PassQueen
a.k.a dumpsqueen.com

Check Point Certified Security Expert R82

Checkpoint 156-315.82

Version Demo

Total Demo Questions: 14

Total Premium Questions: 141

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Advanced Management Server Configuration	75
Topic 2, Advanced Security Gateway Configuration	23
Topic 3, Advanced Firewall Configuration	5
Topic 4, Advanced VPN Configuration	21
Topic 5, Monitoring and Optimization	17
Total	141

QUESTION NO: 1

Which of the following is a trigger for synchronization between Active and Standby servers?

A. Publishing a session in SmartConsole.

B. Making a change in a network object and clicking OK.

C. Running the Save operation from the SmartConsole toolbar or menu.

D. After 10 seconds of inactivity in SmartConsole.

ANSWER: A

Explanation:

Publishing a session in SmartConsole is a trigger for synchronization from the Active Security Management Server to its Standby peer in a Management High Availability deployment. SmartConsole uses a session-based workflow: an administrator can make and review changes privately within a session, but those changes do not become part of the shared, committed management database until the session is published. When publication occurs, the Active server synchronizes the relevant database updates with the Standby server, helping ensure that the standby management environment has a current copy of the configuration and can take over if needed.

This behavior is important operationally because it separates work in progress from the configuration that is available to other administrators and to the high-availability peer. Administrators should therefore publish completed sessions intentionally, especially before maintenance or a planned management-server role change. In addition to publication, management synchronization also occurs periodically; however, among the listed actions, publishing is the applicable synchronization trigger. Check Point documents that unpublished sessions are not synchronized to Standby servers and describes the synchronization behavior in its Management High Availability guidance. See Check Point R82 Security Management Administration Guide: Management High Availability and Check Point R82 Security Management Administration Guide: Publishing Sessions.

QUESTION NO: 2

How many members are supported by an ElasticXL Cluster?

A. Maximum three members per site with a maximum of three sites.

B. Three members per site with a maximum of two sites.

C. Maximum two members per site with a maximum of three sites.

D. Up to four members per site with a maximum of two sites.

ANSWER: B

Explanation:

Three members per site with a maximum of two sites. is correct. In Check Point R82, an ElasticXL Cluster supports up to three ElasticXL Cluster Members at an individual site. ElasticXL can be deployed across two sites, resulting in a maximum supported total of six members in the complete ElasticXL Cluster. This limit is important when designing both capacity and resiliency: a site can use up to three members for scale, while the second site provides the supported multi-site deployment structure. Therefore, the supported answer expresses the documented per-site limit and the maximum number of sites rather than only stating the aggregate total. Administrators should plan the ElasticXL topology within these bounds, including any required inter-site connectivity, synchronization, and traffic-distribution design. Where an environment needs to scale beyond the ElasticXL member limit, the appropriate Check Point scale-out architecture should be evaluated rather than extending an ElasticXL Cluster beyond its supported design. Check Point documents the member and site limitation in the R82 ElasticXL Administration Guide's Important Notes section. See the ElasticXL Important Notes and the R82 ElasticXL Overview.

QUESTION NO: 3

A company requires branch-to-branch VPN traffic to pass through a central Security Gateway, where it can be inspected and controlled. The branch gateways must not establish direct VPN tunnels with one another. Which VPN Community topology should the administrator configure?

A. A Meshed Community containing all branch and central Security Gateways

B. A Route-based VPN without a VPN Community

C. A Domain-based VPN without a Center Gateway

D. A Star Community with the central Security Gateway configured as the Center Gateway

ANSWER: D

Explanation:

A Star Community with the central Security Gateway configured as the Center Gateway is correct. In a Star VPN Community, Satellite Gateways establish VPN tunnels with the Center Gateway. Traffic between two satellites is sent through the center, which allows the central gateway to enforce its security policy and inspect the traffic.

A Meshed Community permits each participating gateway to establish direct tunnels with every other peer. It would not meet the requirement to force branch-to-branch traffic through a central inspection point. VPN Domains define protected networks, but do not determine whether the community uses a star or meshed topology.

QUESTION NO: 4

Where can a Firewall administrator configure VPN routes between Security Gateways?

A. vpn_route.conf on the Security Management Server

B. Via Gaia Portal or CLI on the Security Gateway

C. VTI_route.conf on the Security Management Server

D. vpn_route.conf on the Security Gateway

ANSWER: A

Explanation:

`vpn_route.conf on the Security Management Server` is correct because Check Point VPN routing for managed Security Gateways is centrally defined in the `vpn_route.conf` configuration file on the Security Management Server. This file is used to specify routing behavior for traffic that must traverse VPN tunnels between Security Gateways, including cases where the normal VPN-domain configuration does not provide the required forwarding path. The administrator defines the applicable source and destination networks and the relevant VPN gateway routing details in the centrally managed configuration, then installs the policy so the configuration is distributed and applied as part of the managed environment. This central approach ensures that VPN routing remains consistent with the management policy and topology rather than being maintained as an unrelated local operating-system route on an individual gateway. It is particularly important in more complex VPN deployments, such as hub-and-spoke or overlapping routing scenarios, where traffic must be directed through a specific VPN peer. Check Point's Security Management documentation covers VPN administration and centrally managed gateway configuration in the R82 Security Management Administration Guide. Additional official Check Point knowledge resources can be located through the Check Point Support knowledge-base search for vpn_route.conf.

QUESTION NO: 5

Which process is responsible for the code generation and compilation of Legacy Dump files?

A. FWM

B. CPM

C. Stateful Compiler

D. Inspect Engine

ANSWER: A

Explanation:

FWM is responsible for code generation and compilation of Legacy Dump files. The FWM process is the legacy Firewall Management component that remains part of the Security Management Server architecture for functions that depend on the traditional policy-compilation workflow. During policy installation, it participates in producing the compiled policy representation used by the legacy management and gateway mechanisms. These compiled artifacts are commonly referred to as dump files in Check Point documentation and troubleshooting contexts.

Although modern Check Point management relies heavily on the CPM process and the management database, FWM continues to provide compatibility with legacy policy-management operations and compilation components. This distinction matters when examining management-server processes, reviewing policy-installation behavior, or troubleshooting issues involving legacy dump-file creation. Check Point's Security Management documentation describes the roles of management-server processes, including the legacy FWM process, while the R82 documentation provides the current administrative context for Security Management Server architecture and policy operations. See the Check Point R82 Security Management Administration Guide and the R81.20 Security Management process reference.

QUESTION NO: 6

What is correct regarding the target device for deploying SmartEvent components?

A. SmartEvent is just a blade on the Security Management Server and can be activated on a Primary or Secondary SMS only.

B. SmartEvent works by correlating logs; hence, it has to be deployed on each Log Server. If any Log Server does not include SmartEvent components, then its logs will not be correlated.

C. SmartEvent is always a dedicated standalone exclusive device.

D. SmartEvent can be integrated with the Security Management Server or deployed on a dedicated Log or SmartEvent Server.

ANSWER: D

Explanation:

SmartEvent can be integrated with the Security Management Server or deployed on a dedicated Log or SmartEvent Server. This reflects SmartEvent's supported distributed architecture: smaller or lower-volume environments can run SmartEvent functions on the Security Management Server, while environments with greater logging, correlation, retention, or reporting requirements can separate those functions onto dedicated servers. This flexibility permits an administrator to size the deployment according to the number of gateways, expected event rate, log-storage requirements, and the processing load created by event correlation and views.

In a dedicated design, the SmartEvent Server manages the SmartEvent configuration and provides the SmartEvent interface, while a SmartEvent Correlation Unit can process logs and identify correlated security events. Log Server functionality can also be part of the architecture, allowing log collection and analysis responsibilities to be distributed as needed. The Security Management Server remains connected to and manages the overall SmartEvent deployment rather than requiring every source of logs to host all SmartEvent functions. Check Point's administration guidance describes both SmartEvent architecture and deployment choices, including use with management and dedicated components. See the Check Point R82 SmartEvent architecture documentation and the R82 SmartEvent deployment documentation.

QUESTION NO: 7

What are the key components of an Access Role object?

A. Name, Subnet, Mask-length, User Group, LDAP Account Unit

B. Name, IP Address, Mask-Length, LDAP Account Unit, Remote Access Client

C. Name, LDAP Account Unit, Remote Access Client, Subnet, Host Object Type

D. Name, Networks, Users, Machines, Remote Access Clients

ANSWER: D

Explanation:

Name, Networks, Users, Machines, Remote Access Clients is correct because an Access Role is the Check Point policy object used to combine identity-based and network-context criteria into a reusable source or destination definition for Access Control rules. The Name identifies the object in SmartConsole. Networks specifies the network locations from which the connection originates, while Users identifies individual users or groups obtained through configured identity sources. Machines supplies endpoint or computer-based identity criteria, enabling policy matching based on managed or identified computers as well as the user identity. Remote Access Clients provides the context needed for users connecting through Check Point Remote Access VPN.

These dimensions let an administrator express policy requirements such as allowing a specified user group access to a service only when members connect from approved networks, use designated corporate machines, or arrive through the remote-access environment. Rather than being a simple address or directory-reference object, the Access Role object brings these identity and location attributes together so that a single policy rule can match the required combination. Check Point documents Access Roles as containing criteria for networks, users and user groups, computers and computer groups, and Remote Access VPN clients. See the R82 Identity Awareness Administration Guide: Access Roles and the R82 Security Management Administration Guide: Access Roles.

QUESTION NO: 8

What is the CLI command to check the Deployment Agent Build Number?

A. show deployment agent -v

B. show installer version

C. show deployment agent --version

D. show installer status build

ANSWER: D

Explanation:

The correct command is `show installer status build`. This is the Gaia Clish CPUSE command that displays the build number of the CPUSE Deployment Agent installed on the Security Management Server or Security Gateway. CPUSE, the Check Point Upgrade Service Engine, uses the Deployment Agent to obtain, manage, and deploy Check Point software packages and hotfixes. Checking the agent build is useful when validating that the environment has the expected CPUSE components before downloading packages, performing upgrades, or troubleshooting deployment behavior. The `show installer status` command reports the installer's general status, while adding the `build` keyword specifically requests the Deployment Agent build-number information required by the question. Check Point documents CPUSE command usage in its Installation and Upgrade guide, including the installer status commands: CPUSE CLI commands. Additional context on CPUSE operation and package deployment is available in the Check Point R82 Installation and Upgrade Guide.

QUESTION NO: 9

Choose the correct names for the bonding interfaces that are present by default in an ElasticXL configuration.

A. Mgmt, eth1-Sync

B. magg1, Sync

C. magg1, eth1

D. Mgmtagg1, Syncagg1

ANSWER: B

Explanation:

magg1, Sync is correct because these are the default logical bonding-interface names created by ElasticXL for its management and synchronization connectivity. During ElasticXL configuration, Check Point automatically transforms the applicable physical interfaces into subordinate members of the relevant bonds. The management-side bond is named `magg1`, while the synchronization bond is named `Sync`. These logical bond names are the interfaces an administrator should recognize when reviewing the ElasticXL network configuration, operational status, or interface-related output in Gaia and Check Point commands.

The distinction between a bond and its member interfaces is important in ElasticXL. A physical management interface can be incorporated as a member of the `magg1` bond, and a physical synchronization interface can be renamed and incorporated beneath the `Sync` bond. Consequently, the requested names are the top-level bond interfaces rather than the physical or subordinate interface identifiers. Check Point documents ElasticXL interface behavior and configuration in its R82 Scalable Platforms Administration Guide. Additional ElasticXL deployment context is available in the ElasticXL overview.

QUESTION NO: 10

Two Management Servers in a Management High Availability deployment lose communication with each other. An administrator manually promotes the Standby server while the original Active server is still accepting SmartConsole connections. Which condition must the administrator avoid?

A. Collision Mode, with both Management Servers operating as Active.

B. A normal load-sharing condition, with both servers processing policy installations.

C. A standard scheduled synchronization condition, with both servers operating as Standby.

D. A Gateway ClusterXL state-synchronization condition.

ANSWER: A

Explanation:

The administrator must avoid making independent management changes on both servers. If both Management Servers identify themselves as Active, the environment is in Collision Mode. Separate policy or object changes can create divergent management databases and complicate recovery.

The appropriate recovery objective is to restore one authoritative Active server and return the other server to Standby status with a synchronized database. Maintaining both servers as Active is not a load-sharing design for management operations.

QUESTION NO: 11

What is crucial in translating services, specifically destination ports, in a NAT rule?

A. This can only be accomplished with the Automatic NAT Rule with “Translate Destination on Server Side” enabled.

B. This can only be accomplished with Automatic NAT Rule in conjunction with Bi-Directional NAT.

C. This can only be accomplished with the Automatic NAT Rule with “Automatic ARP Configuration” enabled.

D. This has to be done with a Manual NAT Rule.

ANSWER: D

Explanation:

This has to be done with a Manual NAT Rule. Service translation changes the destination service in addition to, or independently of, address translation—for example, translating traffic received on a public IP address at TCP port 8443 to an internal server's TCP port 443. This requires explicit control of the Original Service and Translated Service fields in the NAT rule. A Manual NAT rule supports this packet-level definition, allowing the administrator to specify original source, destination, and service values and their corresponding translated values. The rule can therefore match the inbound service exposed to clients and rewrite it to the actual service on the protected host while maintaining the necessary address translation and rule order. Check Point's NAT guidance identifies translating services, including destination-port translation, as a scenario for manually defined NAT rules. This is important operationally because the NAT rule must be evaluated with the intended original service and must specify the translated service precisely, ensuring that traffic reaches the application listener on the correct internal port. See the Check Point R82 Security Management Administration Guide NAT introduction and the NAT rule configuration documentation.

QUESTION NO: 12

Check Point Security Gateways support two methods of identifying traffic to include in the VPN. What are the two methods?

A. Domain-based and Community-based

B. Domain-based and Route-based

C. Kernel-based and INSPECT-based

D. Community-based and Route-based

ANSWER: B

Explanation:

Domain-based and Route-based is correct because Check Point supports these two models to determine which traffic enters a site-to-site VPN tunnel. In a domain-based VPN, the gateway identifies interesting traffic using the VPN Domain configured for each participating gateway. A VPN Domain represents the internal IP addresses, networks, or host objects located behind that gateway that are protected through the VPN. Traffic whose source and destination match the relevant VPN Domains is selected for encryption according to the applicable VPN community and its settings.

In a route-based VPN, the tunnel is associated with a Virtual Tunnel Interface (VTI). The VTI operates as a routable interface, so the gateway uses routing decisions to forward qualifying packets into the VPN tunnel. This model is particularly useful when routing protocols, dynamic routing, or more flexible network topologies are required. Thus, domain-based VPN uses defined protected domains to identify VPN traffic, while route-based VPN uses routing through a VTI. Check Point documents these VPN routing and configuration approaches in its R82 Security Management Administration Guide and describes VPN-domain concepts in the VPN Domains documentation.

QUESTION NO: 13

To which directory does CPTA transfer policy files to the Security Gateway?

A. $FWDIR/state/_tmp/FW1

B. $FWDIR/state/local/FW1

C. $CPDIR/state/tmp/FW1

ANSWER: A

Explanation:

`$FWDIR/state/_tmp/FW1` is the staging directory used when CPTA transfers Security Policy files to a Security Gateway. During policy installation, the management components prepare and transmit the compiled policy artifacts, while gateway-side processes use the temporary `_tmp/FW1` location as part of the controlled transfer and installation workflow. Using a temporary directory prevents a partially received policy from being treated as an active local policy. After the transfer and validation stages complete, the gateway can process the policy and apply it to the relevant inspection components. The `$FWDIR` variable represents the Check Point Firewall installation directory, so this path is relative to the installed Security Gateway software rather than to a generic operating-system temporary directory. This behavior is consistent with Check Point's policy-management architecture, in which Security Management Server components install a policy on managed gateways and the gateways locally receive and load the generated policy package. See Check Point's R82 Security Management Administration Guide: Policy Management and the R82 Security Gateway Guide for the management-to-gateway policy-installation model and gateway file-system context.

QUESTION NO: 14

Any VPN Gateway that can establish a direct VPN tunnel with any peer Gateway is a member of which VPN Community?

A. Direct Community

B. Any Community

C. Star Community

D. Mesh Community

ANSWER: D

Explanation:

Mesh Community is correct because it implements a fully meshed site-to-site VPN topology. In a Check Point Meshed VPN Community, every participating Security Gateway is a peer of every other participating Security Gateway. Consequently, a gateway can negotiate and establish a direct IPsec VPN tunnel to any other gateway in the same community when traffic matches the relevant encryption domain and VPN configuration. This is the topology described by the question's requirement that any VPN Gateway be able to establish a direct tunnel with any peer Gateway.

Check Point uses VPN Communities to define the gateways, topology, and shared VPN properties for site-to-site VPN deployments. A meshed design is appropriate when branch or site gateways must communicate directly with one another, rather than requiring their traffic to traverse a designated central gateway. The Security Management documentation identifies Star and Meshed as the supported VPN-community topology models and describes the direct peer-to-peer connectivity provided by a meshed community. See the Check Point R82 Security Management Administration Guide: VPN Communities and the Check Point Security Management overview.