PassQueen
a.k.a dumpsqueen.com

Check Point Certified Security Expert R81.20

Checkpoint 156-315.81

Version Demo

Total Demo Questions: 58

Total Premium Questions: 581

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Advanced System Management	149
Topic 2, Advanced Security Management	93
Topic 3, Advanced Security Gateway Management	161
Topic 4, Advanced VPN	51
Topic 5, Threat Prevention	60
Topic 6, Advanced Troubleshooting	66
Topic 7, Mix Questions	1
Total	581

QUESTION NO: 1

Which statements about HTTPS Inspection are correct? (Choose three.)

A. Outbound inspection uses certificates signed by the outbound inspection CA.

B. Clients must trust the outbound inspection CA certificate.

C. Inbound inspection requires the server certificate and private key.

D. The SIC certificate is used as the outbound inspection CA certificate.

ANSWER: A B C

Explanation:

For outbound HTTPS Inspection, the Security Gateway creates replacement certificates signed by the configured outbound inspection CA. Client devices must trust that CA certificate to avoid certificate warnings. For inbound inspection, the gateway requires access to the protected server certificate and corresponding private key so it can decrypt and inspect traffic destined for that server.

HTTPS Inspection rules can also bypass inspection for selected sites, categories, or traffic types. The gateway does not use its SIC certificate as a replacement for the outbound HTTPS Inspection CA certificate. See the R81 HTTPS Inspection Administration Guide.

QUESTION NO: 2

What is the command to see cluster status in cli expert mode?

A. fw ctl stat

B. clusterXL stat

C. clusterXL status

D. cphaprob stat

ANSWER: D

Explanation:

The command **cphaprob stat** is the standard ClusterXL command for viewing the current status of a Check Point cluster from the command line, including when working in Expert mode. It displays the local member's ClusterXL state, such as *ACTIVE*, *STANDBY*, or *READY*, and presents the status of the other cluster members as detected through the Cluster Control Protocol. This makes it the primary quick-status command for confirming which member currently handles traffic and whether the cluster has formed correctly. Administrators commonly run it during operational checks and incident troubleshooting before using more detailed ClusterXL diagnostic commands. The command is available on Security Gateway cluster members and is part of the Check Point High Availability/ClusterXL command set. Check Point's ClusterXL administration material describes monitoring the cluster and its member state through the `cphaprob` utility. See the Check Point R81 ClusterXL Monitoring documentation and the Check Point ClusterXL monitoring knowledge article for additional operational guidance.

QUESTION NO: 3

What API command below creates a new host with the name "New Host" and IP address of "192.168.0.10"?

A. new host name "New Host" ip-address "192.168.0.10"

B. set host name "New Host" ip-address "192.168.0.10"

C. create host name "New Host" ip-address "192.168.0.10"

D. add host name “New Host” ip-address “192.168.0.10”

ANSWER: D

Explanation:

The command `add host name "New Host" ip-address "192.168.0.10"` is correct because the Check Point Management API uses the `add` verb to create a new object. The object type immediately follows the verb, so `host` specifies that the new object is a Host network object. The remaining parameters define the object’s properties: `name` assigns its display name and `ip-address` assigns its IPv4 address. This command syntax is used with the Web Services API through `mgmt_cli`, or as the command value for an API request after authenticating to the Security Management Server. Names containing spaces must be quoted, as shown, so that “New Host” is treated as one value. The command creates the object in the current API session; to make the change effective, the session must subsequently be published, either with `mgmt_cli publish` or through the applicable API workflow. Check Point documents the `add host` API command and its `name` and `ip-address` parameters in the Management API reference. See the Check Point Management API Reference and the add-host command documentation.

QUESTION NO: 4

What statement best describes the Proxy ARP feature for Manual NAT in R81.10?

A. Automatic proxy ARP configuration can be enabled

B. Translate Destination on Client Side should be configured

C. fw ctl proxy should be configured

D. local.arp file must always be configured

ANSWER: A

Explanation:

“Automatic proxy ARP configuration can be enabled” is correct because R81.10 supports automatic Proxy ARP handling for applicable Manual NAT configurations. Proxy ARP lets a Security Gateway answer ARP requests on an external network for an IP address used in a NAT translation, even when that address is not assigned directly to a gateway interface. This makes the translated address reachable by neighboring devices, which then send the relevant traffic to the gateway for NAT processing.

When automatic Proxy ARP is enabled, the gateway can derive and install the required Proxy ARP behavior from the Manual NAT policy rather than requiring a separate, static ARP-entry workflow for every applicable translation address. This reduces administrative effort and helps keep ARP behavior aligned with the deployed NAT rule base, especially when published-address mappings change. The setting must still be used in a topology where the gateway is expected to answer ARP on the relevant connected network and where the translated address is appropriate for that interface.

Check Point documents NAT behavior and gateway handling in the R81.10 Security Management Administration Guide and the R81.10 Gaia Administration Guide.

QUESTION NO: 5

Which conditions can cause a connection to be handled in the Firewall Medium Path instead of the SecureXL accelerated path? (Choose two.)

A. Traffic that requires IPS inspection

B. Traffic that requires Threat Prevention inspection

C. A connection that matches an eligible Accept template

D. Traffic forwarded through an interface with link state up

E. A packet that has a valid destination route

ANSWER: A B

Explanation:

Traffic that requires IPS inspection and **traffic that requires Threat Prevention inspection** are correct. SecureXL can accelerate eligible traffic, but inspection requirements that cannot be fully handled in the accelerated path cause traffic to be forwarded to the Firewall Medium Path for additional processing. The Medium Path preserves required security inspection while avoiding some of the overhead of the full Firewall Path where applicable.

Acceleration eligibility depends on the policy, enabled blades, connection characteristics, and gateway configuration. See the R81 Traffic Acceleration documentation.

QUESTION NO: 6

Which statements about manual and automatic NAT rules are correct? (Choose three.)

A. Automatic NAT rules are generated from object NAT settings.

B. Manual NAT rules are evaluated before automatic NAT rules.

C. Automatic NAT rules are visible in the NAT Rule Base.

D. Manual NAT rules are created automatically from object NAT settings.

ANSWER: A B C

Explanation:

Automatic NAT rules are generated from the NAT settings configured on network objects. Manual NAT rules are evaluated before automatically generated NAT rules, allowing an administrator to define explicit exceptions or more specific translations. In SmartConsole, automatic NAT rules are visible in the NAT Rule Base but are derived from the relevant object configuration rather than edited directly as ordinary manual rules.

Manual NAT rules are not generated automatically from object settings; they must be created and maintained by the administrator. See the R81 Security Management Administration Guide: Network Address Translation.

QUESTION NO: 7

Which conditions are required for a Site-to-Site VPN tunnel to establish successfully between two Check Point gateways? (Choose three.)

A. Reachability between the peer external VPN addresses

B. Compatible IKE and IPsec settings

C. Correctly defined VPN domains

D. Identical internal address space behind both gateways

ANSWER: A B C

Explanation:

The gateways must be able to reach each other's external VPN addresses, and their IKE and IPsec proposals must have compatible settings. Each gateway must also have a correctly defined VPN domain that identifies the protected networks for encryption and decryption.

The internal networks behind the two gateways do not need to use the same address space. In fact, overlapping encryption domains require special design considerations and are not a normal VPN requirement. See the R81.20 Site-to-Site VPN Administration Guide.

QUESTION NO: 8

Which actions are recommended before upgrading a Security Management Server to a later R81 release? (Choose two.)

A. Create and verify a management backup

B. Validate the upgrade path and compatibility requirements

C. Delete all existing policies before the upgrade

D. Permanently disable SIC on every gateway

E. Remove all licensed Software Blades

ANSWER: A B

Explanation:

Create and verify a management backup and **validate the upgrade path and compatibility requirements** are correct. A verified backup provides a recovery option if the upgrade does not complete as expected. Administrators must also confirm that the source version, target version, hardware, operating system, Jumbo Hotfix Accumulator level, and deployed products are supported by the intended upgrade path.

Stopping SIC permanently or deleting existing policies is neither required nor recommended. Consult the R81 Installation and Upgrade Guide.

QUESTION NO: 9

Is it possible to establish a VPN before the user login to the Endpoint Client?

A. yes, you had to set neo_remember_user_password to true in the trac.defaults of the Remote Access Client or you can use the endpoint_vpn_remember_user_password
attribute in the trac_client_1 .ttm file located in the SFWDIR/conf directory on the Security Gateway

B. no, the user must login first.

C. yes. you had to set neo_always_connected to true in the trac.defaults of the Remote Access Client or you can use the endpoint_vpn_always_connected attribute in the
trac_client_1 .ttm file located in the SFWDIR/conf directory on the Security Gateway

D. yes, you had to enable Machine Authentication in the Gateway object of the Smart Console

ANSWER: D

Explanation:

Yes, enabling Machine Authentication in the Security Gateway object supports establishing a Remote Access VPN connection before a user authenticates to the Endpoint Client. Machine Authentication gives the endpoint a way to authenticate at the device level, using its machine identity or certificate, rather than requiring an interactive user authentication event as the prerequisite for the tunnel. This is particularly important when an organization needs network connectivity before user sign-in, such as to contact domain services, apply policies, retrieve updates, or support logon processes for managed Windows devices. In SmartConsole, the administrator enables and configures this capability on the applicable Security Gateway's Remote Access VPN settings, then deploys the corresponding endpoint configuration and trusted machine credentials. The gateway can therefore validate the managed computer and establish the protected connection during the pre-logon phase; user authentication can occur afterward according to the configured access-control and authentication policy. Check Point's Remote Access VPN administration documentation describes gateway-side Remote

Access VPN configuration and endpoint connectivity features at Remote Access VPN Administration Guide. Additional Check Point product documentation is available through the Check Point documentation portal.

QUESTION NO: 10

You can access the ThreatCloud Repository from:

A. R81.10 SmartConsole and Application Wiki

B. Threat Prevention and Threat Tools

C. Threat Wiki and Check Point Website

D. R81.10 SmartConsole and Threat Prevention

ANSWER: D

Explanation:

R81.10 SmartConsole and Threat Prevention is correct because the ThreatCloud Repository is integrated into the Threat Prevention workflow in SmartConsole. Administrators use the Threat Prevention area to work with protections and obtain repository-based information that supports the configuration, investigation, and operational use of Check Point threat-prevention capabilities. ThreatCloud is Check Point's cloud intelligence service, supplying continuously updated information about threats, indicators, and protections to the security-management and gateway ecosystem.

In practice, SmartConsole is the management interface through which an administrator configures Threat Prevention policy and reviews the protections relevant to the organization's security posture. The repository connection in this context allows the administrator to use Check Point's current threat intelligence while managing protections, rather than treating ThreatCloud as a separate on-premises database. This integration is a core element of the R81.10 Threat Prevention architecture and administration workflow. Check Point's Threat Prevention Administration Guide documents ThreatCloud and the management of Threat Prevention protections in SmartConsole. See the ThreatCloud Repository documentation and the Check Point ThreatCloud overview.

QUESTION NO: 11

Fill in the blank: The "fw monitor" tool can be best used to troubleshoot _____.

A. network traffic issues

ANSWER: A

Explanation:

The correct completion is **network traffic issues**. The `fw monitor` utility captures packets as they traverse the Check Point Security Gateway's inspection path in the firewall kernel. Unlike a conventional interface-level capture, it can show whether a packet reached the gateway, entered firewall inspection, passed through policy processing, and left the gateway. This visibility makes it particularly useful when investigating connectivity failures, unexpected packet drops, asymmetric flows, NAT behavior, and traffic associated with VPN processing.

Administrators can apply filters to focus the capture on relevant source and destination addresses, protocols, or ports, reducing noise while following a specific connection. By examining packets at the kernel inspection points, an administrator can establish where traffic stops appearing and narrow the problem to traffic arrival, policy/kernel processing, translation, or egress. This is why the tool is a primary diagnostic utility for path and packet-flow investigations on a Security Gateway. Check Point documents `fw monitor` as a packet-capture command that displays traffic passing through the firewall and supports filtering and inspection-point output. See the Check Point R81 CLI Reference for fw monitor and the Check Point R81 traffic-monitoring documentation.

QUESTION NO: 12

Fill in the blanks: A _____ license requires an administrator to designate a gateway for attachment whereas a _____ license is automatically attached to a Security Gateway.

A. Formal; corporate

B. Local; formal

C. Local; central

D. Central; local

ANSWER: D

Explanation:

Central; local is correct because Check Point distinguishes these license types by how they are associated with a Security Gateway. A central license is installed on, and managed by, the Security Management Server. The administrator must explicitly attach it to the intended gateway, selecting the gateway object that will consume the license. This model supports centralized license administration and makes it possible to reassign an eligible license when a gateway is replaced or its management arrangement changes.

A local license is tied directly to an individual Security Gateway. It is installed locally and is automatically associated with that gateway rather than requiring the administrator to designate a separate gateway attachment through centralized license management. The distinction is important operationally: central licensing uses an explicit attachment step, while local licensing remains bound to the gateway for which it is issued and installed.

Check Point's Security Management Administration Guide describes the management and attachment of central licenses and the use of local licenses on gateways. See the R81.10 Security Management Administration Guide: Licenses and the Check Point Licensing resources.

QUESTION NO: 13

When configuring SmartEvent Initial settings, you must specify a basic topology for SmartEvent to help it calculate traffic direction for events. What is this setting called and what are you defining?

A. Network, and defining your Class A space

B. Topology, and you are defining the Internal network

C. Internal addresses you are defining the gateways

D. Internal network(s) you are defining your networks

ANSWER: D

Explanation:

Internal network(s) you are defining your networks is correct. During SmartEvent's initial configuration, the Internal Networks setting establishes which IP networks belong to the organization's protected, internal environment. SmartEvent uses this topology context when it processes logs and correlates events, allowing it to determine the direction of a connection relative to those internal networks. For example, it can identify whether traffic is entering an internal network, leaving it, moving between internal systems, or remaining external. This directional classification is important because SmartEvent event views, severity assessment, correlation, and reporting need to distinguish activity affecting internal assets from activity that merely passes between external addresses. The setting is not a requirement to define a particular address class or to enumerate gateways; it is a definition of the internal network address ranges that represent the organization's environment. Administrators should therefore include all applicable private or routed internal subnets that SmartEvent must recognize as internal. Check Point documents SmartEvent configuration and event analysis in the R81.10 Logging and Monitoring Administration Guide and provides related product documentation through its documentation portal.

QUESTION NO: 14

Which of these statements describes the Check Point ThreatCloud?

A. Blocks or limits usage of web applications

B. Prevents or controls access to web sites based on category

C. Prevents Cloud vulnerability exploits

D. A worldwide collaborative security network

ANSWER: D

Explanation:

“A worldwide collaborative security network” correctly describes Check Point ThreatCloud. ThreatCloud is Check Point’s cloud-based threat-intelligence infrastructure. It collects and correlates threat indicators, malicious-file information, malware behavior, URLs, botnet activity, and other telemetry from a broad global sensor network, then distributes continuously updated intelligence to Check Point security products. This shared intelligence enables protections such as Anti-Bot, Anti-Virus, Threat Emulation, Threat Extraction, URL Filtering, and other Threat Prevention capabilities to recognize and prevent newly identified threats rapidly.

The key idea is collaboration at global scale: information observed by security gateways, endpoints, research systems, and Check Point’s threat-analysis services contributes to a common intelligence source. Check Point describes ThreatCloud as the brains behind its threat-prevention architecture, using big-data threat intelligence and threat discovery to identify emerging attacks and provide protections across the environment. This makes “worldwide collaborative security network” the accurate high-level characterization rather than a description of one individual security-control feature. See Check Point’s ThreatCloud overview and the ThreatCloud documentation.

QUESTION NO: 15

Which statements about ClusterXL State Synchronization are correct? (Choose three.)

A. It helps established connections survive a failover.

B. A dedicated synchronization network is recommended.

C. Synchronization interfaces must be configured in the cluster topology.

D. CCP eliminates the need for State Synchronization.

ANSWER: A B C

Explanation:

State Synchronization shares connection-state information between ClusterXL members so that established connections can continue after a failover. Check Point recommends a dedicated synchronization network to isolate synchronization traffic and avoid contention with production traffic. The interfaces used for synchronization must be defined correctly in the cluster-member topology.

CCP and State Synchronization are separate ClusterXL functions. CCP communicates member health and cluster-control information; it does not replace connection-state synchronization. See the R81.20 ClusterXL Administration Guide: State Synchronization.

QUESTION NO: 16

The CPD daemon is a Firewall Kernel Process that does NOT do which of the following?

A. Secure Internal Communication (SIC)

B. Restart Daemons if they fail

C. Transfers messages between Firewall processes

D. Pulls application monitoring status

ANSWER: D

Explanation:

Pulls application monitoring status is the correct answer. The Check Point Daemon Process (CPD) is a core gateway and management communication service. Its responsibilities center on Check Point control-plane functions, including Secure Internal Communication, inter-process communication, and supervision of Check Point services through the watchdog framework. In practical troubleshooting, CPD is therefore relevant when a gateway cannot establish or maintain trusted communication with the management server, or when essential Check Point processes require monitoring and recovery.

Application monitoring status is not a CPD responsibility. Application-level monitoring is associated with the applicable software blade, application-control functionality, and the monitoring or reporting components that collect and present application-related information. CPD provides the foundational daemon and communication services on which Check Point components rely, rather than polling application-monitoring results as a dedicated function. Check Point's process documentation describes CPD among the principal processes used for gateway and management operations, while the watchdog documentation describes process monitoring and restart behavior. See the Check Point Processes documentation and the cpwd_admin CLI reference.

QUESTION NO: 17

NO: 240

You notice that your firewall is under a DDoS attack and would like to enable the Penalty Box feature, which command you use?

A. sim erdos -e 1

B. sim erdos -m 1

C. sim erdos -v 1

D. sim erdos -x 1

ANSWER: A

Explanation:

`sim erdos -e 1` is the correct command because it enables the SecureXL ERDOS (Early Drop / denial-of-service protection) mechanism, which includes use of the Penalty Box. ERDOS is designed to protect a Security Gateway when it is receiving excessive traffic, such as during a flood or distributed denial-of-service event. After it is enabled, SecureXL can identify traffic sources that exceed relevant thresholds and place them in the Penalty Box, causing their packets to be dropped for a defined period. This protects gateway resources by reducing the work required for traffic that is behaving like a flood source and allows legitimate traffic a better opportunity to be processed. The `-e` parameter controls whether ERDOS is enabled, and the value `1` turns it on. This is a runtime SecureXL command and should be applied with an understanding of the gateway's traffic patterns and operational requirements. Check Point describes SecureXL acceleration and performance-protection configuration in its R81.10 Performance Tuning Administration Guide; related command syntax is documented in the R81.10 CLI Reference Guide.

QUESTION NO: 18

Which one is not a valid Package Option In the Web GUI for CPUSE?

A. Clean Install

B. Export Package

C. Upgrade

D. Database Conversion to R81.10 only

ANSWER: B

Explanation:

Export Package is not a CPUSE package deployment option. In the Gaia Portal CPUSE workflow, package options describe the operation that the appliance performs with a selected CPUSE upgrade package. These operations determine the installation or migration path, such as applying the package as an upgrade, performing a clean installation, or—in the applicable R81.10 management-server scenario—performing database conversion only. They are choices made as part of executing the package on the local Check Point system.

Exporting, by contrast, concerns making a package available for transfer or reuse and does not define how CPUSE installs, upgrades, or converts the local system. It is therefore not part of the Package Option selection used to choose the package deployment behavior in the CPUSE Web GUI. Check Point's installation and upgrade guidance distinguishes CPUSE package-management tasks from the upgrade and clean-install procedures executed by CPUSE. See the Check Point R81.10 CPUSE documentation and the Upgrade with CPUSE procedure.

QUESTION NO: 19

What are the methods of SandBlast Threat Emulation deployment?

A. Cloud, Appliance and Private

B. Cloud, Appliance and Hybrid

C. Cloud, Smart-1 and Hybrid

D. Cloud, OpenServer and Vmware

ANSWER: A

Explanation:

Cloud, Appliance and Private is correct because Check Point Threat Emulation is offered through three principal deployment models. The cloud deployment model sends files to Check Point's cloud-based emulation service for analysis, which provides sandboxing capacity without an organization operating its own emulation infrastructure. The appliance model uses a dedicated Check Point Threat Emulation appliance deployed in the customer environment, providing local analysis capacity and control. Private deployment enables Threat Emulation to run in a customer-controlled private-cloud or virtualized infrastructure, supporting organizations that need to retain submitted files and analysis processing within their own environment.

These deployment choices allow an organization to align malware-emulation operations with its security, performance, connectivity, and data-residency requirements. A cloud service can simplify operations and scale analysis resources, while appliance and private deployments support local control and internal processing. Check Point describes Threat Emulation as a sandboxing technology that identifies previously unknown threats by executing suspicious files in an isolated environment before they reach users or systems. See Check Point's Threat Emulation overview and the Threat Prevention Administration Guide for deployment and operational context.

QUESTION NO: 20

Which statements describe Access Roles used in Access Control rules? (Choose three.)

A. They can include users and user groups.

B. They can include machines and network locations.

C. They can be used in Access Control rules for identity-based enforcement.

D. They require HTTPS Inspection for every connection.

ANSWER: A B C

Explanation:

An Access Role is an identity-based object that can combine users, groups, machines, and network locations. Administrators can place an Access Role in an Access Control rule to apply access decisions according to the identity learned by Identity Awareness rather than only according to IP addresses.

Access Roles do not require all traffic to undergo HTTPS Inspection. HTTPS Inspection may be needed to identify users or applications in certain web scenarios, but it is not an inherent requirement for using Access Roles in policy rules. See the R81 Identity Awareness Administration Guide: Access Roles.

QUESTION NO: 21

In R81, where do you manage your Mobile Access Policy?

A. Access Control Policy

B. Through the Mobile Console

C. Shared Gateways Policy

D. From the Dedicated Mobility Tab

E. In SmartConsole, under Security Policies > Mobile Access

ANSWER: E

Explanation:

In R81, the Mobile Access policy is managed in SmartConsole under **Security Policies > Mobile Access**. This dedicated policy area is where an administrator defines the rules that govern access by Mobile Access users to internal resources. The policy is installed on the applicable Mobile Access gateway after it is configured and published.

Mobile Access is an integrated Software Blade, so its access rules are administered through the SmartConsole policy-management workflow rather than a separate legacy console. From the Mobile Access policy view, administrators can create and order rules, specify sources such as users or groups, define protected internal resources and services, and apply actions appropriate to the organization's remote-access requirements. This separation lets the Mobile Access rule base be managed specifically for remote and mobile users while retaining centralized administration, revision control, publishing, and installation in SmartConsole.

Check Point's R81 Mobile Access Administration Guide describes Mobile Access configuration and policy administration through SmartConsole. See the R81 Mobile Access Administration Guide and the R81 Security Management Administration Guide.

QUESTION NO: 22

When installing a dedicated R81 SmartEvent server. What is the recommended size of the root partition?

A. Any size

B. Less than 20GB

C. More than 10GB and less than 20GB

D. At least 20GB

ANSWER: D

Explanation:

At least 20GB is the recommended root-partition size for a dedicated R81 SmartEvent Server. The root partition contains the Gaia operating system, installed Check Point software packages, configuration files, logs, temporary files, upgrades, and other system-level data. Allocating at least this amount provides the required capacity for the base installation and normal operating overhead, while helping prevent system instability caused by root-filesystem exhaustion.

SmartEvent can process and retain substantial event data, but its event database and related high-volume data should be planned separately according to the deployment's retention, correlation, and reporting requirements. The root partition recommendation is therefore a baseline operating-system and application requirement rather than a substitute for sizing the overall SmartEvent storage design. During installation, administrators should also account for future Jumbo Hotfix Accumulators, upgrades, diagnostic data, and maintenance operations, all of which can temporarily require additional root filesystem space.

Check Point's R81 installation documentation lists the applicable platform and disk-space requirements, and the SmartEvent administration documentation provides deployment and storage-planning guidance: R81 Installation and Upgrade Guide – Requirements and R81 SmartEvent Administration Guide.

QUESTION NO: 23

CoreXL is NOT supported when one of the following features is enabled: (Choose three)

A. Route-based VPN

B. IPS

C. IPv6

D. Overlapping NAT

ANSWER: A C D

Explanation:

CoreXL compatibility must be evaluated against the specific feature combination enabled on the Security Gateway. For the CoreXL limitations represented by this question, **Route-based VPN**, **IPv6**, and **Overlapping NAT** are the relevant unsupported feature cases. Enabling one of these functions requires processing that is not compatible with the CoreXL acceleration/worker distribution model covered by the referenced compatibility guidance. Consequently, an administrator planning to use any of these features should verify the resulting CoreXL status and account for the performance impact of operating without that CoreXL support.

These restrictions are important in design and troubleshooting: CoreXL is intended to distribute traffic processing across multiple CPU cores, so an unsupported configuration can change the gateway's effective performance characteristics. Feature support evolves between releases, hotfix levels, and supported Jumbo Hotfix Accumulators; therefore, for a production R81 deployment, always validate the exact gateway version in the current release documentation and applicable SK articles before making a design decision. The historical Check Point CoreXL limitation list is available in the Performance Tuning Administration Guide; Check Point's current documentation portal is available at sc1.checkpoint.com/documents.

QUESTION NO: 24

The "MAC magic" value must be modified under the following condition:

A. There is more than one cluster connected to the same VLAN

B. A firewall cluster is configured to use Multicast for CCP traffic

C. There are more than two members in a firewall cluster

D. A firewall cluster is configured to use Broadcast for CCP traffic

ANSWER: A

Explanation:

“There is more than one cluster connected to the same VLAN” is correct. ClusterXL uses a MAC-magic value as an input when deriving the multicast MAC address used by the Cluster Control Protocol (CCP). The default value is appropriate when only one ClusterXL cluster exists on a Layer-2 segment. When separate clusters share the same VLAN, however, they can otherwise derive and use the same CCP multicast destination MAC address. Assigning a distinct MAC-magic value to each additional cluster makes its derived CCP multicast MAC address unique on that VLAN, preventing control traffic from one cluster from being received and processed by another cluster on the same broadcast domain.

The setting is therefore determined by the presence of multiple independent clusters in a common Layer-2 VLAN, rather than by the number of members in one cluster. It is a ClusterXL design consideration intended to isolate CCP communications where network topology places multiple clusters on the same segment. Check Point’s ClusterXL Administration Guide describes CCP operation and its multicast communications in the Cluster Control Protocol (CCP) section. See also the Check Point MAC magic guidance for the configuration scenario involving multiple clusters on the same network.

QUESTION NO: 25

What is the valid range for VRID value in VRRP configuration?

A. 1 - 254

B. 1 - 255

C. 0 - 254

D. 0 - 255

ANSWER: B

Explanation:

The valid VRRP Virtual Router Identifier range is **1 - 255**. A VRID identifies a virtual router within a LAN or broadcast domain, allowing VRRP routers that participate in the same redundancy group to recognize advertisements for that specific virtual router. The value must be configured consistently on every participating router for the same VRRP instance, while remaining unique among separate virtual routers operating on the same network segment. Check Point’s VRRP configuration uses this standard VRID range when defining a virtual router.

This range follows the VRRP protocol specification: the VRID field is an 8-bit identifier, with values from 1 through 255 assigned to virtual routers. Selecting an identifier in this range is therefore required for a valid VRRP configuration and enables the correct association of the virtual IP and its master/backup router state. See the VRRP standard in RFC 5798 and Check Point’s ClusterXL Administration Guide: VRRP.

QUESTION NO: 26

Bob has finished io setup provisioning a secondary security management server. Now he wants to check if the provisioning has been correct. Which of the following Check Point command can be used to check if the security management server has been installed as a primary or a secondary security management server?

A. cpprod_util MgmtIsPrimary

B. cpprod_util FwIsSecondary

C. cpprod_util MgmtIsSecondary

D. cpprod_util FwIsPrimary

ANSWER: D

Explanation:

`cpprod_util FwIsPrimary` is the appropriate command for confirming the High Availability role assigned to a Security Management Server. It queries the management installation state and reports whether the server was configured as the primary member. This makes it useful immediately after provisioning: a primary result confirms that the server owns the primary management role, while a non-primary result confirms that it was installed as the secondary management server. The command name is specifically `FwIsPrimary`, with a capital `I` in `Is`; this distinguishes it from similarly written but invalid command names. Check Point uses the primary/secondary designation in a Management High Availability deployment to establish the active management role and support synchronization between management servers. Verifying that designation at the operating-system command line is a direct way to validate that the intended secondary server was provisioned with the expected role before proceeding with operational checks such as synchronization and policy-management validation. Check Point documents the management HA architecture in its Security Management Administration Guide and lists command-line utilities in the CLI Reference Guide.

QUESTION NO: 27

Which statements about Check Point Threat Emulation are correct? (Choose two.)

A. It analyzes suspicious files in a virtual environment

B. It can use cloud-based emulation services when configured

C. It is used only to categorize websites by reputation

D. It removes macros and delivers a sanitized document without analysis

E. It is a replacement for Access Control policy enforcement

ANSWER: A B

Explanation:

Threat Emulation analyzes suspicious files in a virtual environment and **it can use cloud-based emulation services when configured** are correct. Threat Emulation identifies malicious behavior by executing or analyzing files in an isolated environment. Depending on the profile and deployment, analysis can be performed by a local emulation appliance, a cloud service, or both.

Threat Emulation is not limited to URL categorization, and it does not simply remove active content from files; that rapid sanitization function is associated with Threat Extraction.

See the Check Point Threat Emulation documentation.

QUESTION NO: 28

Which statements about Identity Awareness are correct? (Choose two.)

A. It can associate network traffic with users and groups

B. Identity Collector can obtain identity information from Active Directory domain controllers

C. It replaces Secure Internal Communication between gateways and management

D. It requires every user to authenticate with a VPN client

E. It automatically permits traffic for all identified users

ANSWER: A B

Explanation:

Identity Awareness can associate network traffic with users and groups and **Identity Collector can obtain identity information from Active Directory domain controllers** are correct. Identity Awareness allows policy rules to use users,

groups, and machines as sources. Identity Collector is an external component that collects identity information from Active Directory and shares it with Check Point management and gateways.

Identity Awareness does not replace SIC, and enabling it does not eliminate the need for correctly configured Access Control rules. Refer to the R81 Identity Awareness Administration Guide.

QUESTION NO: 29

Which characteristics apply to a Check Point Access Control policy layer? (Choose two.)

A. Layers are evaluated in their configured order

B. An inline layer is evaluated when its parent rule matches

C. Each layer must be installed separately from the policy package

D. Every packet is evaluated by all layers regardless of matching actions

E. Layers can contain only source and destination columns

ANSWER: A B

Explanation:

Layers are evaluated in their configured order and **a layer can be configured as an inline layer that is evaluated only when a parent rule matches** are correct. Ordered layers provide a structured policy flow, while inline layers enable detailed inspection rules to be evaluated conditionally from a parent rule.

An Access Control policy layer is not installed independently of its policy package, and rule evaluation does not always continue into every layer after a matching action terminates inspection. See the R81 Access Control Policy Layers documentation.

QUESTION NO: 30

Which of the following is NOT a VPN routing option available in a star community?

A. To satellites through center only.

B. To center, or through the center to other satellites, to Internet and other VPN targets.

C. To center and to other satellites through center.

D. To center only.

ANSWER: A D

Explanation:

In a Check Point star VPN community, VPN routing is defined around the center gateway and its satellite gateways. The routing choices that are provided for a satellite gateway include reaching the center and, where required, reaching other satellite gateways through that center. Check Point also supports the broader routing selection that sends traffic to the center or, through the center, to other satellites, Internet destinations, and other VPN targets. These choices explicitly describe the center as the required hub for traffic between satellites.

“To satellites through center only” is not an available routing selection because it omits connectivity to the center itself and does not match one of the defined star-community routing modes. “To center only” is a valid routing restriction in a star community: it allows the satellite to establish VPN connectivity solely with the center gateway and prevents satellite-to-satellite connectivity. Review the Star Community VPN-routing configuration guidance in the Check Point Security Management Administration Guide and the related Star Community documentation.

QUESTION NO: 31

After making modifications to the $CVPNDIR/conf/cvpnd.C file, how would you restart the daemon?

A. cvpnd_restart

B. cvpnd_restart

C. cvpnd restart

D. cvpnrestart

ANSWER: B

Explanation:

`cvpnd_restart` is the appropriate command to restart the Check Point VPN daemon after changing `$CVPNDIR/conf/cvpnd.C`. The `cvpnd.C` file contains settings used by the CVPN/Mobile Access VPN daemon, and restarting that daemon causes it to terminate its running instance and launch a new instance that reads the updated configuration. This is important because editing the file alone does not necessarily make the active daemon reload its settings.

The command should be run in the applicable Check Point environment, typically from an expert shell on the Security Gateway or Mobile Access gateway where the CVPN service is installed. Operationally, plan the restart during an appropriate maintenance window when possible, because active remote-access or Mobile Access connections can be affected while the daemon restarts. Check Point's Mobile Access documentation describes the Mobile Access gateway components and their operation: R81 Mobile Access Administration Guide. For command-line operational context on Gaia systems, consult the R81 CLI Reference Guide.

QUESTION NO: 32

How many versions, besides the destination version, are supported in a Multi-Version Cluster Upgrade?

A. 1

B. 3

C. 2

D. 4

ANSWER: B

Explanation:

3 is correct. In a Check Point Multi-Version Cluster Upgrade, the destination version can coexist with as many as three other supported software versions during the upgrade process. This capability allows an administrator to replace cluster members gradually rather than requiring every member to be upgraded in one maintenance event. Members running the destination release are introduced into the existing cluster while members on earlier versions are removed or upgraded in a controlled sequence.

The ability to support three versions in addition to the destination version is particularly useful for large or distributed deployments, where hardware replacement, staged rollout procedures, and validation periods may require multiple generations of members to remain part of the same cluster temporarily. Check Point's Multi-Version Cluster Upgrade design preserves service continuity by enabling the cluster to operate through these transitional states, provided that the documented compatibility, configuration, and upgrade-path requirements are met. The mixed-version state is an upgrade mechanism and should be completed according to the planned migration procedure rather than treated as a permanent operating configuration.

See Check Point's R81.20 Installation and Upgrade Guide: Multi-Version Cluster Upgrade and sk107042: Multi-Version Cluster Upgrade.

QUESTION NO: 33

Vanessa is firewall administrator in her company. Her company is using Check Point firewall on a central and several remote locations which are managed centrally by R77.30 Security Management Server. On central location is installed R77.30 Gateway on Open server. Remote locations are using Check Point UTM-1570 series appliances with R75.30 and some of them are using a UTM-1-Edge-X or Edge-W with latest available firmware. She is in process of migrating to R81.

What can cause Vanessa unnecessary problems, if she didn't check all requirements for migration to R81?

A. Missing an installed R77.20 Add-on on Security Management Server

B. Unsupported firmware on UTM-1 Edge-W appliance

C. Unsupported version on UTM-1 570 series appliance

D. Unsupported appliances on remote locations

ANSWER: D

Explanation:

Unsupported appliances on remote locations is correct because a Security Management Server upgrade must account for the management compatibility of every deployed gateway and appliance, not only the centrally installed Open Server gateway. The described estate includes legacy UTM-1 Edge-X and Edge-W devices. These appliances use the legacy Edge platform and firmware train, which is not supported for central management by an R81 Security Management Server. Installing the newest firmware available for an appliance does not make the appliance compatible with a newer management release when that platform has reached its management-support boundary.

Before migrating the management server, Vanessa must identify such devices, plan their replacement or retirement, and remove or migrate their management relationship as appropriate. Otherwise, the management upgrade can leave remote sites unmanaged or prevent normal policy administration for those devices. This assessment is particularly important in environments upgraded over multiple Check Point releases, where older appliances can remain operational on legacy releases but are outside the compatibility scope of the target management version. Check Point's upgrade documentation requires administrators to validate supported upgrade paths and compatible products before performing an upgrade. See the R81 Installation and Upgrade Guide and the R81 Release Notes.

QUESTION NO: 34

John is using Management HA. Which Smartcenter should be connected to for making changes?

A. secondary Smartcenter

B. active SmartCenter

C. connect virtual IP of Smartcenter HA

D. primary Smartcenter

ANSWER: B

Explanation:

In a Check Point Management High Availability deployment, administrators make configuration changes on the active SmartCenter (Security Management Server). The active server is the authoritative management peer: it processes administrator sessions, maintains the current management database state, and is the server from which policy installation and management operations are performed. Management HA continuously synchronizes the management database and relevant configuration data to the standby peer, so that the standby server can assume the active role if a failure occurs. Connecting to the active SmartCenter therefore ensures that changes to objects, rules, gateways, policy packages, and other management settings are committed through the server currently responsible for management operations and are replicated as part of HA synchronization. Before making changes, an administrator should confirm which member currently has the Active status, particularly after maintenance or a failover event. Check Point documents the active/standby operation

and synchronization behavior in its R81.10 Security Management Administration Guide. Additional deployment and operational guidance is available through the Check Point Management High Availability documentation.

QUESTION NO: 35

Which file gives you a list of all security servers in use, including port number?

A. $FWDIR/conf/conf.conf

B. $FWDIR/conf/servers.conf

C. $FWDIR/conf/fwauthd.conf

D. $FWDIR/conf/serversd.conf

ANSWER: C

Explanation:

`$FWDIR/conf/fwauthd.conf` is the Security Gateway configuration file that defines the Security Server processes managed by `fwauthd`, the FireWall-1 authentication daemon. Its entries associate a Security Server service with the TCP port on which that service operates, making it the appropriate local file to consult when identifying the Security Servers enabled on a gateway and their port numbers. Typical entries cover services handled through the Security Server framework, such as FTP, HTTP, SMTP, and CVP-related processing, depending on the gateway configuration and installed features.

This file is particularly useful during gateway-level troubleshooting because it reflects the daemon's service/port configuration rather than merely presenting policy objects in SmartConsole. Administrators should inspect it from the applicable gateway installation, using the `$FWDIR` environment path, and treat any manual modifications cautiously: configuration should be backed up and validated because an invalid daemon configuration can affect Security Server operation. Check Point's Security Management documentation describes Security Servers and their gateway role in traffic inspection and protocol handling: Security Servers. Additional Check Point R81 documentation is available through the Security Management Administration Guide.

QUESTION NO: 36

A user complains that some Internet resources are not available. The Administrator is having issues seeing it packets are being dropped at the firewall (not seeing drops in logs). What is the solution to troubleshoot the issue?

A. run fw unloadlocal" on the relevant gateway and check the ping again

B. run "cpstop" on the relevant gateway and check the ping again

C. run ''fw log" on the relevant gateway

D. run ''fw ctl zdebug drop" on the relevant gateway

ANSWER: D

Explanation:

Running `fw ctl zdebug drop` on the affected Security Gateway is the appropriate live troubleshooting method when traffic is suspected of being dropped but no corresponding drop logs are visible. This command enables a temporary kernel debug focused on packet drops and prints drop information directly to the command-line session as it occurs. The output can show the source and destination details, protocol, ports, interface direction, and—most importantly—the kernel drop reason. That lets the administrator correlate a failed connectivity test, such as a ping or attempt to reach the unavailable Internet resource, with the exact enforcement or processing point responsible for the drop.

This is especially useful because not every packet drop necessarily produces an ordinary log entry. Drops can occur at different points in the Security Gateway inspection path, including SecureXL, the firewall kernel, anti-spoofing checks, policy

enforcement, or connection-state processing. By generating the relevant traffic while the debug is active, the administrator obtains immediate diagnostic evidence without disabling security enforcement or stopping Check Point services. The command is intended for targeted, short-duration investigation and should be stopped after gathering the required output. Check Point documents `fw ctl zdebug` as a kernel-debug utility in its R81.10 CLI Reference Guide; related packet-flow and troubleshooting guidance is available in the Security Management Administration Guide.

QUESTION NO: 37

Which of the following describes how Threat Extraction functions?

A. Detect threats and provides a detailed report of discovered threats.

B. Proactively detects threats.

C. Delivers file with original content.

D. Delivers PDF versions of original files with active content removed.

ANSWER: D

Explanation:

Threat Extraction delivers PDF versions of original files with active content removed. It is a prevention technology intended to let users receive a safe, usable document immediately rather than waiting for full malware-analysis results. For supported document types, Threat Extraction reconstructs the file as a PDF and excludes active or potentially exploitable elements, such as macros, embedded objects, scripts, and hyperlinks where applicable. The resulting sanitized PDF preserves the document's viewable content while removing components that could execute code or be used as an attack vector.

This approach is especially useful for files received through common delivery channels such as email and web downloads. A user can continue working with the extracted document while the original file may be sent to Threat Emulation for deeper inspection. If the organization's policy permits and subsequent analysis determines that the original is safe, the original file can be made available according to the configured workflow. Check Point describes Threat Extraction as delivering a safe version of a file by removing potentially malicious content, with conversion to PDF as its standard extraction method for supported files. See the Check Point Threat Extraction Administration Guide and the Check Point Threat Prevention overview.

QUESTION NO: 38

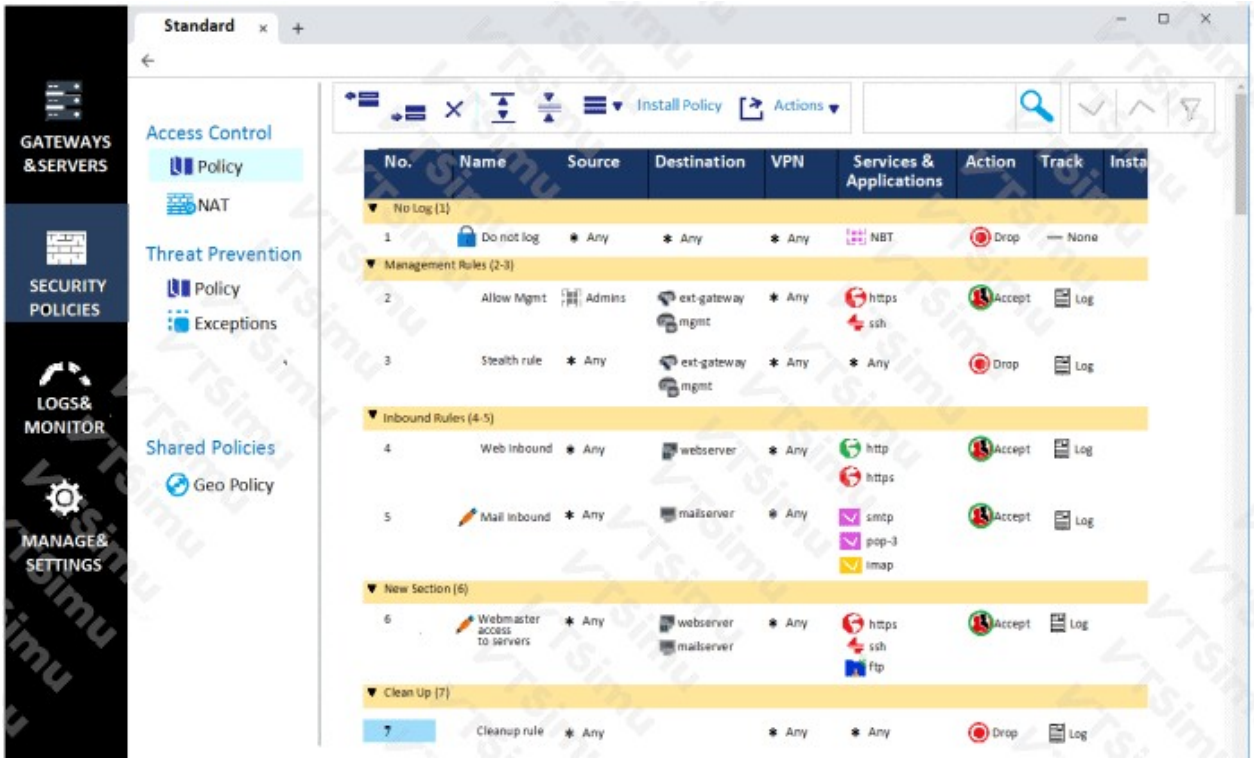

What can we infer about the recent changes made to the Rule Base?

A. Rule 7 was created by the 'admin' administrator in the current session

B. 8 changes have been made by administrators since the last policy installation

C. The rules 1, 5 and 6 cannot be edited by the 'admin' administrator

D. Rule 1 and object webserver are locked by another administrator

ANSWER: D

Explanation:

Rule 1 and object webserver are locked by another administrator is correct. SmartConsole supports concurrent administration, but it protects configuration consistency by placing locks on items that an administrator is actively editing. In the Rule Base, a lock indicator on a rule shows that the rule is currently reserved by a different administrator's session. Likewise, a lock shown for an object indicates that another administrator is editing that specific object. These locks are item-level controls: they identify the affected rule or object and prevent conflicting simultaneous edits while the other administrator has the item checked out or is modifying it.

The displayed indicators therefore establish that both Rule 1 and the object named *webserver* are under another administrator's edit lock. This is a current collaboration state in the Security Management Server database, rather than a record of who created a rule, a count of all changes since a policy installation, or a general authorization restriction on the local administrator. Check Point documents concurrent administration and locking behavior in its Security Management Administration Guide. The complete R81 documentation set is also available from the Check Point Documentation portal.

QUESTION NO: 39

Office mode means that:

A. SecurID client assigns a routable MAC address. After the user authenticates for a tunnel, the VPN gateway assigns a routable IP address to the remote client.

B. Users authenticate with an Internet browser and use secure HTTPS connection.

C. Local ISP (Internet service Provider) assigns a non-routable IP address to the remote user.

D. Allows a security gateway to assign a remote client an IP address. After the user authenticates for a tunnel, the VPN gateway assigns a routable IP address to the remote client.

ANSWER: D

Explanation:

“Allows a security gateway to assign a remote client an IP address. After the user authenticates for a tunnel, the VPN gateway assigns a routable IP address to the remote client.” correctly describes Check Point Office Mode. Office Mode is a Remote Access VPN feature through which the Security Gateway allocates an internal virtual IP address to a connecting client after successful authentication and tunnel establishment. The address is drawn from an Office Mode address pool configured for the gateway and represents the remote user within the protected network. This virtual address lets internal resources identify, route return traffic to, and apply Access Control policy to the remote client as though it were connected through the corporate network. The client retains its existing physical network connection and ISP-provided address for Internet transport, while the Office Mode address is used for traffic carried inside the encrypted VPN tunnel. Administrators can configure address allocation through a dedicated pool, DHCP, or other supported assignment mechanisms, depending on the deployment. Check Point documents Office Mode as the mechanism used to assign IP addresses to Remote Access VPN clients in the R81 VPN Administration Guide. Related Remote Access VPN configuration concepts are also covered in the R81 Security Management Administration Guide.

QUESTION NO: 40

GAiA Software update packages can be imported and installed offline in situation where:

A. Security Gateway with GAiA does NOT have SFTP access to Internet

B. Security Gateway with GAiA does NOT have access to Internet.

C. Security Gateway with GAiA does NOT have SSH access to Internet.

D. The desired CPUSE package is ONLY available in the Check Point CLOUD.

ANSWER: B

Explanation:

“Security Gateway with GAiA does NOT have access to Internet.” is correct because Check Point supports an offline workflow for obtaining and applying Gaia software updates. In this workflow, an administrator uses a separate internet-connected computer to download the applicable update package from Check Point’s download resources, transfers the package to the target Gateway or Management Server, and then imports and installs it locally through the relevant Gaia or CPUSE update interface. This allows the appliance to remain isolated from the public internet while still receiving supported updates.

The offline method is particularly important for restricted, air-gapped, or tightly controlled environments, where direct outbound connectivity from security infrastructure is prohibited by policy. Package compatibility must still be verified for the installed Gaia and Check Point release, and administrators should follow the documented installation procedure, including any required pre-installation checks and reboot planning. Check Point documents software-update operations in the Gaia Administration Guide; its product downloads are available through the Check Point Download Center guidance.

QUESTION NO: 41

Which Check Point feature enables application scanning and the detection?

A. Application Dictionary

B. AppWiki

C. Application Library

D. CPApp

E. Application Control

ANSWER: E

Explanation:

Application Control is the Check Point feature that identifies and controls network applications. It performs application-aware inspection to recognize applications and application functions regardless of the port, protocol, evasive technique, or encryption context used by the traffic. This detection capability lets an administrator create policy rules based on the actual application in use, rather than relying only on IP addresses, services, and ports. Application Control uses Check Point's continuously maintained application signatures and categorization data to classify traffic, making application visibility and enforcement available directly in the Access Control policy. It can also be combined with URL Filtering to apply policy to web applications and destinations. AppWiki is a useful Check Point online reference for researching applications, categories, risks, and supported controls, but it is not the gateway feature that scans traffic and performs application detection. See Check Point's Application Control overview and the R81.10 Security Management Administration Guide for configuration and policy details.