



Welcome to Cert007 - Your Ultimate IT Certification Partner



➤ Real Exam Questions

➤ Instant Access

➤ Free Updates

➤ Money-Back Guarantee

➤ Expert Support



Visit us at <https://www.cert007.com/> for more information

Exam : **NSE6_FNC_AD-7.6**

Title : Fortinet NSE 6 - FortiNAC-F
7.6 Administrator

Version : DEMO

1.Scenario: A large university is deploying FortiNAC-F to enforce a captive portal for unregistered student devices across its sprawling campus. The campus network architecture consists of over 200 distinct building subnets. The network engineering team has a strict administrative constraint: they absolutely cannot manage 200 unique DHCP scopes on the FortiNAC appliance itself, nor can they manually map 200 individual logical networks within the FortiNAC GUI due to operational overhead. The enforcement mechanism must be highly scalable.

A. Manually create 200 unique Logical Networks within the FortiNAC GUI, each bound to a specific building's VLAN ID, and configure FortiNAC's internal DHCP server with discrete scopes for every subnet to ensure accurate L3 routing to the portal.

B. Enable the "Layer 3 HA Distributed Portal" feature on the FortiNAC-F Manager, which dynamically provisions decentralized captive portals directly onto the edge access switches, entirely bypassing the need for centralized DHCP.

C. Configure the edge switches to use DHCP Relay for the Registration VLAN pointing to the FortiNAC appliance, and map a single global Isolation Network in FortiNAC to handle all captive portal redirections without needing per-subnet DHCP scopes.

D. Use a Python script via Security Automation to update the FortiGate firewall tags, forcing the upstream FortiGate to inject a captive portal page into the user's browser session regardless of their local L2 subnet.

Answer: C

Explanation:

Option C is the correct and most scalable architectural approach for Isolation Networks. FortiNAC can utilize a single global Isolation Network (e.g., a single Registration VLAN ID used consistently across all edge switches). By configuring DHCP Relay on the edge switches to point to FortiNAC, the appliance can manage the IP assignment and DNS redirection for the captive portal centrally without requiring the administrator to build and maintain 200 discrete subnets and DHCP scopes within the FortiNAC GUI. Option A explicitly violates the administrative constraint by requiring the manual creation of 200 scopes and networks.

Option B is incorrect because no such "Layer 3 HA Distributed Portal" feature exists on the Manager to push portals to edge switches.

Option D is incorrect because injecting portal pages via FortiGate firewall tags is not a valid or supported method for FortiNAC native captive portal enforcement, and it bypasses FortiNAC's isolation network capabilities entirely.

2.Scenario: A financial institution is designing a resilient FortiNAC architecture to protect its critical infrastructure. They currently operate three primary active data centers, each with its own standalone FortiNAC-F appliance managing local Layer 2 MAC Authentication Bypass (MAB) enforcement. Compliance regulations dictate a maximum 5-minute Recovery Time Objective (RTO) if any single appliance suffers a catastrophic hardware failure. The IT budget strictly limits the purchase to exactly one additional FortiNAC-F appliance for redundancy purposes.

A. Deploy the single newly purchased appliance in an N+1 High Availability architecture, configuring it as the secondary server to actively back up all three primary appliances simultaneously.

B. Configure Hot Standby mode by connecting the new appliance to the main data center and using FortiNAC-F Manager to stretch the Layer 2 heartbeat across the WAN to the other two remote sites.

C. Install the single appliance in a centralized cloud VPC and configure it as a standard Fabric

Connector. In the event of a failure, manually update the RADIUS server IP addresses on all edge switches via CLI to point to this connector.

D. Utilize the built-in Database Clustering feature to synchronize the MariaDB instances across the three existing primary nodes in an active-active-active mesh, eliminating the need to deploy the newly purchased hardware.

Answer: A

Explanation:

Option A is the correct design. FortiNAC's N+1 High Availability architecture is specifically designed for this scenario. It allows a single secondary appliance to act as a backup for multiple (up to 3 or more, depending on sizing) primary appliances. If any primary fails, the secondary assumes its identity and IP, meeting the strict 5-minute RTO and the budget constraint.

Option B is incorrect because Hot Standby HA is strictly a 1:1 redundancy model; one secondary cannot serve as a Hot Standby for three separate primary active appliances.

Option C violates the 5-minute RTO and the intent of automated HA, as manual CLI updates on edge switches would cause significant downtime.

Option D is incorrect because FortiNAC does not support an active-active-active mesh database clustering architecture for L2 enforcement redundancy.

3.Scenario: A manufacturing plant heavily relies on FortiNAC for endpoint compliance and threat containment. The CISO demands that if an engineering workstation initiates outbound SSH traffic to a known malicious IP (detected by a third-party perimeter Intrusion Detection System via syslog), the workstation must immediately have its access revoked at the switch port level. The solution must not require manual administrator intervention and must execute within 15 seconds.

A. Configure a FortiOS Automation Stitch on the third-party firewall to send a REST API call directly to the edge access switch to shut down the port, bypassing the FortiNAC database entirely.

B. Map the third-party firewall as a generic SNMP device in the FortiNAC topology tree, configure FortiNAC to poll the firewall's ARP table every 10 seconds, and use a device profiling rule to strip the workstation of its active network access policy.

C. Assign the workstation to a strict Host Profile that forces a mandatory FortiClient posture check every 5 seconds to continuously verify that no malicious outbound SSH sessions are actively established.

D. Create a custom Security Event Parser in FortiNAC to extract the compromised host's IP address from the third-party syslog payload, then build a Security Rule that matches this event to automatically change the host state to "At Risk" and trigger an isolation VLAN reassignment.

Answer: D

Explanation:

Option D correctly addresses the requirements using native Security Automation. FortiNAC can ingest syslogs from third-party devices. Because syslog formats vary, a custom Security Event Parser (often using Regex) is required to extract the offending endpoint's IP or MAC address. Once parsed, a Security Rule triggers, changing the host state to "At Risk" and instructing the switch to change the VLAN via RADIUS CoA or SNMP, easily meeting the 15-second SLA.

Option A violates the scenario constraints because the perimeter device is a third-party IDS, not a FortiGate (cannot use FortiOS stitches), and bypassing FortiNAC breaks the central visibility mandate.

Option B is incorrect because polling an ARP table does not provide intrusion detection data or identify malicious SSH traffic; Profiling Rules classify devices, they do not respond to real-time security events.

Option C is incorrect because running posture checks every 5 seconds creates crippling overhead, and posture checks evaluate endpoint configuration (like AV status or registry keys), not real-time outbound network flow analysis.

4.Scenario: A regional hospital utilizes FortiGate for remote contractor VPN access and FortiNAC for internal access control. The IT department wants FortiNAC to dynamically apply access restrictions for remote VPN users. When a contractor logs in via FortiClient, FortiNAC must instruct the FortiGate to restrict the user to the "Medical_Records" subnet if the endpoint passes an antivirus posture check, or restrict them to a highly limited "Remediation" subnet if the check fails. The network architecture board strictly prohibits the use of separate, dedicated IPsec VPN tunnels or static IP pools to manage different compliance states.

A. Configure FortiNAC to terminate the VPN sessions directly using its internal RADIUS server, applying different L2 VLAN tags based on the posture check results before routing the traffic upstream to the FortiGate.

B. Integrate FortiNAC with the Fortinet Security Fabric to dynamically pass Firewall Tags based on the endpoint's posture state, allowing the FortiGate to apply Dynamic Address Objects in its firewall policies to permit or deny access.

C. Use a Security Automation rule in FortiNAC to execute an SSH script against the FortiGate, dynamically modifying the static routing table to blackhole the contractor's IP address if the posture check fails.

D. Establish a Syslog trap listener on the FortiGate that parses FortiNAC log events. When a failure is detected, the FortiGate triggers a local automation script to force a TCP RST packet to the contractor, terminating the session entirely.

Answer: B

Explanation:

Option B represents the correct and officially supported method for FortiGate VPN integration with FortiNAC. By acting as a Fabric Connector, FortiNAC evaluates the endpoint (posture check) and assigns an abstract Firewall Tag (e.g., "Compliant" or "Remediation"). This tag is shared with the FortiGate via the Security Fabric. The FortiGate uses these tags in Dynamic Address Objects within its firewall policies to enforce L3/L4 access, perfectly satisfying the requirement without needing static IP pools or multiple tunnels.

Option A is technically impossible; FortiNAC is an access control appliance, not a VPN gateway, and cannot terminate IPsec/SSL VPN tunnels.

Option C violates best practices and operational stability; relying on SSH scripts to dynamically alter core routing tables based on endpoint state is brittle and unsupported.

Option D is incorrect because terminating the session with a TCP RST does not fulfill the requirement of moving the user to a "Remediation" subnet; it merely disconnects them.

5.Scenario: A corporate environment is rapidly deploying thousands of IoT smart bulbs across its facilities. These bulbs do not support 802.1X supplicants, lack an HTTP web interface, and utilize MAC randomization (they generate a new random MAC address every time they are power-cycled). The security team requires these devices to be automatically profiled and placed into an isolated "IoT-Lighting" logical network immediately upon connection. The IT team is prohibited from purchasing or relying on external vendor OUI databases to identify the hardware.

- A. Utilize the "Manual Host Registration" feature by importing a CSV file containing the expected MAC addresses provided by the lighting vendor prior to the physical deployment.
- B. Enable automated host registration based strictly on Endpoint Fingerprints derived from HTTP User-Agent strings intercepted during the captive portal redirection process.
- C. Configure a Device Profiling Rule that utilizes DHCP fingerprinting (such as examining DHCP Option 55 or Option 60) to uniquely identify the smart bulbs regardless of their randomized MAC addresses.
- D. Deploy FortiNAC-F Manager to actively probe the smart bulbs using Nmap OS fingerprinting scripts every 60 seconds, correlating the responses with the FortiGuard threat intelligence database to assign the correct logical network.

Answer: C

Explanation:

Option C is the only viable solution. Because the MAC addresses are randomized upon reboot, traditional profiling based on the MAC OUI will fail completely. However, when the device requests an IP address, its DHCP request parameters (like Option 55 - Parameter Request List, or Option 60 - Vendor Class Identifier) remain static and specific to the device's operating system/firmware. FortiNAC can use these DHCP fingerprints in a Device Profiling Rule to accurately identify and categorize the bulbs immediately.

Option A is incorrect because MAC randomization renders the vendor's CSV list of MAC addresses useless.

Option B violates the scenario constraints because the smart bulbs lack an HTTP web interface and cannot generate HTTP User-Agent strings.

Option D is incorrect because active Nmap probing against lightweight, headless IoT devices is highly unreliable, creates massive network overhead, and fails the requirement to authorize the device "immediately upon connection."