

Check Point Certified Troubleshooting Expert - R81.20 (CCTE)

Checkpoint 156-587

Version Demo

Total Demo Questions: 12

Total Premium Questions: 128

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

QUESTION NO: 1

What function receives the AD log event information?

- A. FWD
- B. CPD
- C. PEP
- D. ADLOG

ANSWER: D

Explanation:

ADLOG is the Check Point Identity Awareness process responsible for receiving Active Directory log-event information. In an AD Query deployment, the Security Gateway obtains identity information by monitoring domain-controller security events, particularly logon and logoff activity. ADLOG handles the reception and processing of those events so that user-to-IP address mappings can be made available to the Identity Awareness infrastructure and subsequently used when evaluating Access Control policy rules based on users, groups, or identity attributes.

This function is important because AD log events provide the evidence needed to associate a network connection with an authenticated AD user. The resulting identity data enables the gateway to enforce policies using user identity rather than relying solely on source IP addresses. Check Point documents AD Query as a method that acquires identities from Active Directory domain-controller event logs and identifies the ADLOG process as part of the Identity Awareness architecture. See the [Check Point R81.20 Identity Awareness Administration Guide: AD Query](#) and the [Identity Awareness Processes](#).

QUESTION NO: 2

For Identity Awareness, what is the PDP process?

- A. Identity server
- B. Log Sifter
- C. Captive Portal Service
- D. UserAuth Database

ANSWER: A

Explanation:

Identity server is correct. In Check Point Identity Awareness, the PDP (Policy Decision Point) is the Security Gateway component that obtains, correlates, and maintains identity information for users and machines. It receives identities from the Identity Awareness identity sources configured for the environment, which can include AD Query, Identity Agents, Captive Portal authentication, Terminal Servers, and RADIUS Accounting. The PDP associates the learned identity with network context, such as an IP address, so that the gateway has the information needed to apply access-control rules that reference users, groups, or Access Roles. The PDP works with the PEP (Policy Enforcement Point) on the gateway: the PDP supplies identity-related policy decisions and identity mappings, while enforcement is applied to connections traversing the gateway. This architecture lets security policy be based on the authenticated user rather than solely on a source address, and identity mappings can be updated as users log in, log out, or move between endpoints. Check Point documents the PDP and PEP roles in its [R81.20 Identity Awareness Administration Guide](#). Additional deployment and identity-source context is available in Check Point's [Identity Awareness configuration guidance](#).

QUESTION NO: 3

You modified kernel parameters and after rebooting the gateway, a lot of production traffic gets dropped and the gateway acts strangely. What should you do?"

- A. Run command `fw ctl set int fw1_kernel_all_disable=1`
- B. Restore `fwkern.conf` from backup and reboot the gateway
- C. run `fw unloadlocal` to remove parameters from kernel
- D. Remove all kernel parameters from `fwkem.conf` and reboot

ANSWER: B

Explanation:

Restore `fwkern.conf` from backup and reboot the gateway is the appropriate recovery action when instability and widespread traffic loss begin immediately after modified kernel parameters take effect. Check Point kernel parameters can alter low-level Firewall Kernel behavior, including connection handling, inspection, acceleration, and resource processing. Because these settings are loaded during the gateway startup process, a bad, unsupported, or incompatible parameter can produce persistent and unpredictable behavior after every reboot.

Restoring the known-good backup returns the gateway to the exact previously working kernel-parameter baseline rather than attempting to diagnose or remove settings individually while production traffic is affected. Rebooting is required so the Firewall Kernel starts cleanly and loads the restored configuration, ensuring that settings introduced by the recent change are no longer active. This approach also preserves any valid pre-existing custom parameters that may have been required in the production environment. After service is stable, the original change should be reviewed, validated against the applicable Check Point version and hotfix level, and tested in a controlled environment before being reintroduced. Check Point documents kernel-parameter management in its [R81.20 Security Management Administration Guide](#); additional product documentation is available through the [Check Point Support Center](#).

QUESTION NO: 4

A gateway is expected to accelerate high-volume traffic, but a specific application flow appears to be processed by the Firewall Kernel instead. The administrator needs to view the connections currently accelerated by SecureXL. Which command is most appropriate?

- A. `fwaccel stats`
- B. `fw tab -t connections -s`
- C. `fwaccel conns`
- D. `fw ctl chain`

ANSWER: C

Explanation:

`fwaccel conns` displays connections that are currently accelerated by SecureXL. It is useful for determining whether the affected flow is present in the acceleration path before investigating features or policy conditions that can require non-accelerated processing.

`fwaccel stats` provides aggregate acceleration statistics, but it does not list the individual accelerated connections. `fw tab -t connections -s` shows connection-table utilization rather than SecureXL acceleration state for flows.

QUESTION NO: 5

Where will the usermode core files be located?

- A. `$FWDIRVar/log/dump/usermode`

- B. /var/suroot
- C. /var/log/dump/usermode
- D. \$CPDIR/var/log/dump/usermode
- E. \$FWDIR/log/dump/usermode

ANSWER: E

Explanation:

\$FWDIR/log/dump/usermode is the standard Check Point directory for user-mode core-dump files. A core file is written when a user-space Check Point process terminates unexpectedly, preserving process memory and state for subsequent crash analysis. \$FWDIR identifies the installed Firewall product directory, so using this variable is important because its physical location can vary by release and installation. The log/dump/usermode subdirectory is specifically intended for dumps from user-mode processes, such as management and gateway daemons. During troubleshooting, an administrator can inspect this location to establish whether a process crash occurred, identify the affected process from the dump filename, and collect the relevant artifacts for Check Point Support. Core files can be large and may contain sensitive process-memory data; they should therefore be handled according to the organization's support-collection and data-protection procedures. Check Point's official documentation and support resources provide the current procedures for collecting and analyzing diagnostic files: [Check Point Support Center](#) and [Gaia Administration Guide](#).

QUESTION NO: 6

Like a Site-to-Site VPN between two Security Gateways, a Remote Access VPN relies on the Internet Key Exchange (IKE) what types of keys are generated by IKE during negotiation?

- A. Produce a symmetric key on both sides
- B. Produce an asymmetric key on both sides
- C. Symmetric keys based on pre-shared secret
- D. Produce a pair of public and private keys

ANSWER: A

Explanation:

"Produce a symmetric key on both sides" is correct because IKE negotiates security associations and derives shared symmetric keying material that both VPN peers can use to protect the VPN connection. During IKE negotiation, the peers perform a Diffie-Hellman exchange and exchange nonces. Those inputs establish a shared secret without sending that secret across the network. IKE then applies its key-derivation process to create the symmetric keys used for integrity protection and encryption of IKE messages and IPsec Encapsulating Security Payload traffic. In practice, separate directional keys can be derived so that each peer has the corresponding material needed to encrypt outbound traffic and authenticate or decrypt inbound traffic. The important principle is that the resulting IPsec protection relies on mutually derived symmetric session keys, rather than a single party distributing a session key. Authentication methods such as a pre-shared secret or certificates validate the peers during the exchange, while the negotiated Diffie-Hellman secret and nonces feed the session-key derivation. RFC 7296 defines the IKE keying material, including the derived keys used to protect IKE and Child SAs: [RFC 7296, Internet Key Exchange Protocol Version 2](#). Check Point also describes IKE as the protocol used to establish VPN security associations: [Check Point R81.20 Security Management Administration Guide](#).

QUESTION NO: 7

What is the buffer size set by the fw ctl zdebug command?

- A. 8GB
- B. 1 MB

C. 1 GB

D. 8 MB

ANSWER: B

Explanation:

The correct answer is **1 MB**. The `fw ctl zdebug` command is a simplified, temporary kernel-debug command intended for quickly collecting debugging output without manually configuring all kernel-debug parameters. As part of its predefined setup, it configures the kernel debug buffer to 1 MB and enables the relevant debugging behavior for the requested module or modules. This fixed buffer size is important operationally: the command is designed for short-lived, focused troubleshooting, so administrators should begin capturing output immediately and, where appropriate, redirect it to a file for later review. The command is useful when investigating traffic processing, kernel events, and module-level behavior because it reduces the amount of manual `fw ctl debug` configuration required. For more extensive traces, an administrator can use the underlying kernel-debug commands to explicitly control flags and buffer allocation, but the standard behavior of `fw ctl zdebug` itself uses a 1 MB buffer. Check Point documents the command syntax and behavior in its [R81.20 CLI Reference Guide](#). See also Check Point's [kernel debug troubleshooting guidance](#) for safe collection and handling of diagnostic output.

QUESTION NO: 8

Which kernel process is used by Content Awareness to collect the data from contexts?

A. PDP

B. cpemd

C. dlpsda

D. CMI

ANSWER: C

Explanation:

dlpsda is the correct kernel component. Its name refers to the Data Loss Prevention Download Agent, and it is used by the Content Awareness functionality to obtain the relevant traffic data identified through contexts. Content Awareness does not make inspection decisions from basic connection information alone: it requires the applicable content to be made available for deeper analysis, including recognition of configured file and data types. The Context Management Infrastructure supplies and manages the context information associated with traffic flows, while **dlpsda** acts on the relevant context to collect the data required for Content Awareness inspection. This makes **dlpsda** the component that directly matches the question's focus on collecting data *from contexts* rather than simply maintaining context metadata or applying a policy decision. Check Point documents **dlpsda** among the gateway components supporting content inspection and identifies it as the Download Agent used for Content Awareness. This role is particularly relevant during troubleshooting because process and kernel-component status can help establish whether Content Awareness has the required data path for content analysis. See Check Point's [R81.20 Content Awareness documentation](#) and the [Content Awareness Gateway configuration guide](#).

QUESTION NO: 9

You receive complains that Guest Users cannot login and use the Guest Network which is configured with Access Role of Guest Users. You need to verify the Captive Portal configuration. Where can

you find the config file?

A. on the gateway at \$NACPORTAL_HOME/conf/httpd2_nac.conf

B. on the management at SCPNAC_HOME/conf/httpd_nac.conf

C. on the management at SNACPORTAL_HOME/conf/httpd_nac.conf

D. on the gateway at \$CPNAC_HOME/conf/httpd_nac.conf

ANSWER: A

Explanation:

The Captive Portal web-server configuration is stored locally on the Security Gateway in \$NACPORTAL_HOME/conf/httpd2_nac.conf. NACPORTAL_HOME is the installation-home environment variable for the gateway's Network Access Control/Captive Portal component. This Apache-based configuration file contains the settings used to provide the portal service, including relevant listener and virtual-host behavior. Therefore, when Guest Users cannot authenticate through a Guest Network protected by an Access Role, an administrator troubleshooting the portal service should inspect this gateway-local file, along with the associated gateway logs and service status. It is important to perform such troubleshooting in Expert mode and to preserve a backup before making any manual change, because unsupported or incorrect modifications can prevent the portal from starting or functioning correctly. Check Point's Identity Awareness documentation describes Captive Portal as a gateway-based component used to identify and authenticate users before access is granted. See the [R81.20 Identity Awareness Administration Guide: Captive Portal](#) and the [R81.20 Installation and Upgrade Guide](#) for gateway and Expert-mode administration context.

QUESTION NO: 10

A gateway is expected to use SecureXL acceleration, but throughput is substantially lower than expected after a configuration change. Which command should be used first to verify whether SecureXL acceleration is enabled on the Security Gateway?

- A. fw ctl multik stat
- B. fwaccel stat
- C. fw ctl pstat
- D. cphaprob state

ANSWER: B

Explanation:

fwaccel stat is correct because it reports the current SecureXL acceleration status. Before investigating CoreXL distribution, interface throughput, or policy behavior, the administrator should confirm whether acceleration is enabled and operational on the affected gateway.

fw ctl multik stat shows CoreXL Firewall Kernel instances. fw ctl pstat reports kernel memory information, and cphaprob state reports ClusterXL member state.

QUESTION NO: 11

What is the shorthand reference for a classification object?

- A. classobj
- B. CLOB
- C. COBJ
- D. class.obj

ANSWER: C

Explanation:

COBJ is the shorthand reference used for a classification object. In Check Point terminology and internal object references, this abbreviation directly combines “classification” with “object,” making it the recognized concise identifier for this object type. Classification objects are part of the Security Management object model and allow policy components to identify and handle traffic or entities according to their assigned classification. Knowing these abbreviations is useful during troubleshooting because object references can appear in management output, policy-related diagnostics, and technical documentation in abbreviated form. The abbreviation is therefore not merely a descriptive label; it is the expected shorthand used when referring to this specific Check Point object category. For broader context on managing objects and policy components in R81.20, consult the [Check Point R81.20 Security Management Administration Guide](#) and the [Check Point certification information](#).

QUESTION NO: 12

You need to run a kernel debug over a longer period of time as the problem occurs only once or twice a week. Therefore you need to add a timestamp to the kernel debug and write the output to a file. What is the correct syntax for this?

- A. `fw ctl debug -T -f > filename debug`
- B. `fw ctl kdebug -T -f -o filename debug`
- C. `fw ell kdebug -T > filename debug`
- D. `fw ctl kdebug -T -f > filename.debug`

ANSWER: B

Explanation:

`fw ctl kdebug -T -f -o filename debug` is the correct syntax because it uses the dedicated `fw ctl kdebug` command and combines the required kernel-debug controls. The `-T` switch adds timestamps to debug output, which is essential when correlating infrequent events with logs, traffic, or other system activity. The `-f` switch keeps the command running and continuously follows kernel-debug output, allowing capture to remain active until the intermittent issue occurs. The `-o filename` parameter directs the collected output to the specified file rather than relying on an interactive terminal session or shell redirection. The trailing `debug` value represents the debug flags argument supplied to the command. This approach is appropriate for extended troubleshooting because the capture is timestamped, persistent, and written directly to a chosen output file. Check Point documents the available `fw ctl kdebug` parameters, including timestamping, follow mode, output-file selection, and debug flags, in its CLI Reference Guide: [fw ctl kdebug command reference](#). Additional kernel-debug troubleshooting context is available in the [Check Point Support Center](#).