, allowing them to continue operations even if the central node is temporarily unavailable.

Further reading:

Cisco ISE Deployment Models: https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/install_guide/b_ise_InstallationGuide27/b_ise_InstallationGuide27_chapter_0100.html

Cisco ISE High Availability: https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/admin_guide/b_ise_admin_guide_27/b_ise_admin_guide_27_chapter_01111.html

Question: 69

A network administrator has just added a front desk receptionist account to the Cisco ISE Guest Service sponsor group. Using the Cisco ISE Guest Sponsor Portal, which guest services can the receptionist provide?

A.Keep track of guest user activities.

B.Create and manage guest user accounts.

C.Configure authorization settings for guest users.

D.Authenticate guest users to Cisco ISE.

Answer: B

Explanation:

The correct answer is **B. Create and manage guest user accounts.**

Here's why: Cisco ISE's Guest Sponsor Portal is designed with specific roles and permissions to delegate guest management tasks. A user added to the "Guest Service sponsor group" is granted the ability to manage guest accounts, which primarily includes creating new guest user credentials (usernames and passwords), setting their duration, and potentially revoking access.

Options A, C, and D are incorrect.

A. Keep track of guest user activities: While ISE can track guest activities, this function is typically accessed through reporting and monitoring tools within the ISE admin interface, not the Guest Sponsor Portal. Monitoring often requires higher administrative privileges than a front desk receptionist would typically have.

C. Configure authorization settings for guest users: Authorization policies are configured by network administrators within the ISE policy sets and are not within the scope of a sponsor's duties. A front desk receptionist typically would not be configuring network-level access permissions.

D. Authenticate guest users to Cisco ISE: Authentication is the process by which guests verify their identity using the credentials provided. While the receptionist creates those credentials, they do not directly authenticate users, but rather enable the users to do so when logging into the network.

In summary, the primary responsibility delegated to a "Guest Service sponsor" is the creation and management of guest accounts, not advanced configuration or monitoring. They are essentially the gatekeepers for guest network access, issuing the credentials required to connect.

Authoritative Links:

Cisco ISE Guest Access:https://www.cisco.com/c/en/us/td/docs/security/ise/3-1/guest_access/b_ISE_Guest_Access_3_1_Admin_Guide/m_guest_portal.html (Refer to the "Guest Sponsor Portal" section).

Cisco ISE Admin Guides:https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-installation-and-configuration-guides-list.html (This provides access to documentation for various ISE versions.)

These resources from Cisco provide detailed information on the functionalities and permissions associated with different user roles within the Cisco ISE system.

Question: 70

What is needed to configure wireless guest access on the network?

A. endpoint already profiled in ISE
B. WEBAUTH ACL for redirection
C. Captive Portal Bypass turned on
D. valid user account in Active Directory

Answer: B

Explanation:

The correct answer is **B. WEBAUTH ACL for redirection**. To facilitate wireless guest access, a crucial component is the redirection of unauthenticated users to a captive portal. This redirection is achieved using a WEBAUTH Access Control List (ACL). The WEBAUTH ACL, applied at the network device (e.g., wireless controller), intercepts initial web requests from guest users. Instead of allowing access to the requested website, the ACL redirects the traffic to the Cisco ISE's (Identity Services Engine) captive portal. This portal presents the guest user with a login page or terms of service agreement. Once the user interacts with the captive portal (e.g., authenticates or accepts terms), ISE authorizes their access and the redirected traffic is replaced with regular network traffic. Option A is incorrect because endpoint profiling, while useful for other ISE functions, isn't a prerequisite for the basic functionality of guest access. Option C, captive portal bypass, is the opposite of what is needed for guest access as it would allow traffic through without authentication.

Option D, while crucial for domain users, isn't mandatory for generic guest access; guest users often don't need Active Directory accounts. The WEBAUTH ACL is therefore the most fundamental requirement to redirect unauthenticated traffic to a captive portal for guest wireless access.

For further reading on Cisco ISE and WebAuth, consider these authoritative links:

Cisco Identity Services Engine Configuration Guides: https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-configuration-guides-list.html
Cisco ISE Guest Access Configuration:https://www.cisco.com/c/en/us/td/docs/security/ise/2-6/ise_guest_access/b_ISE_Guest_Access/b_ISE_Guest_Access_chapter_010.html
Web Authentication on Cisco Devices:https://www.cisco.com/c/en/us/support/docs/security/identity-services-engine/117611-technote-ise-00.html

Question: 71

Which two methods should a sponsor select to create bulk guest accounts from the sponsor portal? (Choose two.)

A.Known
B.Monthly
C.Daily
D.Imported
E.Random

Answer: DE

Explanation:

The correct answer is **D. Imported** and **E. Random**. Cisco ISE's sponsor portal facilitates guest account creation for network access. To create multiple accounts efficiently, sponsors can use the 'Imported' method, where a CSV file containing guest details is uploaded to the system. This allows for pre-defining user

attributes like usernames, passwords, and expiration dates, saving time and effort. The 'Random' option, on the other hand, automatically generates usernames and passwords based on configured criteria. This is suitable when specific user details aren't necessary. Options 'Known', 'Monthly', and 'Daily' aren't methods for creating bulk accounts within the ISE sponsor portal. 'Known' might refer to specific individuals known by the sponsor, but it doesn't indicate a bulk creation method. 'Monthly' and 'Daily' are related to scheduling, not account generation itself. Bulk account creation is crucial for managing large guest populations, providing scalable and efficient network access provisioning. ISE simplifies this through import functionalities and automated username and password generation.

Relevant Cloud Computing Concepts:

Scalability: Bulk account creation addresses the need to efficiently provision access for a large number of users.
Automation: The random username/password generator automates the account creation process.
Centralized Management: ISE acts as a central point for managing all guest accounts.

Authoritative Links for Further Research:

1. **Cisco Identity Services Engine (ISE) Guest Management:** https://www.cisco.com/c/en/us/td/docs/security/ise/3-1/guest_admin/b_ise_guest_admin_3_1/b_ise_guest_admin_3_1_chapter_01010.html (This Cisco documentation provides information about guest management, specifically details on how to utilize the sponsor portal for various guest access features).
2. **Cisco ISE Configuration Guides**https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-configuration-guides-list.html (Browse the various configuration guides for ISE and find relevant documentation on sponsor portals and guest access).

Question: 72

What is a valid guest portal type?

A.Sponsor
B.Sponsored-Guest
C.Captive-Guest
D.My Devices

Answer: B

Explanation:

The correct answer is **B. Sponsored-Guest**. Cisco ISE offers various portal types for guest access, and the "Sponsored-Guest" portal is a standard offering. This portal allows an internal user (sponsor) to create guest accounts. The guest then authenticates using these credentials. Option A, "Sponsor," refers to the user who creates the guest accounts, not the portal type itself. Option C, "Captive-Guest," is not a recognized term within Cisco ISE's guest portal definitions; typically, a captive portal is used in conjunction with various authentication methods, including sponsor-based guest access. Option D, "My Devices," relates to user self-registration of personal devices for network access, not guest access. A sponsored-guest portal is designed to provide controlled access to visitors or temporary users, enhancing network security and control. The sponsor is responsible for approving and managing guest accounts, limiting potential abuse. This mechanism is common in corporate environments and educational institutions where a layer of oversight for guest network access is required. In summary, the Sponsored-Guest portal type is specifically designed to leverage internal users to grant guest access via accounts they create.

Further research on this topic can be done on the official Cisco ISE documentation:

Cisco Identity Services Engine Guest Access:https://www.cisco.com/c/en/us/td/docs/security/ise/2-4/guest_access/b_ise_guest_access_24/b_ise_guest_access_24_chapter_010.html
Cisco ISE Guest Portals Overview:https://www.cisco.com/c/en/us/td/docs/security/ise/2-2/guest_access/b_guest_access_22/b_ise_guest_access_22_chapter_010.html (This documentation is for ISE 2.2, but the concepts are broadly similar across versions.)

Question: 73

What is the purpose of the ip http server command on a switch?

A.It enables the https server for users for web authentication.
B.It enables dot1x authentication on the switch.
C.It enables MAB authentication on the switch.
D.It enables the switch to redirect users for web authentication.

Answer: D

Explanation:

The ip http server command on a Cisco switch is crucial for enabling HTTP services, specifically facilitating web-based authentication redirection. It does not directly enable HTTPS or authentication methods like Dot1x or MAB. Instead, it sets up the switch to act as an HTTP server, listening for incoming web requests. When a client attempts to access the network without proper authentication, the switch uses this server to redirect the client's browser to a specific web authentication portal. This portal, often hosted on an external server like Cisco ISE, prompts the user for credentials. Once authenticated, the user is permitted network access. The ip http server command lays the foundation for this redirection process, ensuring users are guided towards the authentication mechanism. While HTTPS provides encrypted communication, the ip http server command itself only handles basic HTTP, often used for the initial redirection step. This command is fundamental for implementing web authentication solutions where the switch plays a critical role in directing users to the authentication point. Without this command, redirection to a web-based authentication portal would not be possible.

Authoritative Links:

Cisco Command Reference: (Search for "ip http server" within Cisco IOS Command References) https://www.cisco.com/c/en/us/support/ios-nx-os-software/index.html
Cisco Identity Services Engine (ISE) Documentation: (Specifically, documentation related to web authentication configuration) https://www.cisco.com/c/en/us/support/security/identity-services-engine/tsd-products-support-series-home.html

Question: 74

Which advanced option within a WLAN must be enabled to trigger Central Web Authentication for Wireless users on AireOS controller?

A.DHCP server
B.override Interface ACL
C.static IP tunneling
D.AAA override

Answer: D

Explanation:

The correct answer is **D. AAA override**.

Here's a detailed justification:

Central Web Authentication (CWA) requires the wireless controller to redirect users to a web page for authentication before granting network access. This redirection mechanism is controlled through the "AAA override" option within a WLAN configuration on Cisco AireOS controllers. When "AAA override" is enabled, the controller will consult an external AAA (Authentication, Authorization, and Accounting) server, like Cisco ISE, to determine how to handle client authentication. It instructs the controller to deviate from the default local authentication method for the WLAN. When a client associates with a WLAN that has AAA Override enabled, the controller, under normal circumstances, is configured to send an authentication request to an external RADIUS server (such as ISE). ISE would then send back the appropriate authorization attributes, such as the ACL to apply to the client. When the web auth option is set on ISE, and the client is not yet
authenticated, ISE would instead send the redirection parameters back to the controller. This redirection is what triggers the web page to appear and allows the user to authenticate using the web browser. Options like a DHCP server (A), override Interface ACL (B), and static IP tunneling (C) do not initiate the necessary client redirection to a web portal for authentication. A DHCP server allocates IP addresses, while interface ACL overrides define access control based on network interfaces. Static IP tunneling is used in different scenarios, such as remote access or mobile data optimization, not authentication redirection. Therefore, the 'AAA override' option is the explicit trigger needed to communicate with the AAA server and signal that the client needs to be redirected for web authentication.

For further research:

1. **Cisco Wireless LAN Controller Configuration Guide**: This document provides a comprehensive guide on configuring various features on Cisco WLCs, including AAA override and web authentication.https://www.cisco.com/c/en/us/support/wireless/wireless-lan-controller-software/products-installation-and-configuration-guides-list.html
2. **Cisco Identity Services Engine (ISE) Documentation**: Cisco ISE documentation provides detailed explanations of how ISE works in conjunction with wireless controllers for authentication and authorization.https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-technical-reference-list.html

Question: 75

Which configuration is required in the Cisco ISE authentication policy to allow Central Web Authentication?

A.MAB and if user not found, continue

B.MAB and if authentication failed, continue

C.Dot1x and if authentication failed, continue

D.Dot1x and if user not found, continue

Answer: A

Explanation:

The correct answer is **A. MAB and if user not found, continue**. Here's why:

Central Web Authentication (CWA) in Cisco ISE relies on a specific authentication flow. Initially, endpoints lacking 802.1X capabilities or credentials will typically perform MAC Address Bypass (MAB). MAB uses the

endpoint's MAC address to look up authorization policies. However, CWA involves redirecting the user to a web portal for login. If a MAB lookup fails to identify the user (user not found), the ISE must continue to the next step, which is the CWA redirect. Therefore, the authentication policy should be configured to use MAB initially, and if the MAC address is not found in the internal or external identity sources, it should trigger a continuation, not a failed authentication. Options B (failed authentication) and C/D (Dot1x) are not appropriate.

Dot1x is for scenarios where supplicants like laptops perform 802.1X authentication. Option C (Dot1x and if authentication failed, continue) is incorrect because CWA follows MAB failure not a Dot1x failure. Option D (Dot1x and if user not found, continue) is incorrect because CWA cannot be triggered directly after a Dot1x user not found. For a successful CWA workflow, the authentication policy should be configured to use MAB as the first step, and if user information is not found, trigger the continuation to the authorization profile that includes the redirect to the CWA portal.

Further Research:

1. **Cisco ISE Configuration Guide - Guest Services**: https://www.cisco.com/c/en/us/td/docs/security/ise/2-7/config_guide/b_ise_config_guide_27/b_ise_config_guide_27_chapter_01101.html#id_104326 (Look for the section regarding guest access and central web authentication)
2. **Cisco ISE: Central Web Authentication (CWA) - How to Configure**: https://www.youtube.com/watch?v=bHk5Q7rJ_9I (A video guide detailing CWA configuration, showcasing the required steps) 3. **Cisco Identity Services Engine (ISE) At-a-Glance**: https://www.cisco.com/c/en/us/products/collateral/security/identity-services-engine/at-a-glance-c45-743135.html (General information about Cisco ISE and its capabilities.)

Question: 76

An engineer is configuring web authentication using non-standard ports and needs the switch to redirect traffic to the correct port.
Which command should be used to accomplish this task?

A.permit tcp any any eq <port number>

B.ip http port <port number>

C.aaa group server radius

D.aaa group server radius proxy

Answer: B

Explanation:

The correct command is **B. ip http port** . This command configures the specific TCP port that the switch's embedded HTTP server listens on for web authentication requests. When using non-standard ports for web authentication, the switch needs to know which port to redirect incoming HTTP traffic to. The ip http port command achieves this by explicitly defining the listening port for the switch's HTTP server which is responsible for handling redirection to the web authentication portal. Option A (permit tcp any any eq <port number>) is an access control list (ACL) command and is used to allow or deny traffic based on source, destination and port. While it plays a crucial role in securing network access, it doesn't directly instruct the switch about which port to use for web authentication. Option C (aaa group server radius) and D (aaa group server radius proxy) are related to configuring RADIUS servers for authentication and authorization, but do not dictate the web authentication port. They deal with communication between the switch and an external authentication server, not with the local HTTP server that handles redirects. Therefore, option B is the only command that directly addresses the task of configuring the switch to use a specific port for web
authentication redirection.

Further Reading:

Cisco IOS HTTP Server Configuration: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/http/configuration/15-mt/http-15-mt-book/http-server.html
Cisco ISE Web Authentication Configuration: https://www.cisco.com/c/en/us/td/docs/security/ise/2-4/admin_guide/b_ise_admin_guide_24/b_ise_admin_guide_24_chapter_0110.html

Question: 77

An engineer is using Cisco ISE and configuring guest services to allow wireless devices to access the network. Which action accomplishes this task?

A.Create the redirect ACL on Cisco ISE and add it to the Cisco ISE Policy.
B.Create the redirect ACL on the WLC and add it to the WLC policy.
C.Create the redirect ACL on Cisco ISE and add it to the WLC policy.
D.Create the redirect ACL on the WLC and add it to the Cisco ISE policy.

Answer: D

Explanation:

Okay, here's a detailed justification for why option D is the correct answer for configuring guest wireless access using Cisco ISE, along with supporting information:

Justification:

The correct answer is **D. Create the redirect ACL on the WLC and add it to the Cisco ISE policy.** When a guest device connects to a wireless network using Cisco ISE, it usually needs to be redirected to a captive portal (hosted by ISE) for authentication and authorization. This redirection process involves a key component called a "redirect ACL" or "pre-authentication ACL."

This ACL is essential in this process because the Wireless LAN Controller (WLC), which controls wireless access points, is responsible for enforcing this access control. It is the first point of entry for any client connection, including guest devices. The WLC needs to know which traffic to allow initially (such as DHCP and DNS) and which traffic to redirect for the captive portal experience. The redirect ACL on the WLC specifies the traffic that will be allowed before authentication.

Following this logic, we must then tell the ISE Policy that we are using a WLC that needs to have the ACL rules applied to it. The ISE policy then uses this pre-authentication ACL information to determine when to intercept the initial request and subsequently redirect the unauthenticated guest to the guest portal.

Specifically, the redirect ACL typically permits DNS and DHCP traffic and blocks all other traffic. Once the user authenticates via the captive portal, the redirect ACL is removed, and the user gets full access to the network. The ISE policy is then adjusted to allow full network access, based on the authentication/authorization results.

Therefore, the redirect ACL must first exist on the WLC (where the network access is enforced), and then the ISE policy must be configured to leverage the pre-authentication acl that is stored on the WLC.

Why other options are incorrect:

Option A, B and C: These options suggest that the redirect ACL is created on Cisco ISE, which is incorrect. The redirect ACL is a network access control, which must be created on the network enforcement point. The WLC is responsible for the initial enforcement of traffic, hence the redirect ACL needs to reside on the WLC.

Authoritative Links for further research:

1. **Cisco ISE Guest Access Overview:**https://www.cisco.com/c/en/us/td/docs/security/ise/3-3/guest-access/b_ise_guest_access_3_3_guide/b_ise_guest_access_3_3_guide_chapter_010.html 2. **Cisco ISE and WLC Integration:**https://www.cisco.com/c/en/us/td/docs/security/ise/3-3/compatibility/ise_3_3_compatibility.html#id_130317
3. **Cisco WLC ACL Configuration:**https://www.cisco.com/c/en/us/td/docs/wireless/controller/8-8/config-guide/b_cg_88/access_control_lists.html

These links provide detailed information on Cisco ISE, guest access configurations, and how it integrates with wireless controllers (WLCs). By understanding the interaction between ISE and WLCs, especially regarding the traffic management, you will be able to better understand why option D is the correct answer.

Question: 78

An engineer is configuring web authentication and needs to allow specific protocols to permit DNS traffic. Which type of access list should be used for this configuration?

A. extended ACL
B. reflexive ACL
C. numbered ACL
D. standard ACL

Answer: A

Explanation:

The correct answer is **A. extended ACL**. Extended Access Control Lists (ACLs) are required for this scenario because they offer the granularity needed to specify both source and destination IP addresses, along with protocol types and port numbers. Web authentication, specifically involving DNS, necessitates this level of control. Standard ACLs only filter traffic based on source IP addresses, lacking the ability to filter by destination and protocols. Reflexive ACLs, while stateful, aren't designed for initial filtering based on protocol type. Numbered ACLs are just a naming convention and do not affect functionality; they can be either standard or extended. To allow DNS traffic (UDP or TCP port 53) during web authentication, an extended ACL would define a rule permitting traffic from the client to the DNS server and vice-versa using the appropriate protocol and port. Extended ACLs provide the flexibility needed to fine-tune which specific types of traffic are permitted before and during authentication, enabling granular security policies. This is crucial to avoid unnecessarily blocking crucial services like DNS while still enforcing access control.

For further research, refer to:

Cisco's official documentation on ACLs:https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/security/acl/configuration/15-sy/sec-acl-15-sy-book.html
Cisco's Identity Services Engine (ISE) configuration guides:
https://www.cisco.com/c/en/us/support/security/identity-services-engine/products-configuration-examples-list.html

Question: 79

An administrator is adding a switch to a network that is running Cisco ISE and is only for IP Phones. The phones do not have the ability to authenticate via
802.1X.
Which command is needed on each switch port for authentication?

A. dot1x system-auth-control

B. enable bypass-MAC

C. enable network-authentication

D. mab

Answer: D

Explanation:

The correct answer is **D. mab**. Let's break down why. The scenario describes IP phones that lack 802.1X capabilities, meaning they cannot perform traditional authentication with usernames and passwords. In such cases, MAC Authentication Bypass (MAB) becomes essential. MAB is a port-based authentication method where the switch sends the MAC address of the device as the identifier to the authentication server (Cisco ISE in this case). Cisco ISE then checks if this MAC address is authorized based on its configured policy. If the MAC address is recognized, the device is granted network access. Option A dot1x system-auth-control enables 802.1X authentication, which isn't applicable for our non-802.1X capable IP phones. Option B enable bypass-MAC does not exist as a valid command related to Cisco ISE configuration and bypass MAC mechanisms.

Option C enable network-authentication is a generic command and does not specifically enable MAB, which is required in this scenario. Therefore, MAB is the most appropriate method, and the mab command on each switch port enables it, allowing the IP phones to connect and authenticate based on their MAC addresses through Cisco ISE.

Authoritative Links for further research:

Cisco Identity Services Engine (ISE) Configuration Guide:
https://www.cisco.com/c/en/us/td/docs/security/ise/2-6/configuration_guide/b_ise_config_guide_2_6/b_ise_config_guide_26_chapter_01001.html (Search for "MAC Authentication Bypass")
Cisco Documentation on MAB:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960x/software/15_2_e/configuration/guide/b_2960x

Question: 80

A network engineer needs to ensure that the access credentials are not exposed during the 802.1X authentication among components.
Which two protocols should be configured to accomplish this task? (Choose two.)

A.PEAP

B.EAP-TLS

C.EAP-MD5

D.EAP-TTLS

E.LEAP

Answer: AD

Explanation:

The correct answer is **A. PEAP and D. EAP-TLS**. Let's break down why:

The core requirement is to protect access credentials during 802.1X authentication. This means preventing the transmission of usernames and passwords in a way that they can be easily intercepted and deciphered.

PEAP (Protected Extensible Authentication Protocol) achieves this by establishing an encrypted TLS tunnel before the actual authentication exchange. Within this secure tunnel, a weaker inner authentication method can be used (like EAP-MSCHAPv2), but the critical aspect is that the credentials are protected by TLS

encryption from the client to the authentication server (like Cisco ISE). This approach prevents eavesdropping on the authentication data in transit.

EAP-TLS (Extensible Authentication Protocol – Transport Layer Security) provides the strongest security for 802.1X. It employs mutual authentication based on certificates. The client and the authentication server exchange certificates to verify each other's identity before the session is granted. Credentials, in the traditional password sense, are not even used; instead, authentication is based on the validation of certificates, making it highly resistant to password interception attacks.

Both PEAP and EAP-TLS secure the transmission of authentication data by using encryption or certificate validation.

EAP-MD5 (Extensible Authentication Protocol – Message Digest 5) is an older and less secure method that is not recommended for credential protection. It simply hashes the password, which can be cracked using readily available tools. EAP-TTLS (Extensible Authentication Protocol Tunneled TLS) encrypts the communication with TLS after the initial handshake. However, like EAP-MD5, a username and password are sent in the tunnel, making it a lesser option than EAP-TLS. LEAP (Lightweight Extensible Authentication Protocol) is a Cisco proprietary protocol that is considered insecure and vulnerable to attacks. Therefore, it should not be considered for use.

Therefore, PEAP and EAP-TLS are the ideal choices to fulfill the security requirements in the given scenario.

Authoritative Links:

Cisco - Understanding EAP Methods for 802.1X Authentication: https://www.cisco.com/c/en/us/support/docs/security-vpn/extensible-authentication-protocol-eap/14098-eap-types.html
National Institute of Standards and Technology (NIST) - Guidelines for Securing Wireless Networks Using 802.11 Authentication and Encryption: https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-48r1.pdf (While not Cisco-specific, it discusses EAP methods generally)