

DUMPS ARENA

Palo Alto Networks XDR Analyst

Palo Alto Networks XDR-Analyst

Version Demo

Total Demo Questions: 11

Total Premium Questions: 115

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cortex XDR Fundamentals	19
Topic 2, Cortex XDR Data Collection	13
Topic 3, Cortex XDR Detection	26
Topic 4, Cortex XDR Investigation	19
Topic 5, Cortex XDR Response	28
Topic 6, Cortex XDR Threat Hunting	6
Topic 7, Mix Questions	4
Total	115

QUESTION NO: 1

An analyst reviews a suspicious executable that was submitted for cloud analysis. Which two statements correctly describe the WildFire role in a Cortex XDR investigation? (Choose two.)

- A. WildFire accepts and analyzes submitted samples.
- B. WildFire provides a verdict for the analyzed sample.
- C. WildFire runs locally on the Cortex XDR agent to detect behavioral threats.
- D. WildFire aggregates related alerts into Cortex XDR incidents.
- E. WildFire replaces Behavioral Threat Protection on endpoints.

ANSWER: A B

Explanation:

WildFire accepts and analyzes submitted samples and **WildFire provides a verdict for the analyzed sample** are correct. The verdict and report can enrich Cortex XDR investigation evidence and help an analyst determine whether the artifact is malicious.

WildFire is not a local Cortex XDR agent engine and does not replace endpoint Behavioral Threat Protection. It analyzes submitted samples rather than directly correlating all endpoint alerts into incidents.

QUESTION NO: 2

A Windows Malware profile is configured to respond automatically when Cortex XDR identifies a malicious causality chain. Which two actions can the profile perform to contain the activity? (Choose two.)

- A. Automatically kill the processes involved in malicious activity.
- B. Automatically block the IP addresses involved in malicious traffic.
- C. Automatically close the connections involved in malicious traffic.
- D. Automatically terminate the threads involved in malicious activity.

ANSWER: A B

Explanation:

Automatically killing processes involved in malicious activity and **automatically blocking IP addresses involved in malicious traffic** are correct. These actions are available through the response to malicious causality chains setting and help contain both execution on the endpoint and associated malicious network communication.

Closing individual connections and terminating individual threads are not the configured causality-chain response actions. The response is based on the identified malicious process chain and associated IP addresses.

QUESTION NO: 3

What is by far the most common tactic used by ransomware to shut down a victim's operation?

- A. preventing the victim from being able to access APIs to cripple infrastructure
- B. denying traffic out of the victims network until payment is received
- C. restricting access to administrative accounts to the victim
- D. encrypting certain files to prevent access by the victim

ANSWER: D

Explanation:

Encrypting certain files to prevent access by the victim is the most common operational-disruption tactic associated with ransomware. In a typical ransomware incident, attackers deploy malware that encrypts business-critical data, such as documents, databases, application files, and shared network storage. The victim can often still power on affected systems, but essential data and services become unavailable because the organization cannot read or use the encrypted content. This loss of availability can halt normal operations, including order processing, patient care, manufacturing, financial transactions, and communications, until systems are restored from reliable backups or files are decrypted. Attackers then present a ransom demand, commonly requesting cryptocurrency in exchange for a purported decryption key. Modern ransomware operations may also steal data and threaten its release, but encryption-based denial of access remains the central and most widely recognized mechanism for disrupting a victim's operations. Palo Alto Networks describes ransomware as malware that encrypts a victim's files and demands payment to restore access. The Cybersecurity and Infrastructure Security Agency likewise identifies encrypting data and systems as a defining ransomware impact. See [Palo Alto Networks: What Is Ransomware?](#) and [CISA: Ransomware Guide](#).

QUESTION NO: 4

A WildFire report confirms that a file is malicious. The file remains on one affected endpoint, and the security team must both preserve the local sample for investigation and prevent the exact known file from executing on other protected endpoints. Which two actions should the analyst take? (Choose two.)

- A. Quarantine the file on the affected endpoint
- B. Add the file SHA-256 hash to the Cortex XDR block list
- C. Create an alert exclusion for the WildFire verdict
- D. Change the incident status to resolved

ANSWER: A B

Explanation:

Quarantine the file on the affected endpoint and **add the file SHA-256 hash to the Cortex XDR block list** are correct. Quarantine preserves the local artifact in a protected folder while preventing execution. A block-list entry applies a prevention decision for the exact known malicious file hash across protected endpoints.

An alert exclusion only suppresses matching alerts and does not prevent execution. Changing the incident status is an administrative workflow action and does not remediate or prevent the malicious file.

QUESTION NO: 5

Where can SHA256 hash values be used in Cortex XDR Malware Protection Profiles?

- A. in the macOS Malware Protection Profile to indicate allowed signers
- B. in the Linux Malware Protection Profile to indicate allowed Java libraries
- C. SHA256 hashes cannot be used in Cortex XDR Malware Protection Profiles
- D. in the Windows Malware Protection Profile to indicate allowed executables

ANSWER: D

Explanation:

SHA256 hash values can be used in the Windows Malware Protection Profile to identify allowed executables. This capability provides a precise file-based allowlist: Cortex XDR compares an executable's SHA256 value with the configured value and

can exclude a known, trusted executable from malware scanning or prevention processing as configured in the profile. Hash-based identification is especially useful when an organization must accommodate a legitimate internally developed application, specialized business software, or another trusted executable that might otherwise generate an unwanted detection. Because a SHA256 hash identifies a specific file version, administrators should validate the file's provenance before allowlisting it and update the entry if the approved executable is modified or replaced. Palo Alto Networks documents SHA256 allowlisting as part of the Windows-specific malware protection profile configuration. This should be managed narrowly and reviewed regularly so that the exception remains limited to the intended trusted executable. See the Palo Alto Networks documentation for [configuring a Windows Malware Protection Profile](#) and the broader [Malware Protection Profiles](#) reference.

QUESTION NO: 6

An analyst is preparing a Technical Assistance Center case for an endpoint security event that requires further investigation. Which two artifacts should be collected from Cortex XDR before opening the case? (Choose two.)

- A. The agent technical support file.
- B. The prevention archive from the alert.
- C. The unique agent ID.
- D. The distribution ID of the agent.
- E. A list of all exceptions currently applied to the agent.

ANSWER: A B

Explanation:

The agent technical support file and **the prevention archive from the alert** are correct. The technical support file provides agent diagnostics and logs, while the prevention archive provides event-specific evidence associated with the prevention alert.

An agent ID and distribution ID can help identify an endpoint, but they do not provide the diagnostic and forensic information required to investigate the security event. A list of exceptions may be useful context, but it is not a substitute for the required support artifacts.

QUESTION NO: 7

When creating a scheduled report which is not an option?

- A. Run weekly on a certain day and time.
- B. Run quarterly on a certain day and time.
- C. Run monthly on a certain day and time.
- D. Run daily at a certain time (selectable hours and minutes).

ANSWER: B

Explanation:

Run quarterly on a certain day and time. is the correct selection because Cortex XDR scheduled-report recurrence supports daily, weekly, and monthly schedules, rather than a quarterly recurrence. When defining a report schedule, an analyst can configure the recurrence parameters relevant to the selected supported interval, including the execution time and, for weekly or monthly scheduling, the applicable day. The scheduling workflow also includes operational settings such as the report time zone, start and end dates, and report recipients. This allows teams to automate regular delivery of report output at useful operational cadences without having to manually run the report each time. A quarterly schedule is not offered as a native recurrence choice in the Cortex XDR scheduled-report configuration. To produce a quarterly deliverable, teams can use an available schedule as appropriate and manage the quarterly reporting process separately. Palo Alto Networks

documents the available reporting and scheduling workflow in its Cortex XDR administrator documentation: [Run or Schedule Reports](#). Additional Cortex XDR reporting documentation is available in the [Reporting](#) section.

QUESTION NO: 8

When reaching out to TAC for additional technical support related to a Security Event; what are two critical pieces of information you need to collect from the Agent? (Choose Two)

- A. The agent technical support file.
- B. The prevention archive from the alert.
- C. The distribution id of the agent.
- D. A list of all the current exceptions applied to the agent.
- E. The unique agent id.

ANSWER: A B

Explanation:

The agent technical support file and the prevention archive from the alert are the key artifacts to collect before engaging Palo Alto Networks Technical Assistance Center for a security-event investigation. The agent technical support file consolidates endpoint diagnostics needed to assess the Cortex XDR agent's operational state, configuration, component health, and relevant logs. This lets TAC correlate agent-side behavior with the reported event and identify issues that may not be visible in the console alone. The prevention archive from the alert preserves event-specific forensic material generated by the prevention action. It provides the investigation context TAC needs, such as the data associated with the alert and the activity leading to the prevention verdict. Together, these artifacts give support both the endpoint diagnostic evidence and the alert-specific evidence needed to efficiently analyze the issue, reproduce relevant conditions where possible, and determine the appropriate remediation path. Cortex XDR administrators can use the Cortex XDR documentation to manage agents and investigate alerts, while the Palo Alto Networks support portal provides the case-management process for submitting diagnostic material to TAC. See the [Cortex XDR documentation](#) and the [Palo Alto Networks Customer Support Portal](#).

QUESTION NO: 9

When using the "File Search and Destroy" feature, which of the following search hash type is supported?

- A. SHA256 hash of the file
- B. AES256 hash of the file
- C. MD5 hash of the file
- D. SHA1 hash of the file

ANSWER: A

Explanation:

SHA256 hash of the file is the supported hash type for Cortex XDR File Search and Destroy searches. This response capability lets an analyst locate a specific file across managed endpoints and, when appropriate, take remediation action to remove it. A SHA-256 value identifies the exact file content, rather than relying only on a filename or path, which may be shared by multiple unrelated files. Using a content-derived identifier is important during incident response because it enables a precise search for the known malicious artifact associated with an alert, investigation, or threat-intelligence indicator. SHA-256 is also the standard file identifier broadly used throughout Cortex XDR investigations and endpoint telemetry, helping analysts pivot consistently from a detected artifact to endpoint-wide response actions. In the File Search and Destroy workflow, the SHA-256 value must correspond to the target file so Cortex XDR can identify matching instances on endpoints and apply the selected action accurately. Palo Alto Networks documents File Search and Destroy as an endpoint response

feature in the Cortex XDR administration documentation; see the [File Search and Destroy documentation](#) and the [Cortex XDR response documentation](#).

QUESTION NO: 10

Several alerts from one endpoint involve a suspicious document, a spawned command interpreter, and outbound connections to a suspicious address. Cortex XDR presents the alerts together as one incident and identifies the most relevant artifacts. Which engine performs this correlation?

- A. Sensor Engine
- B. Causality Analysis Engine
- C. Log Stitching Engine
- D. Causality Chain Engine

ANSWER: B

Explanation:

Causality Analysis Engine is correct because it analyzes the relationships among alerts and artifacts, identifies relevant entities, and aggregates related alerts into an incident. This gives the analyst a connected investigation context rather than a set of unrelated detections.

A sensor collects endpoint telemetry, but it does not perform incident-level correlation. The other options do not represent Cortex XDR engines responsible for incident aggregation.

QUESTION NO: 11

What contains a logical schema in an XQL query?

- A. Bin
- B. Array expand
- C. Field
- D. Dataset

ANSWER: D

Explanation:

Dataset is correct because XQL queries operate on datasets, which are logical collections of records made available for investigation and analytics in Cortex XDR. A dataset defines the logical schema that a query can address: its schema establishes the available fields, their names, and the data structure that XQL stages use when filtering, joining, aggregating, or projecting results. For example, beginning a query with a dataset selection establishes the data source and the fields that can subsequently be referenced in expressions. This distinction is important when building XQL queries because field names are evaluated in the context of the selected dataset's schema. Cortex XDR provides predefined datasets for different telemetry and investigation use cases, and the XQL language uses dataset-oriented query construction to retrieve and process that data. Palo Alto Networks' XQL language documentation describes the dataset stage and its role in selecting the query data source, while the Cortex XDR documentation provides the broader investigation context for using XQL queries. See the [Palo Alto Networks Dataset Stage reference](#) and the [Cortex XDR XQL Query Builder documentation](#).