

VMware Carbon Black XDR User Guide

16 May 2024

VMware Carbon Black Cloud

You can find the most up-to-date technical documentation on the VMware by Broadcom website at:

<https://docs.vmware.com/>

VMware by Broadcom

3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2023-2024 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to <https://www.broadcom.com>. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies. [Copyright and trademark information](#).

Contents

- 1 Carbon Black XDR User Guide** 5
- 2 Carbon Black XDR Overview** 6
 - What are the Benefits of XDR? 7
 - What are the Use Cases of XDR? 8
 - What is the difference between XDR and EDR? 8
- 3 XDR Data and Methodologies** 9
- 4 XDR Search Fields** 14
- 5 Disable XDR by Policy** 15
- 6 Retrieving XDR Data** 16
 - Retrieve XDR Data on the Processes or Observations Page 16
 - Retrieve XDR Data on the Alerts Page 17
- 7 Exploring XDR Data** 19
 - Exploring XDR Data on the Process Analysis Page 19
 - Exploring XDR Data on the Observations Page 20
 - Exploring XDR Data on the Alerts Page 22
 - Exploring XDR Data on the Alert Triage Page 24
- 8 Investigate** 25
 - About Digital Signatures 27
 - Investigate - Processes 29
 - Process Analysis 30
 - Investigate - Observations 35
 - Observations Overview 36
 - Searching for Observations 36
 - Observation Types 38
 - Histogram 39
 - Group By and View By 40
 - Observations Search Results 42
 - Investigate - Auth Events 44
 - Enable Auth Event Collection 45
 - Viewing Auth Events in the Console 45
 - Windows Authentication Events 46

- Authentication Event Details 48
- Investigating Script-Based Attacks 49
 - Investigate Obfuscated PowerShell Scripts 50
- Obtain a Process Hash 51
- Add an Investigate Query to a Watchlist Report 52

9 Alerts 54

- View Alert Details 54
 - Alert Types 56
 - View Specific Alert Types 57
 - Alert and Report Severity 58
 - Alert ID, Event ID, and Threat ID 59
 - Anomaly Classification 59
- Group Alerts 60
 - Group By: Threat ID 61
- View Alerts in a Threat ID 61
 - View Threat ID Details 62
 - Add Notes 63
- Editing the Alert Workflow 63
 - Close Alerts 63
 - Open Alerts 64
 - Mark Alert as In Progress 65
- Search Basics 65
- Alert Triage 67
 - Investigating Alerts 67
 - Take Action on Alerts 68
 - Quarantining an Asset 69
 - Add Determination for Alerts 70
 - True and False Positives 70
 - Visualizing Alerts 71
 - Alert Origin, Behaviors, and TTPs 74
- Script Host Replacement Occurrence 75

Carbon Black XDR User Guide

1

Carbon Black Extended Detection and Response (XDR) greatly enhances lateral security by leveraging telemetry. Security teams can leverage Carbon Black XDR to quickly identify threats across their environment and make better-informed decisions in applying prevention policies.

Carbon Black XDR unifies endpoint and workload security capabilities with critical visibility into the network—reducing blind spots and detecting threats faster.

The Future-Ready SOC

Using XDR to achieve unified visibility and control

Carbon Black



Carbon Black XDR Overview

2

Carbon Black XDR is a consolidation of tools and data that provides extended visibility, analysis, and response across endpoints, workloads, users, and networks.

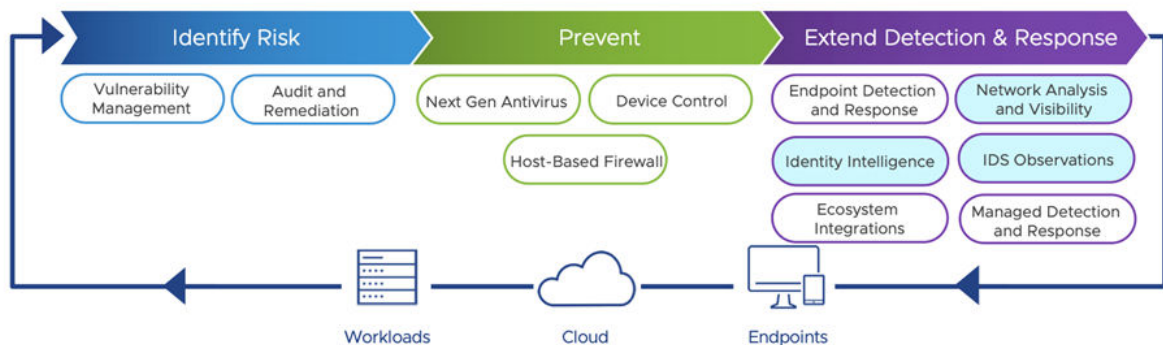
Carbon Black XDR focuses on adding network telemetry for XDR, and provides insight into network packets and processes.

For more information about XDR, see [What is Extended Detection and Response \(XDR\)?](#)

Carbon Black XDR

One agent. One console. One platform.

Go beyond the endpoint to see more and stop more



Carbon Black XDR implements XDR for Carbon Black Cloud Enterprise EDR. This implementation requires the Carbon Black Cloud Windows Sensor 3.9.1 MR1+.

With Carbon Black XDR, you can visualize and analyze relevant network data. For example:

- Signatures of network connections (JA3 and JA3S thumbprints)
- Network intrusion detection
- Security wrapper details (TLS data)
- Signer of certificate (encryption - TLS data)
- HTTP details

Read the following topics next:

- [What are the Benefits of XDR?](#)

- [What are the Use Cases of XDR?](#)
- [What is the difference between XDR and EDR?](#)

What are the Benefits of XDR?

XDR's capabilities above and beyond EDR give it several tangible benefits for securing an organization's IT environment. These benefits include the following.

Greater Visibility and Context

XDR provides a full, 360-degree view of the security environment. It allows security analysts to see threats—even those that leverages legitimate software, ports and protocols to gain entry—on any security layer. It collects the how of an attack, the blueprint, the entry point, who else is affected, where the threat originated, and how it spread. This additional context, as well as the analytics required to make sense of it, is crucial to a speedy response to threats.

Prioritization

IT and security teams often struggle to keep up with endless alerts generated by their security services. XDR's data analysis and correlation capabilities allow it to group related alerts, prioritize them, and surface only the most important ones.

Automation

XDR's use of automation speeds up detection and response and removes manual steps from security processes. This allows IT teams to handle a large volume of security data and carry out complex processes in a repeatable way.

Operational Efficiency

Instead of a fragmented collection of security tools, XDR provides a holistic view of threats throughout the entire managed fleet. It offers centralized data collection and response that is tightly integrated into the environment and broader security ecosystem.

Faster Detection and Response

These advantages add up to an effective security posture. XDR's added efficiency allows it to detect and respond to threats faster.

Sophisticated Responses

XDR's sophisticated capabilities and greater visibility allow it to tailor the response to the specific system and leverage other control points to minimize the overall impact.

What are the Use Cases of XDR?

There are many ways in which you can use XDR to detect and respond to endpoint threats. The following are the most common use cases.

Threat Hunting

Although it is likely that threats already exist in any given network, many security teams struggle to find the time to do proactive threat hunting. XDR's telemetry and automation capabilities allow much of this work to be done automatically, significantly lightening the load on security teams and allowing them to carry out threat hunting alongside their other tasks.

Triage

One of a security team's most important functions is to prioritize or triage alerts and quickly respond to the most crucial ones. XDR helps sift through the noise by using powerful analytics to correlate thousands of alerts into a small number of high-priority ones.

Investigation

XDR's extensive data collection, superior visibility, and automated analysis allow security teams to quickly and easily establish where a threat originated, how it spread, and what other users or devices might be affected. This is crucial to both removing the threat and hardening the network against future threats.

What is the difference between XDR and EDR?

XDR extends the capabilities of EDR across all the security layers in the environment—workloads, devices, users, and networks.

Rather than the single point of view that EDR provides, XDR enables telemetry and behavioral analysis across multiple security layers, allowing security teams to see the big picture.

Bad actors don't limit their attacks to a single security layer; thus, security teams cannot afford to limit their view to one layer. EDR gives security professionals visibility into endpoints that might be compromised—but this is not enough when an attack has moved across the network and into other systems before the security team has become aware of it.

This is where XDR comes in. By providing a holistic view of activity across the system that avoids visibility gaps, XDR allows security teams to understand where a threat comes from and how it is spreading across the environment. XDR offers greater analysis and correlation capabilities and a holistic point of view.

XDR Data and Methodologies

3

This topic describes netconn data that you can retrieve in the Carbon Black Cloud Console. It also introduces Intrusion Detection System (IDS) to identify and classify netconn traffic, and Network Traffic Analysis (NTA) to determine anomalies.

Netconn Data

XDR analyzes network connections (netconn) and makes these analyses visible for threat hunting and investigations. For a list of netconn fields that you can search on, see [Chapter 4 XDR Search Fields](#).

Observation data *types* are available for filtering and sorting. For descriptions of observation types, see [Investigate - Observations](#).

- CB Analytics
- Contextual Activity
- TAU Intelligence
- Tamper
- Blocked Hash
- Intrusion Detection System
- Network Traffic Analysis
- Host-based Firewall
- Indicator of Attack

Intrusion Detection System (IDS)

The MicroIDS network classification engine runs on netconn events and uses IP packets to classify connections. This engine is embedded in the Carbon Black Cloud 3.9+ Windows sensor.

IDS provides the following benefits:

- Monitors inbound and outbound network traffic
- Monitors data that moves between the system and the network

- Detects attacks by capturing and analyzing network packets
- Identifies abnormal network traffic

Intrusion Detection System is a filter option on the Alerts, Processes, and Observations pages.

Network Traffic Analysis (NTA)

Network Traffic Analysis (NTA) monitors network availability and activity to identify anomalies. Benefits of NTA include:

- Improved visibility into your network by providing rich context
- Detecting anomalous actions
- Assisting in security triage investigations

Data extracted from network packets can help you track usage and monitor for suspicious behavior.

For example:

TIME ▼	TYPE/REASON ▼
2:19:12 pm Jul 19, 2023	Network Traffic Analysis DO-NOT-UPGRADE-3DOT9-1 received a connection on an unusual port (445) from ::1.
12:32:17 pm Jul 18, 2023	Network Traffic Analysis DO-NOT-UPGRADE-3DOT9-1 received a connection on an unusual port (5357) from [REDACTED]
12:32:16 pm Jul 18, 2023	Network Traffic Analysis DO-NOT-UPGRADE-3DOT9-1 received a connection on an unusual port (5357) from [REDACTED]
12:32:16 pm Jul 18, 2023	Network Traffic Analysis DO-NOT-UPGRADE-3DOT9-1 received a connection on an unusual port (5357) from [REDACTED]
3:30:47 pm Jul 17, 2023	Network Traffic Analysis CBSExcbpserv received a connection on an unusual port (47001) from ::1.

Carbon Black shows the following types of NTA detectors:

Alert Type	Description
IP Profiler	Looks for connections to or from a local host that have an unusual remote IP address, compared to remote IP addresses to which the host typically connects.
User Agent Profiler	Looks for unusual HTTP user agents in connections that are made from a local device, compared to previous device activity.

Alert Type	Description
Port Profiler	Looks for unusual port activity on either end of a network connection. This alert type includes four variations: <ul style="list-style-type: none"> ■ External Port Profiler ■ Internal Port Profiler ■ External Server Port Profiler ■ Internal Server Port Profiler
JA3 Profiler	Looks for unusual JA3 hashes of SSL/TLS connections that are made from a local device, compared to previous device activity.

Network Traffic Analysis is a filter option on the Alerts, Processes, and Observations pages.

MITRE ATT&CK Tactics & Techniques

MITRE ATT&CK® is visible throughout the Carbon Black Cloud console as well as in XDR; it is not XDR-specific.

MITRE ATT&CK is widely adopted by the security industry as a knowledge base of adversarial tactics and techniques. For more information about MITRE ATT&CK, see [MITRE ATT&CK](#) (external link).

A *Tactic* represents the purpose behind a technique. For example, a MITRE ATT&CK Tactic ID of `TA0001` signifies that the adversary is trying to penetrate your network.

A *Technique* expresses how the adversary is attacking. For example, a MITRE ATT&CK Technique ID of `T1548.003` indicates that adversaries are performing sudo caching to elevate privileges.

You can filter, search, and sort on `Tactic` and `Technique` fields on the Observations, Alerts, Processes, and Process Analysis pages. For example:

INVESTIGATE 3 days All results

Observations Processes Auth Events

FILTERS Clear


- Type (5)
 - Contextual Activity 90.9%
 - Indicator of Attack 7.6%
 - CB Analytics 0.7%
 - Tamper 0.5%
 - Intrusion Detection System 0.1%
- Event Type (4)
- Process (50+)
- Effective Reputation (5)
- Process Hash (30+)
- Container
- Container Image
- Device (1)
- Username (1)
- Parent Effective Reputation (4)
- TTP (32)
- Location (50+)
- Application Protocol (1)
 - HTTP 0.1%
- ATT&CK Tactic (2)
 - T1029 - Remote Desktop Protocol 2.9%
 - T1003 - LSASS Memory 0.6%
- ATT&CK Technique (2)
 - T1003.001 - LSASS Memory 7.9%
 - T1003.002 - Remote Desktop 0.1%

Showing max 10,000 results (20,649 total)

TIME	TYPE/REASON	USERNAME	DEVICE	ATT&CK TACTIC	ATTACK TECHNIQUE	ACTIONS
6:50:41 pm May 27, 2023	Intrusion Detection System The application vmtoolsd.exe running on DEV01-38x-1 sent RDP data from 10.100.100.100 to 10.100.100.100 which matched the RDP protocol signature	NT AUTHORITY\NETWORK SERVICE	vmtoolsd-38x-1	TA0008 Lateral Movement	T1021.001 Remote Desktop Protocol	
2:08:02 pm May 29, 2023	Indicator of Attack The application vmtoolsd.exe requested the content of basuss.exe. A Deny policy action was applied.	NT AUTHORITY\NETWORK SERVICE	vmtoolsd-38x-1	TA0006 Credential Access	T1003.001 LSASS Memory	
10:33:22 am May 27, 2023	Indicator of Attack The application vmtoolsd.exe requested the content of basuss.exe. A Deny policy action was applied.	NT AUTHORITY\NETWORK SERVICE	vmtoolsd-38x-1	TA0006 Credential Access	T1003.001 LSASS Memory	
10:19:42 am May 29, 2023	Indicator of Attack The application vmtoolsd.exe requested the content of basuss.exe. A Deny policy action was applied.	NT AUTHORITY\NETWORK SERVICE	vmtoolsd-38x-1	TA0006 Credential Access	T1003.001 LSASS Memory	
8:08:11 pm May 29, 2023	Indicator of Attack The application vmtoolsd.exe requested the content of basuss.exe. A Deny policy action was applied.	NT AUTHORITY\NETWORK SERVICE	vmtoolsd-38x-1	TA0006 Credential Access	T1003.001 LSASS Memory	
2:08:08 pm May 29, 2023	Indicator of Attack The application vmtoolsd.exe requested the content of basuss.exe. A Deny policy action was applied.	NT AUTHORITY\NETWORK SERVICE	vmtoolsd-38x-1	TA0006 Credential Access	T1003.001 LSASS Memory	
2:08:02 pm May 27, 2023	Indicator of Attack The application vmtoolsd.exe requested the content of basuss.exe. A Deny policy action was applied.	NT AUTHORITY\NETWORK SERVICE	vmtoolsd-38x-1	TA0006 Credential Access	T1003.001 LSASS Memory	
10:33:11 am May 28, 2023	Indicator of Attack The application vmtoolsd.exe requested the content of basuss.exe. A Deny policy action was applied.	NT AUTHORITY\NETWORK SERVICE	vmtoolsd-38x-1	TA0006 Credential Access	T1003.001 LSASS Memory	
2:08:11 am May 30, 2023	Indicator of Attack The application vmtoolsd.exe requested the content of basuss.exe. A Deny policy action was applied.	NT AUTHORITY\NETWORK SERVICE	vmtoolsd-38x-1	TA0006 Credential Access	T1003.001 LSASS Memory	
10:33:05 pm May 27, 2023	Indicator of Attack The application vmtoolsd.exe requested the content of basuss.exe. A Deny policy action was applied.	NT AUTHORITY\NETWORK SERVICE	vmtoolsd-38x-1	TA0006 Credential Access	T1003.001 LSASS Memory	

Tip

- To show both the **Tactic** and **Technique** columns in the table on the Observations and Alerts pages, click **Configure Table** at the bottom left of the table.



- Click the **Information** icon to view a brief synopsis of the tactic or technique. Within the **Information** pane, you can click **Learn more** to go to the MITRE web site for more information.
- MITRE ATT&CK tactics and techniques display in the **Observation Details** pane on the Observations, Alerts, and Alert Triage pages.

OBSERVATION DETAILS

Type Intrusion Detection System (IDS) ⓘ

Reason TCP traffic from asset [redacted] matched IDS signatures for threat Kerberos Authentication Failure. SVCHOST.EXE made a connection from 10.30.5.2 to 10.101.1.252

Time 4:59:12 pm Nov 21, 2022

Threat Kerberos Authentication Failure ⓘ

Rule IDS:05628395 ⓘ

MITRE ATT&CK

Tactic TA006 - Credential Access ⓘ

Technique T1110 - Brute Force ⓘ

XDR Search Fields

4

The following list shows search fields that you can use to locate XDR-enhanced netconn events. See the in-product *Search Guide* for a full list, descriptions, and examples of all search fields.

netconn_actions	netconn_application_protocol	netconn_bytes_received
netconn_bytes_sent	netconn_community_id	netconn_domain
netconn_first_packet_timestamp	netconn_ja3_local_fingerprint	netconn_ja3_local_fingerprint_fields
netconn_ja3_remote_fingerprint	netconn_ja3_remote_fingerprint_fields	netconn_last_packet_timestamp
netconn_remote_device_id	netconn_remote_device_name	netconn_request_headers
netconn_request_method	netconn_request_uri	netconn_response_headers
netconn_response_status_code	netconn_server_name_indication	netconn_tls_certificate_issuer_name
netconn_tls_certificate_subject_name	netconn_tls_certificate_not_valid_after	netconn_tls_certificate_not_valid_before
netconn_tls_version		

Disable XDR by Policy

5

XDR network data collection is enabled by default. You can disable XDR network data collection for the sensors to which a policy is assigned. Disabling data collection does not disable Carbon Black XDR; it simply stops the sensor from collecting XDR network data and thus reduces noise.

Procedure

- 1 On the left navigation pane, click **Enforce > Policies**.
- 2 Select the policy.
- 3 Click the **Sensor** tab.
- 4 Deselect the check box for the **Enable XDR network data collection** setting.
- 5 Click **Save**.

Retrieving XDR Data

6

There are several ways you can retrieve XDR data in the Carbon Black Cloud console.

You can retrieve XDR data on the Processes, Observations, and Alerts pages.

You can explore the retrieved data on these pages and on the Process Analysis and Alerts Triage pages.

Read the following topics next:

- [Retrieve XDR Data on the Processes or Observations Page](#)
- [Retrieve XDR Data on the Alerts Page](#)


Retrieve XDR Data on the Processes or Observations Page

The following is an example of how to retrieve some kinds of XDR data on the Processes or Observations page.

Procedure

- 1 On the left navigation pane, click **Investigate**.
- 2 On the Investigate page, click **Processes** or **Observations**.
- 3 In the **Filters** pane on the left, scroll to **Application Protocol**. You can filter by the following protocols:
 - HTTP
 - TLS
 - RDP
 - DNS
 - SMB
 - LDAP
 - Kerberos

Tip Click the vertical 3-dot **Configuration** menu to configure the filters that display in the Console. For example:



- 4 Construct and run your search query. For example, search for `netconn_domain:go.microsoft.com`.

Note See netconn-specific XDR search fields in [Chapter 4 XDR Search Fields](#). See all search fields in the in-product *Search Guide*.

What to do next

See [Chapter 7 Exploring XDR Data](#) for ways to view and investigate your search results.


Retrieve XDR Data on the Alerts Page

The following is an example of how to retrieve XDR data on the Alerts page.

Procedure

- 1 On the left navigation pane, click **Alerts**.
- 2 In the **Filters** pane on the left, scroll to **ATT&CK Tactic** and select `TA0002`.

Tip Click the vertical 3-dot **Configuration** menu to configure the filters that display in the console. For example:



- 3 Construct and run your search query. For example, filter for `Intrusion Detection System`. Alternatively, you could search for `type:INTRUSION_DETECTION_SYSTEM` in the Search bar.

Note See netconn-specific XDR search fields in [Chapter 4 XDR Search Fields](#). See all search fields in the in-product *Search Guide*.

What to do next

See [Chapter 7 Exploring XDR Data](#) for ways to view and investigate your search results.

Exploring XDR Data

7

You can explore XDR and netconn data in the Carbon Black Cloud Console.

You can view and investigate XDR and netconn data in various ways. For example:

- The Process Analysis page displays additional information about certain netconns (protocol, timestamps, and headers).
- The Alert Triage page includes network nodes that highlight IDS-specific netconns.
- MITRE ATT&CK Tactics and Techniques are available on the Alert, Alerts Triage, Observations, and Process Analysis pages. You can use these fields to filter and search on the Processes page as well.

Note MITRE ATT&CK is not specific to XDR. Any Carbon Black Cloud instance will display this information.

- You can use the **Configure Table** option to build process-centric and network-centric views.
- An `Application Protocol` filter is available on the Alert, Observations, and Processes pages.
- You can build Watchlists from reported netconn data.

See the following topics for more details.


Read the following topics next:

- [Exploring XDR Data on the Process Analysis Page](#)
- [Exploring XDR Data on the Observations Page](#)
- [Exploring XDR Data on the Alerts Page](#)
- [Exploring XDR Data on the Alert Triage Page](#)

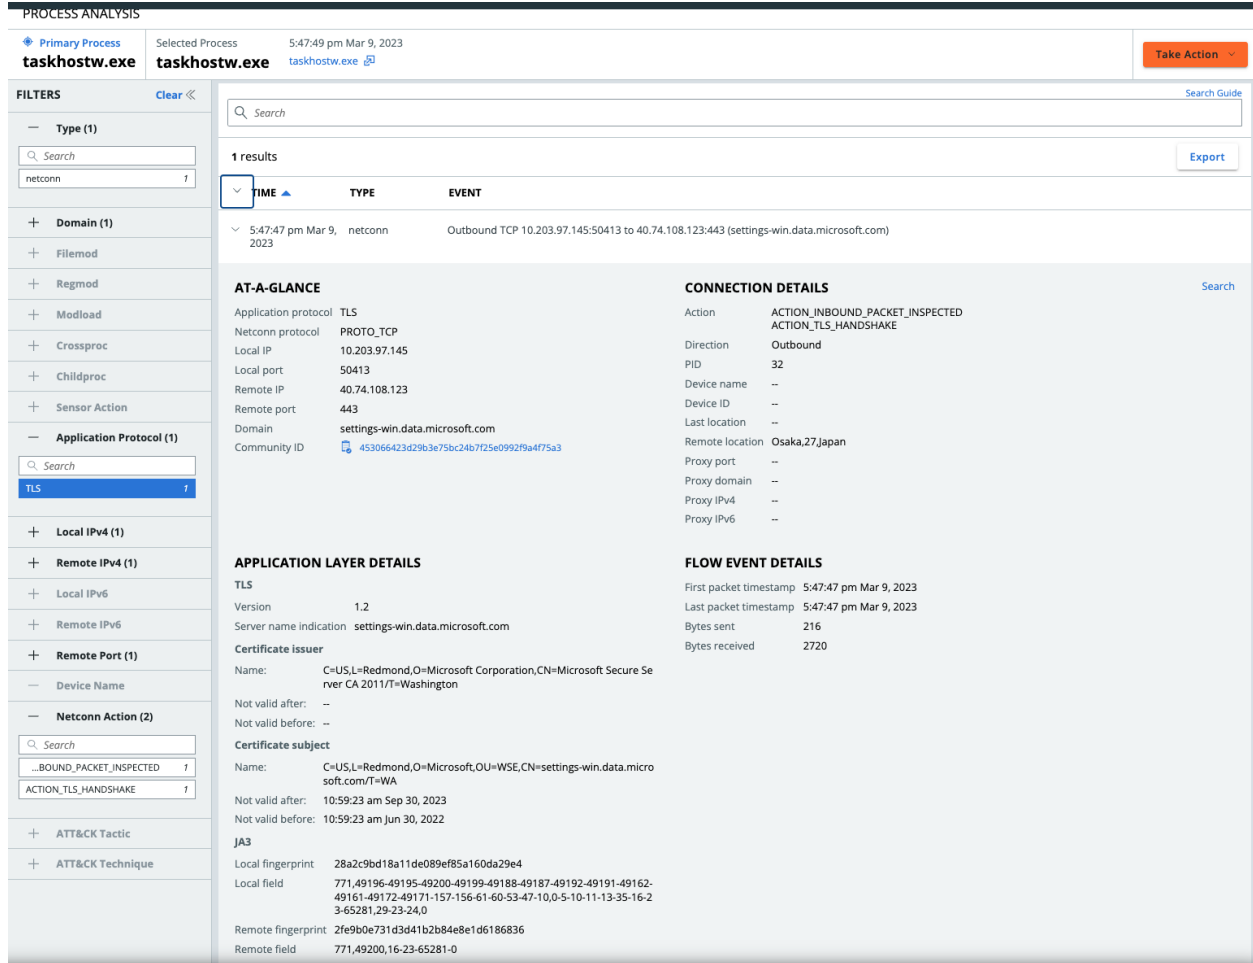
Exploring XDR Data on the Process Analysis Page

After you retrieve an event, you can explore the data on the Process Analysis page.

Note For instructions on retrieving XDR activities see [Retrieve XDR Data on the Processes or Observations Page](#).

On the left navigation pane, click **Investigate** and click the **Processes** or **Observations** tab. Then click the **Process Analysis**  icon for the item of interest.

You can now retrieve raw netconn details. For example, protocol, timestamp, and application header data displays for XDR-enabled systems.



The screenshot displays the 'PROCESS ANALYSIS' interface for the process 'taskhostw.exe' on '5:47:49 pm Mar 9, 2023'. The left sidebar shows filters for 'Type (1)' with 'netconn' selected, and 'Application Protocol (1)' with 'TLS' selected. The main area shows '1 results' for a netconn event at '5:47:47 pm Mar 9, 2023' with the description 'Outbound TCP 10.203.97.145:50413 to 40.74.108.123:443 (settings-win.data.microsoft.com)'. The details are organized into four sections:

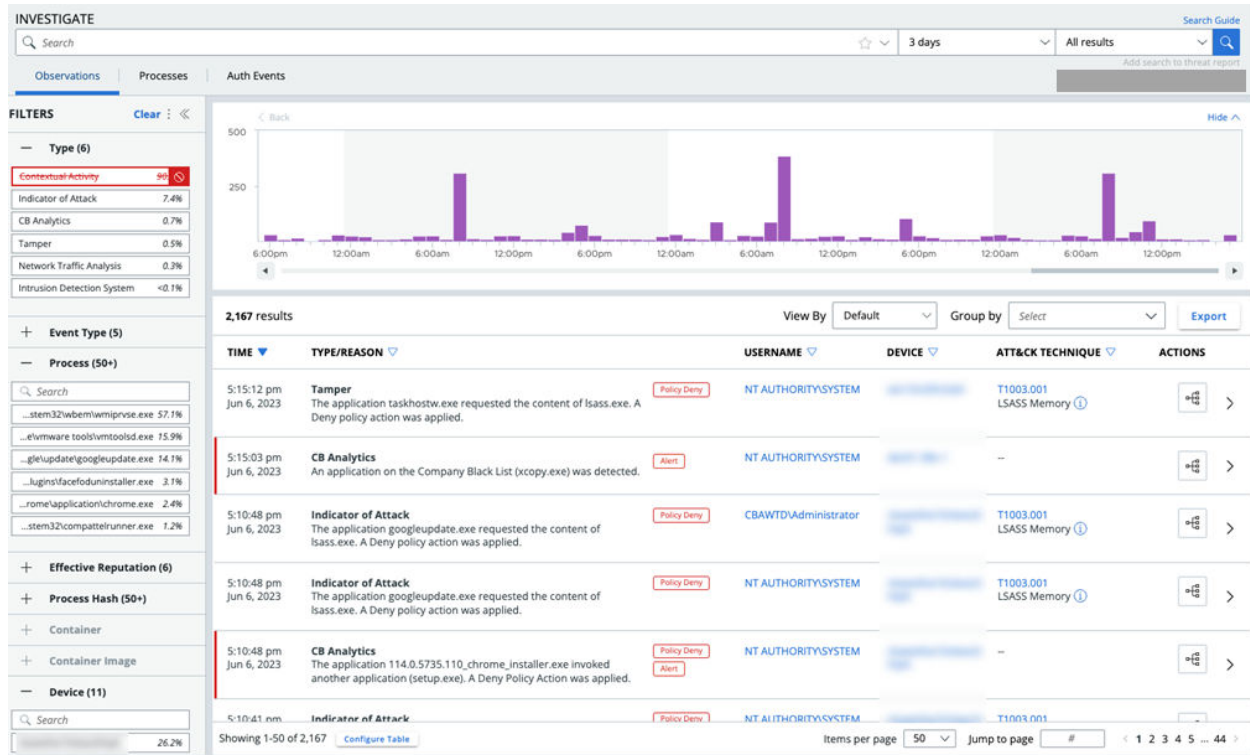
- AT-A-GLANCE:** Application protocol: TLS; Netconn protocol: PROTO_TCP; Local IP: 10.203.97.145; Local port: 50413; Remote IP: 40.74.108.123; Remote port: 443; Domain: settings-win.data.microsoft.com; Community ID: 453066423d29b3e75bc24b7f25e0992f9a47f5a3
- CONNECTION DETAILS:** Action: ACTION_INBOUND_PACKET_INSPECTED, ACTION_TLS_HANDSHAKE; Direction: Outbound; PID: 32; Device name: --; Device ID: --; Last location: --; Remote location: Osaka,27,japan; Proxy port: --; Proxy domain: --; Proxy IPv4: --; Proxy IPv6: --
- APPLICATION LAYER DETAILS:** TLS; Version: 1.2; Server name indication: settings-win.data.microsoft.com; Certificate issuer: Name: C=US,L=Redmond,O=Microsoft Corporation,CN=Microsoft Secure Server CA 2011/T=Washington; Not valid after: --; Not valid before: --; Certificate subject: Name: C=US,L=Redmond,O=Microsoft,OU=WSE,CN=settings-win.data.microsoft.com/T=WA; Not valid after: 10:59:23 am Sep 30, 2023; Not valid before: 10:59:23 am Jun 30, 2022; JA3: Local fingerprint: 28a2c9bd18a11de089ef85a160da29e4; Local field: 771.49196-49195-49200-49199-49188-49187-49192-49191-49162-49161-49172-49171-157-156-61-60-53-47-10,0-5-10-11-13-35-16-2-3-65281,29-23-24,0; Remote fingerprint: 2fe9b0e731d3d41b2b84e8e1d6186836; Remote field: 771.49200,16-23-65281-0
- FLOW EVENT DETAILS:** First packet timestamp: 5:47:47 pm Mar 9, 2023; Last packet timestamp: 5:47:47 pm Mar 9, 2023; Bytes sent: 216; Bytes received: 2720

Note For general information about the Process Analysis page, see [Process Analysis](#).


Exploring XDR Data on the Observations Page

You can retrieve and explore XDR data on the Observations page.

On the left navigation pane, click **Investigate** and click the **Observations** tab.



Tip The following filters are especially useful when searching for observations:



XDR Data on the Observation Details Pane

Click the > icon next to the item of interest.

The following image shows details for an IDS observation. The XDR data is highlighted.

Another example shows the netconn data associated with the observation:

NETCONN
go.microsoft.com

Direction: Outbound
Application Protocol: HTTP

Local IP	[Redacted]
Local Port	58713
Remote IP	23.78.114.164
Remote Port	80
Remote Domain	go.microsoft.com
Remote Location	Toronto ON, Canada

Note For general information about the Observations page, see [Investigate - Observations](#).

Exploring XDR Data on the Alerts Page

You can retrieve and explore XDR data on the Alerts page.

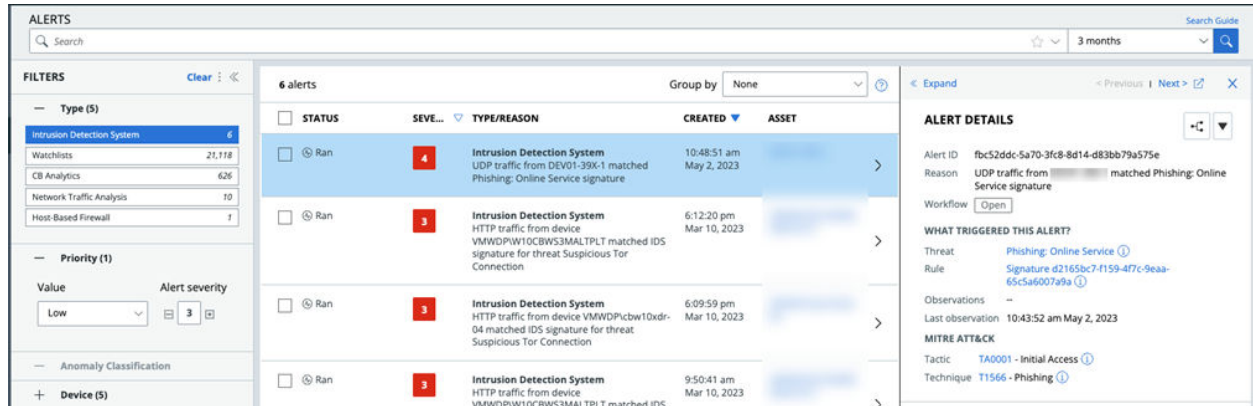
On the left navigation pane, click **Alerts**.

Note

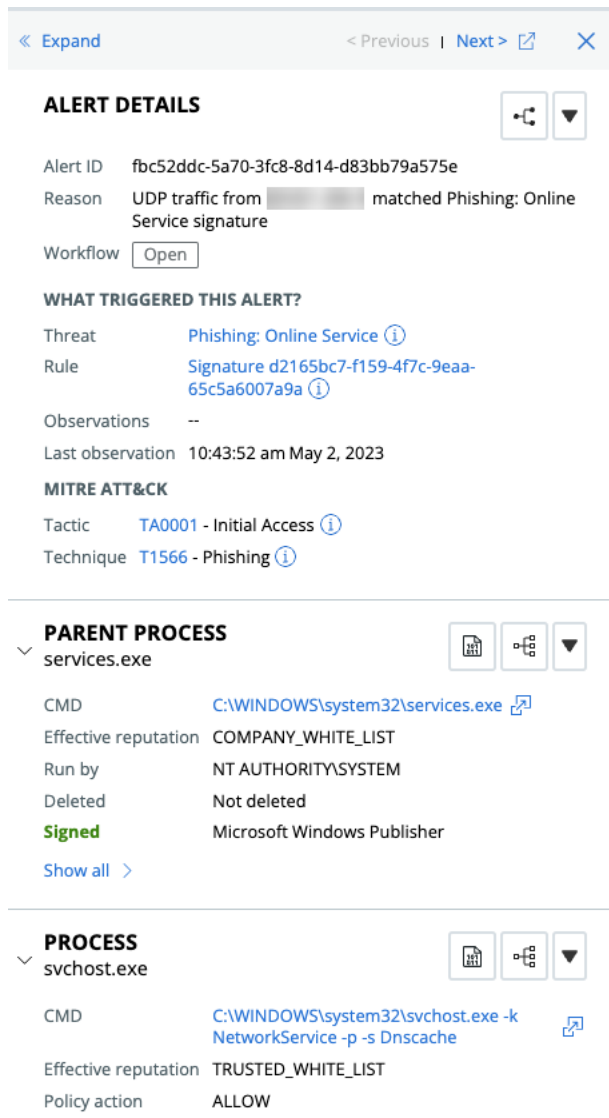
- This topic assumes that you have opened a netconn-generating process on the Alerts page. See [Retrieve XDR Data on the Alerts Page](#).
- For general information about the Alerts page, see [View Alert Details](#).

Click the > icon next to the item of interest.

The following image provides an example of an IDS alert.




The next image shows detailed **Netconn** information for the alert.



Exploring XDR Data on the Alert Triage Page

After you retrieve XDR data, you can explore the data on the Alert Triage page.

Note For instructions on retrieving XDR alert data, see [Retrieve XDR Data on the Alerts Page](#).

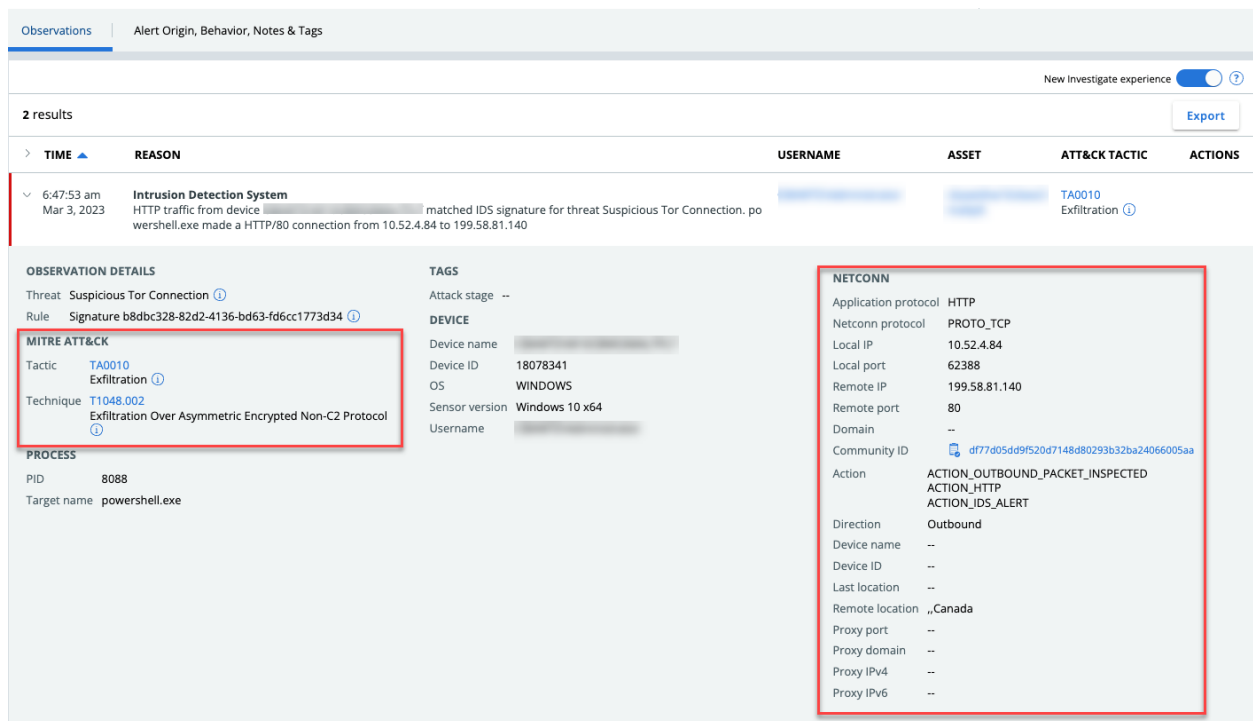
On the left navigation pane, click **Alerts** and then click the **Alert Triage**  icon next to the item of interest.

The Alert Triage diagram displays network nodes (domain or IP) that highlight IDS-specific network connections.

In the bottom pane, select the **Observations** tab to view the following fields:

- **Time** that the observation occurred
- **Reason** for the observation. This is a more detailed view than was previously available with Enriched Events.
- **Username** who executed the process that triggered the observation
- **Asset**
- **ATT&CK Tactic & Technique**
- Available **Actions**

MITRE ATT&CK tactics and techniques and other netconn data display in the **Observations Details** pane that opens when you expand the alert.



The screenshot shows the 'Observations' tab in the Alert Triage interface. It displays a table with 2 results. The first result is expanded, showing the following details:

TIME	REASON	USERNAME	ASSET	ATT&CK TACTIC	ACTIONS
6:47:53 am Mar 3, 2023	Intrusion Detection System HTTP traffic from device [redacted] matched IDS signature for threat Suspicious Tor Connection. powershell.exe made a HTTP/80 connection from 10.52.4.84 to 199.58.81.140	[redacted]	[redacted]	TA0010 Exfiltration	[redacted]

The expanded details pane is divided into several sections:

- OBSERVATION DETAILS:** Threat: Suspicious Tor Connection; Rule: Signature b8dbc328-82d2-4136-bd63-fd6cc1773d34
- MITRE ATT&CK:** Tactic: TA0010 Exfiltration; Technique: T1048.002 Exfiltration Over Asymmetric Encrypted Non-C2 Protocol
- PROCESS:** PID: 8088; Target name: powershell.exe
- TAGS:** Attack stage: --
- DEVICE:** Device name: [redacted]; Device ID: 18078341; OS: WINDOWS; Sensor version: Windows 10 x64; Username: [redacted]
- NETCONN:** Application protocol: HTTP; Netconn protocol: PROTO_TCP; Local IP: 10.52.4.84; Local port: 62388; Remote IP: 199.58.81.140; Remote port: 80; Domain: --; Community ID: df77d05dd9f520d7148d80293b32ba24066005aa; Action: ACTION_OUTBOUND_PACKET_INSPECTED, ACTION_HTTP, ACTION_IDS_ALERT; Direction: Outbound; Device name: --; Device ID: --; Last location: --; Remote location: ,Canada; Proxy port: --; Proxy domain: --; Proxy IPv4: --; Proxy IPv6: --

Note For general information about the Alert Triage page, see [Visualizing Alerts](#).

Investigate



You can investigate and analyze the details of every observation stored in the Carbon Black Cloud, including both failed and successful operations performed by applications and processes on endpoints.

Note

- Tabs do not display on the Investigate page for Carbon Black Cloud Endpoint Standard-only customers. The default view is **Observations**.
 - The **Observations**, **Processes**, and **Auth Events** tabs are only available for Carbon Black Cloud Enterprise EDR customers.
 - As of 26 September 2023, **Enriched Events** is removed from the console. It is replaced with **Observations**.
-

You can collect the data from your search results and, based on the details for your observations and processes, you can take action.

The Investigate page provides an embedded *Search Guide* to assist with creating queries. Use advanced search capabilities to find more detailed information on alerts, conduct investigations, and gain visibility into the prevalence of events, observations, and processes reported from your environment. See also [Advanced Search Techniques](#) in the main section of the *Carbon Black Cloud User Guide*.

Value Search

Use complete values when searching. For example, `powershell` or a trailing wildcard: `power*`.

Search Fields

Form queries that contain search fields: `field:term`. For example, `parent_name:powershell.exe`.

Wildcards

Expand queries using wildcards. `?` matches a single character. For example, `te?t` returns results for "test" and "text". `*` matches zero or more sequential characters. For example, `tes*` returns results for "test," "testing," and "tester".

Leading wildcards are assumed in file extension searches. For example, `process_name:.exe`.

You can use wildcards in a path if you do not quote the value, and if you escape the following special characters with a backslash: `+ - && || ! () { } [] ^ " ~ * ? : / .` For example, to search for **(1+1) :2**, type: `\(1\+1\)\:2`.

Operators

You can refine queries by using operators. Operators must be in uppercase.

- **AND** returns results when both terms are present.
- **OR** returns results when either term is present.
- **NOT** returns results when a term is not present.

Important Be aware of the implicit "AND" operator

There is an implicit "AND" even when the the operator is not used. In the following examples, both queries produce identical results.

- In this example, the "AND" is implied.

```
Process_name:X process_effective_reputation:X
```

- In this example, the "AND" is part of the query.

```
Process_name:X AND process_effective_reputation:X
```

Escaping

Slashes, colons, and spaces must be manually escaped, except when using suggestions and filters.

Date/Time Ranges

You can refine queries by using date/time ranges. For example, `device_timestamp:[2022-10-25T14:00:00Z TO 2022-10-26T15:00:00Z]`.

Count Searches

You can refine queries that include counts together with ranges and wildcards.

- `[3 TO *]` returns count results starting with a value of 3.
- `[* TO 10]` returns counts results up to a value of 10.

Read the following topics next:

- [About Digital Signatures](#)

- [Investigate - Processes](#)
- [Investigate - Observations](#)
- [Investigate - Auth Events](#)
- [Investigating Script-Based Attacks](#)
- [Obtain a Process Hash](#)
- [Add an Investigate Query to a Watchlist Report](#)

About Digital Signatures

This topic describes how Carbon Black Cloud captures and renders the state of digital signatures that are associated with processes and files.

Carbon Black Cloud distinguishes the three ways by which Carbon Black Cloud captures digital signature telemetry and clearly labels the approaches that were used to capture this data.

These distinctions can improve your ability to determine whether the executables and other files observed on your endpoints are legitimate or anomalous, and how you can prevent future unwanted activity.

Signature Definitions

Effective Signature

The **Effective Signature** is the signature at the time the process event occurred. This primary signature corresponds with the process publisher state, the state of the digital signature, or signatures of a Windows or macOS process binary.

Most Prevalent Signature

Carbon Black reviews all prior signatures associated with a particular Windows hash to determine which signature is the most prevalent. This is sourced from the binary details, and can provide additional context when an **Effective Signature** is not available. Because different endpoints and process instances can render differing signature metadata, the signature seen the most for that process is the one shown. This signature has been observed for this file across your fleet of endpoints for the past 90 days (for customers who have Carbon Black Cloud Enterprise EDR with Windows endpoints).

First Seen Signature

What was observed the first time the file was noticed by Carbon Black Cloud on the first endpoint that reported it (for customers who have Carbon Black Cloud Endpoint Standard). The first signature associated with this particular hash within your organization. This field is only available on data from Carbon Black Analytics, and can provide more context when an **Effective Signature** is not available. However, the latest event in the process can have a signature that differs from the first seen signature.

How Carbon Black Cloud Determines which Signature to Display

Carbon Black Cloud shows **Effective Signature** whenever the effective signature data is available, then prioritizes to make **Most Prevalent Signature** available where possible, and only shows **First Seen Signature** when **Most Prevalent Signature** is not available.

Carbon Black Cloud is designed to present consistent signature data when tracking an investigation from page to page in the console, when comparing API responses to the console, or when comparing console and API to Data Forwarder output.

On the Alert Triage page, the **Signature Verification** value mirrors the value that is presented under **Effective Signature**.

In the console, the following pages show one-line summaries and a detailed **Signature** section:

- Alerts
- Alert Triage
- Investigate
- Process Analysis

The following information displays:

- A one-line summary reports the **Effective Signature** state and publisher.
- The **Details** panels show all available properties for the **Effective Signature**, and will make available either the **Most Prevalent Signature** or **First Seen Signature**.
- For each of **Effective Signature**, **Most Prevalent Signature**, or **First Seen Signature**, the console reports four properties:
 - **Signature Status**: Defines whether a signature is present, and if so, whether it has been verified.
 - **Publisher**: Common Name in the certificate used to sign the file.
 - **Certificate Authority**: Common Name in the certificate used to sign the Publisher certificate.
 - **Product**: Product Name property of the Portable Executable header in the Windows binary.

When the console shows a popup that lets you ban or approve a specific hash throughout your organization:

- The message "This hash has been seen on x devices..." obtains that count from the best available API that populates Investigate search.
 - For customers who have Carbon Black Cloud Enterprise EDR, the [Processes Search API](#) determines how many endpoints reported activity for that hash.

- For other customers, the [Observations Search API](#) determines how many endpoints reported activity for that hash.

Note For Carbon Black Cloud Enterprise EDR customers, the prevalence data is based on comprehensive Carbon Black Cloud Enterprise EDR event reporting, and is expected to be accurate. For customers who have Carbon Black Cloud Endpoint Standard only, the prevalence data is based on the Observations data, which is less comprehensive and thus represents the lower bound of prevalence (that is, the binary might have executed on more devices than the count shown).

- The console generally displays **Effective Signature** (where process events are available through the `process_guid`) and either **Most Prevalent Signature** or **First Seen Signature** (except for Linux hashes and Windows SYSTEM process).

When signature data will never exist, Carbon Black Cloud will not show these **Signature** properties in the console. For example, it will always suppress Linux processes and the Windows SYSTEM process.

On the Binary Details page, for Carbon Black Cloud Enterprise EDR customers, the signature data is labeled "Most Prevalent Signature" for consistency. The **Ban** popup does not present the relevant signature data because the only signature data that is relevant in the Binary context is **Most Prevalent**, which is already clearly displayed on the Binary Details page.

Investigate - Processes

Investigate and analyze the details of all processes that have run in your environment.

Note The **Processes** and **Auth Events** tabs are only available for Carbon Black Cloud Enterprise EDR only customers.

On the left navigation pane, click **Investigate** and click the **Processes** tab.

Tip You can also use the [Processes Search API](#) to search through all the data that is reported by your sensors to find one or more processes based on the specific criteria you set.

Search Results

Use the in-product *Search Guide* to access a full list of available search terms to help you create advanced queries.

Results for each process include:

- The latest sensor event and analytics
- Each time a sensor terminated or denied the process
- Each time an event matched a subscribed watchlist

Process Details and Actions

Click the caret to open up additional process, observations, or event information in the right-side panel.

- Click the dropdown arrow next to the process name to take action on the process.
- Click **More** to view additional device details and take action on the device.

Badge indicators can appear next to the process name in the table. Indicators include:

- **Watchlist Hit:** The process has associated watchlist hits. Click the badge for additional information.
- **Alert:** The process has associated alerts. Click the badge for additional information about the highest severity alert. Click the link to view all alerts with the associated process to view on the **Alerts** page.
- **Policy Deny:** A policy action has been taken to keep the process alive, but to deny further operation. This sometimes occurs when the process is denied from loading a banned DLL. Sometimes, this is the case when the process tried to start another process.
- **Policy Terminate:** A policy action has been taken to terminate the process.

Title	Description
Process	The name and path of the process. Click the hyperlinked name to see a visualization of the network connection on the process tree.
Device	The registered name of the device.
Device Time	The device-time of the latest event in a given process segment.
PID	The unique process identifier as defined by the OS.
Username	User context in which the process was executed.
Regmods	The total number of registry modifications associated with the process.
Filemods	The total number of file modifications associated with the process.
Netconns	The total number of network connections associated with the process.
Modloads	The total number of module loads associated with the process.
Childprocs	The total number of child processes associated with the process.

Process Analysis

This section describes the Process Analysis page in the Carbon Black Cloud console.

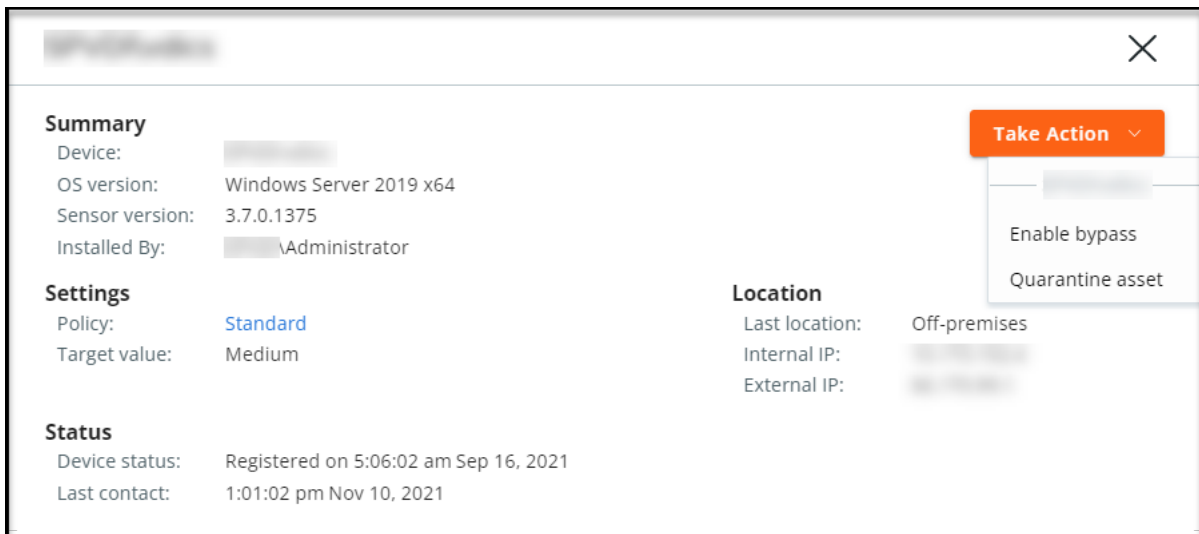
Note If you have Carbon Black XDR, see also [Exploring XDR Data on the Process Analysis Page](#).

At the top right of the Process Analysis page, click the orange **Take Action** button to quickly add a hash to the banned list, enable or disable bypass mode on device, quarantine or unquarantine a device, or view detections in VirusTotal.

The top section of the Process Analysis page contains the following information:

- The primary process that is being analyzed
- The currently selected process (node)
- Date and time
- Process path
- Device details, including:
 - Last logged-in user
 - OS version
 - Device name
 - IP address
 - Location
 - Applied policy

You can click the **More** button to view additional details about this device:



Additional details are included in this view:

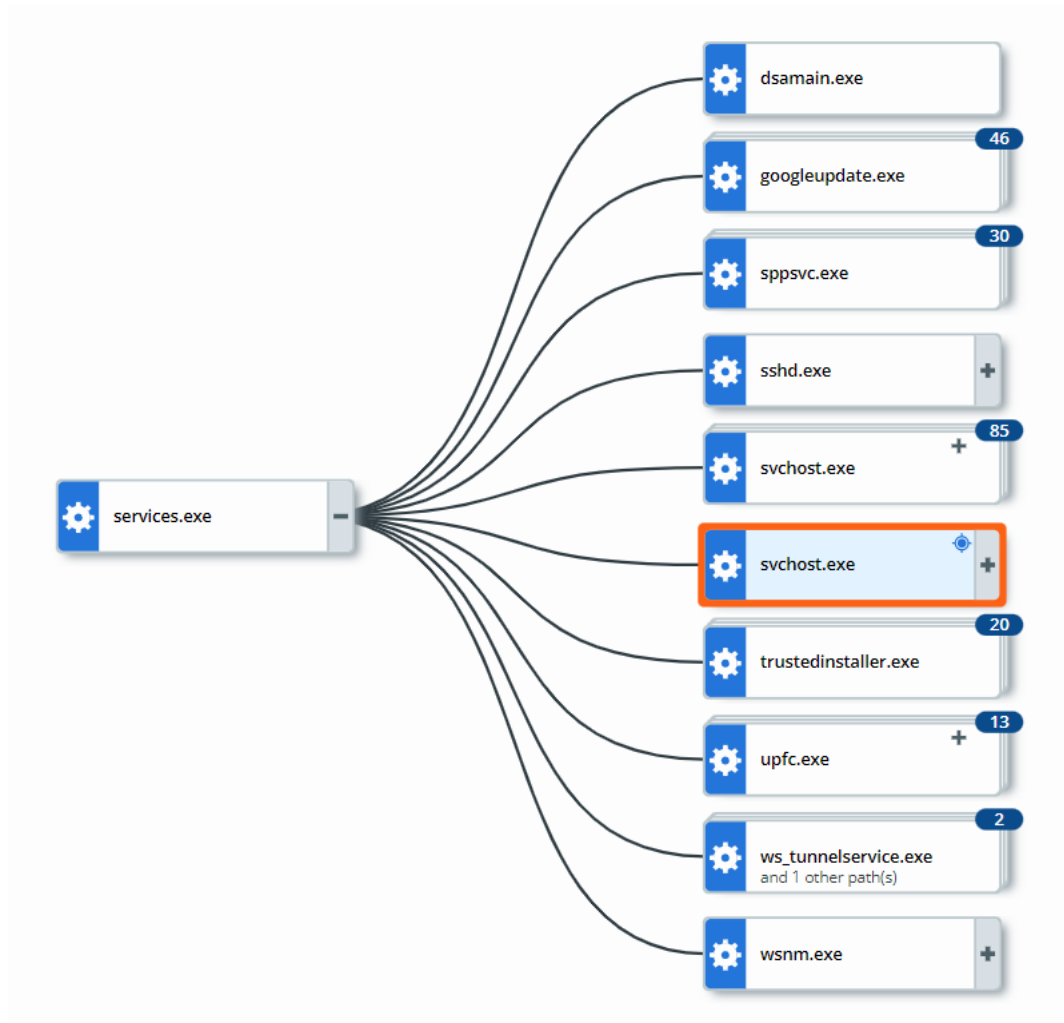
- Sensor version
- Installed by
- Target value
- Device registration date
- Device last contact date

- Last location

You can click the **Take Action** button in this window to enable bypass or quarantine the device.

Visualizing Processes

A visualization of your processes, or a *process tree*, displays in the main section of the Process Analysis page.



Each process in the attack stream is shown in the process tree as a *node* with the attack origin displayed on the left and each subsequent event shown from left to right as the attack progressed. Process trees that have an excessive number of parent or child processes might not display all nodes.

You can group processes by hash by clicking the **Group by hash** toggle. This action causes the process tree to group all processes that have an identical hash, regardless of whether there are child processes or watchlists. The target node is not grouped. Grouping by hash can reduce the number of nodes shown on the page and improve readability.

Selected Node

Click a node to view additional information and take action in the **Selected Node** collapsible panel.

services.exe

CMD [C:\Windows\system32\services.exe](#)

Run by [NT AUTHORITY\SYSTEM](#)

Path [c:\windows\system32\services.exe](#)

MD5 [\[REDACTED\]](#)

SHA-256 [\[REDACTED\]](#)

[Binary Details](#)

REPUTATION [?](#)

Effective [\[REDACTED\]](#)
7:33:57 am Nov 11, 2021

Cloud (Initial) [NOT_LISTED](#)
7:34:28 am Nov 11, 2021

Cloud (Current) [NOT_LISTED](#)
7:39:24 am Nov 11, 2021

PID 636

Start time 5:00:52 am Sep 16, 2021

Process Access Control [?](#)

Elevated --

Integrity --

Privileges --

Signed Microsoft Windows Publisher [ADD](#)

Product --

CA Microsoft Windows Production PCA
2011

Publisher Microsoft Windows Publisher

Binary Details

Select the **Binary Details** button in the **Selected Node** panel to view additional details about a binary.

Note The **Binary Details** button is only available with Carbon Black Cloud Enterprise EDR.

Reputation

Reputation is a given level of trust or distrust.

- **Effective Reputation** is the reputation applied by the sensor at the time the event or observation occurred, based on Carbon Black analytics, cloud intel, and other data.
- **Cloud Reputation (Initial)** is the hash reputation reported by Carbon Black Cloud intel sources at the time that the event or observation was processed by the backend.

- **Cloud Reputation (Current)** is a real-time check of the hash reputation that is reported by Carbon Black Cloud intel sources.

Note **Effective Reputation** is only applicable to users who are running Endpoint Standard.

Signature Definitions

Effective Signature

The **Effective Signature** is the signature at the time the process event occurred. This primary signature corresponds with the process publisher state, the state of the digital signature, or signatures of a Windows or macOS process binary.

Most Prevalent Signature

Carbon Black reviews all prior signatures associated with a particular Windows hash to determine which signature is the most prevalent. This is sourced from the binary details, and can provide additional context when an **Effective Signature** is not available. Because different endpoints and process instances can render differing signature metadata, the signature seen the most for that process is the one shown. This signature has been observed for this file across your fleet of endpoints for the past 90 days (for customers who have Carbon Black Cloud Enterprise EDR with Windows endpoints).

First Seen Signature

What was observed the first time the file was noticed by Carbon Black Cloud on the first endpoint that reported it (for customers who have Carbon Black Cloud Endpoint Standard). The first signature associated with this particular hash within your organization. This field is only available on data from Carbon Black Analytics, and can provide more context when an **Effective Signature** is not available. However, the latest event in the process can have a signature that differs from the first seen signature.

For more information about signatures, see [About Digital Signatures](#).

Process Access Control

- **Elevated:** If “True,” the process is running in an elevated (administrator) context. When a process is elevated, policies that set UAC (user access controls) do not apply.
- **Integrity:** High (administrator), medium (basic user), or low (restricted). Trust is enforced by preventing a process from interacting with processes that have a higher integrity level.
- **Privileges:** Access tokens that encapsulate security identity (privileges) are assigned to each process. Privileges help enforce security boundaries when a process tries to execute.

Watchlist Hits

A process that displays an orange ! indicates that the process has associated watchlist hits. In this case, the **Selected Node** pane also displays:

- Severity score of the latest hit

- Name of the report in which the hit was found
- The query on which the hit occurred
- Time of the occurrence of the event, which was captured as a Watchlist hit

Select the query link to pivot to the Investigate page with the query pre-populated in the Search bar.

Process Event Filtering, Search, and Export

In the bottom section of the Process Analysis page, you can use the filters and the search bar to find process events of interest.

Click **Search Guide** to open the **Process Analysis** section of the in-product *Search Guide*, where the supported search syntax and process fields are documented.

Note The search fields and syntax that are supported on the Process Analysis page can differ from the search fields and syntax that are supported on the **Processes** tab of the Investigate page.

Click the **Export** button to export up to 10,000 process events at a time in CSV format.

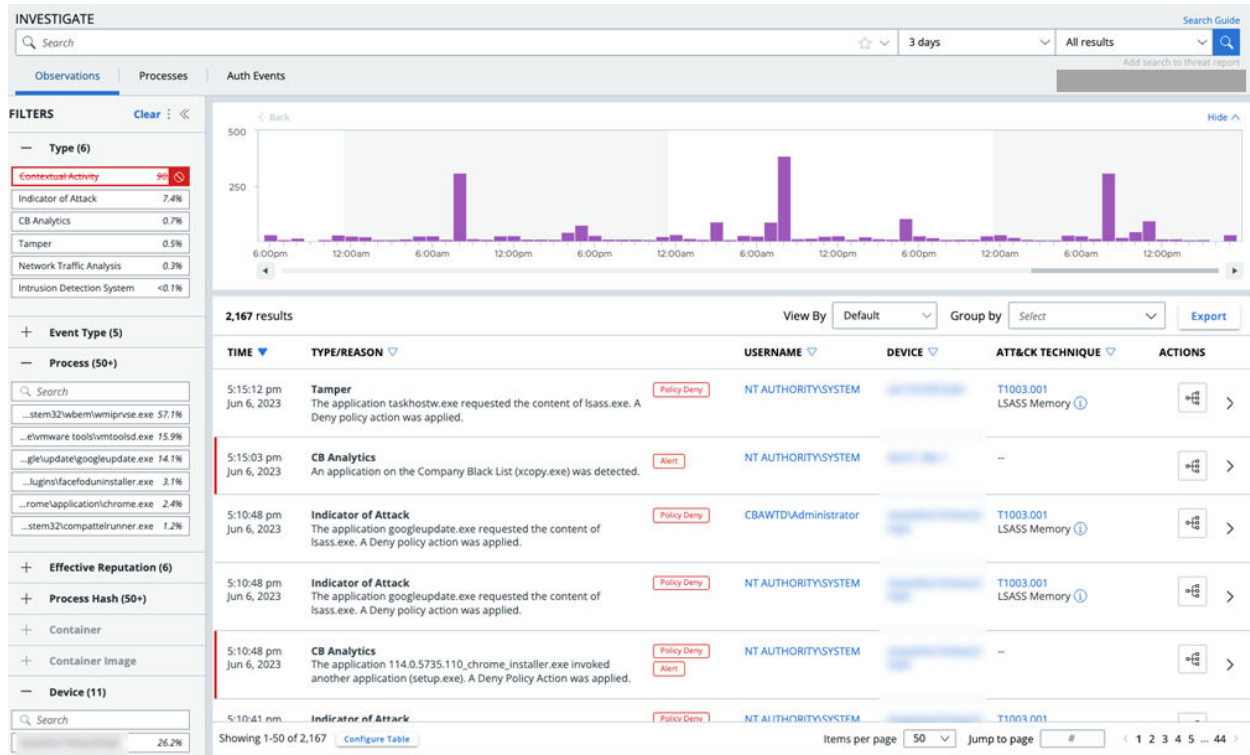
Investigate - Observations

Observations is the default Investigate view for Carbon Black Cloud Endpoint Standard and Carbon Black XDR customers. This page is also visible for Carbon Black Cloud Enterprise EDR customers who have Carbon Black Cloud Endpoint Standard or Carbon Black XDR.

Note

- This section provides a general description of the **Observations** page. For information about XDR-specific data, see [Exploring XDR Data on the Observations Page](#).
 - For Carbon Black Cloud Endpoint Standard customers who do not also have Carbon Black Cloud Enterprise EDR, **Observations** is the default page view.
-

On the left navigation pane, click **Investigate > Observations**.



Tip You can also use the [Observations API](#) to search through all Observations to find one or more specific Observations that match the search criteria.

Observations Overview

The Observations page lets you see interesting or suspicious activity in your environment that does not always reach the importance of generating an alert.

This page lets you search through the stream of notable activities on one or more devices; you can avoid researching all the raw events that are reported by every asset. This page provides a convenient means by which to perform a sweeping search across all your organization's assets.

Observations are the noteworthy, searchable findings across your whole fleet. They complement raw events on Process Analysis page. Not every observation has corresponding raw events; not every observation is truly suspicious.

A smaller subset of observed events are further elevated to Alert status.

Observations are therefore the middle layer of suspicious events.

Searching for Observations

This topic describes ways to filter your searches on the Observations page.

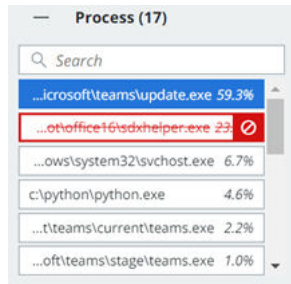
Note Search results are subject to a 10,000 result limit.

You can filter search results in the following ways:

Filter	Examples
<p>Type</p> <hr/> <p>Note View observation <code>Type</code> descriptions in Observation Types.</p>	<ul style="list-style-type: none"> ■ CB Analytics ■ Contextual Activity ■ TAU Intelligence ■ Tamper ■ Blocked Hash ■ Intrusion Detection System ■ Network Traffic Analysis ■ Host-based Firewall ■ Indicator of Attack
<p>Event Type</p>	<ul style="list-style-type: none"> ■ netconn ■ childproc ■ filemod ■ crossproc ■ regmod ■ modload ■ scriptload
<p>Process</p>	<ul style="list-style-type: none"> ■ \system32\svchost.exe ■ system32\services
<p>Effective Reputation</p>	<ul style="list-style-type: none"> ■ TRUSTED_WHITE_LIST ■ LOCAL_WHITE ■ COMPANY_WHITE_LIST ■ ADAPTIVE_WHITE_LIST ■ NOT_LISTED
<p>Process Hash</p>	
<p>Device</p>	<ul style="list-style-type: none"> ■ macOS_workstation ■ Windows11_workstation
<p>Username</p>	<ul style="list-style-type: none"> ■ NETWORK SERVICE ■ SYSTEM ■ LOCAL SERVICE
<p>Parent Effective Reputation</p>	<ul style="list-style-type: none"> ■ TRUSTED_WHITE_LIST ■ LOCAL_WHITE ■ COMPANY_WHITE_LIST ■ ADAPTIVE_WHITE_LIST ■ NOT_LISTED
<p>TTP</p>	<ul style="list-style-type: none"> ■ NETWORK_ACCESS ■ ACTIVE_SERVER ■ RUN_UNKNOWN_APP ■ CODE_DROP ■ POLICY_DENY ■ INTERNATIONAL_SITE

Filter	Examples
Location	<ul style="list-style-type: none"> ■ Seattle,WA,United States ■ San Jose,CA,United States ■ Dublin,L,Ireland
Application Protocol	<ul style="list-style-type: none"> ■ HTTP ■ TLS
ATT&CK Tactic	<ul style="list-style-type: none"> ■ TA0002 ■ TA0004
ATT&CK Technique	<ul style="list-style-type: none"> ■ T1003.0001 ■ T1036.0005 ■ T1105

Tip You can exclude search results by clicking the **Exclude** icon to the right of a filter. For example:



Observation Types

This topic describes observation types.

You can filter your queries by `Type`, as described in [Searching for Observations](#). The following table describes these types.


Type	Description
Blocked Hash	This observation type only applies to Carbon Black Cloud Enterprise EDR environments. It is composed of observations and alerts that surface when processes load hashes that appear on the hash ban list.
CB Analytics	Observations and alerts that were created using Carbon Black Cloud Analytics, which monitors behavioral patterns of processes running on the endpoint. CB Analytics alerts detect attacks but do not prevent them.

Type	Description
Contextual Activity	Contextual Activity are events that were captured by the sensor, but do not match any Carbon Black detections. These events can help to establish context on what else was happening at the endpoint during the same time when potential attack was observed. When Contextual Activity observations are elevated to alerts, they are re-categorized as CB Analytics observations.
Host-based Firewall	Observation that is generated when network traffic matches a Host-based Firewall rule on the endpoint.
Indicator of Attack (IOA)	Observations that arise from endpoint behavior that matches known indicators of attack and are almost always tied to a known MITRE ATT&CK Technique. Indicators of attack are not always malicious in nature, but should be reviewed.
Intrusion Detection System (IDS)	Observations and alerts resulting from network traffic that exhibits known malicious or suspicious patterns on a single network flow. In most cases, these behaviors will map to a known MITRE ATT&CK Technique. Carbon Black monitors for suspicious network traffic against known signatures. When such a signature is found, an IDS observation is generated.
Network Traffic Analysis (NTA)	NTA monitors network availability and activity to identify anomalies.
Tamper	Observations and Alerts that capture evidence of processes that are tampering with the Operating System or the Carbon Black Cloud Sensor. These observations and alerts can result from policy rules that detect and prevent Sensor tampering attempts.
TAU Intelligence	Observations that are generated by specific research findings from the Carbon Black Threat Intelligence Unit (TAU). This category includes observations and alerts that were created using analysis of behavioral patterns on the sensor. These observations and alerts frequently result in prevention.

Histogram

This topic describes the histogram at the top of the Observations page.

The histogram is interactive. You can show or hide the histogram.



You can manipulate the histogram and its associated data in the following ways:

- Click anywhere within the histogram to focus on events that occurred at that particular time.
- Use the scroll bar at the bottom of the histogram to go forward or backwards in time.
- Hover over the histogram to change the cursor to a + sign. Click-and-drag the + to focus on a specific time segment.
- Click the **<Back** button to return to the previous histogram view (before you changed its aspect).
- Group events that have transpired within a certain time period — for example, 1 minute, 10 minutes, 1 hour, or 1 day.

Group By and View By

This topic describes the **Group by** and **View by** features on the Observations page.

In addition to grouping events by time period, you can group by the following fields:

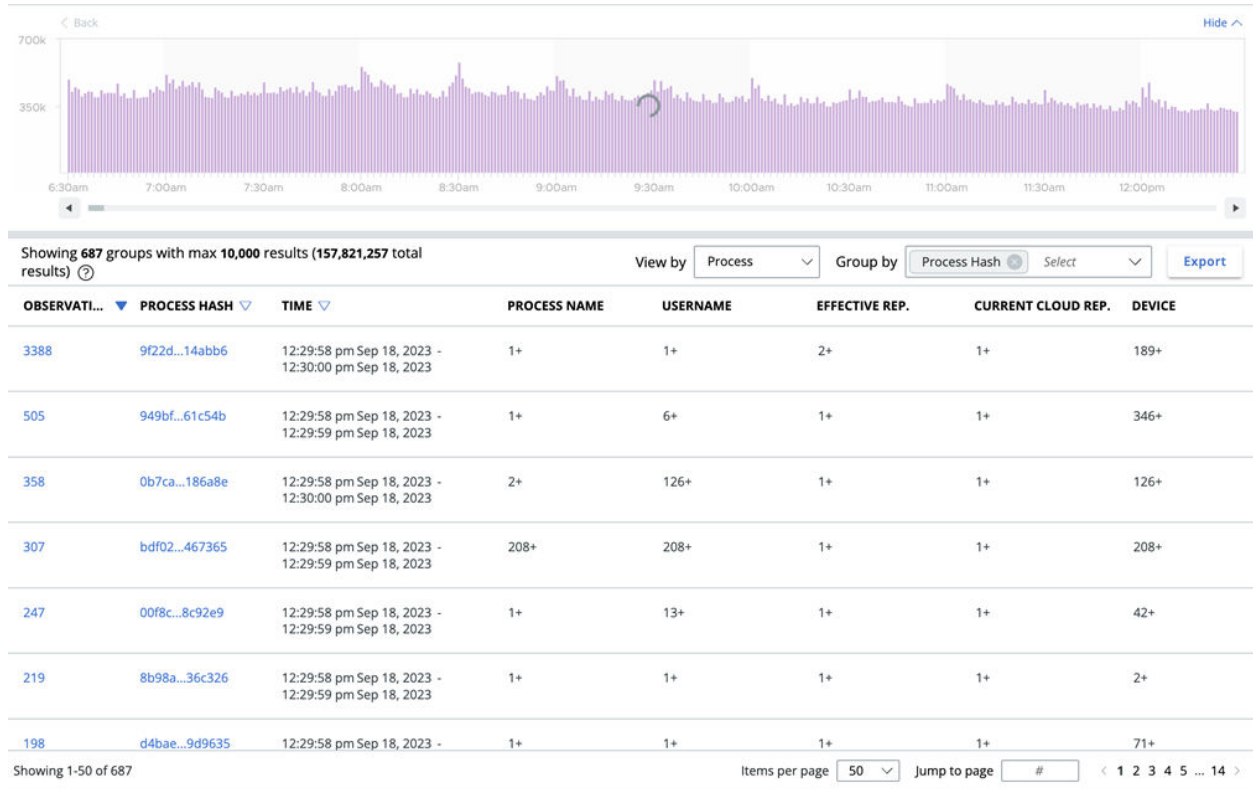
- Type
- Device
- Username
- Local IP
- Remote IP
- ATT&CK Tactic
- Process Hash

Note You can only group by fields that are shown in the table. You can configure which columns (fields) display in the table by clicking the **Configure Table** button at the bottom of the page before you perform a **Group by** action.

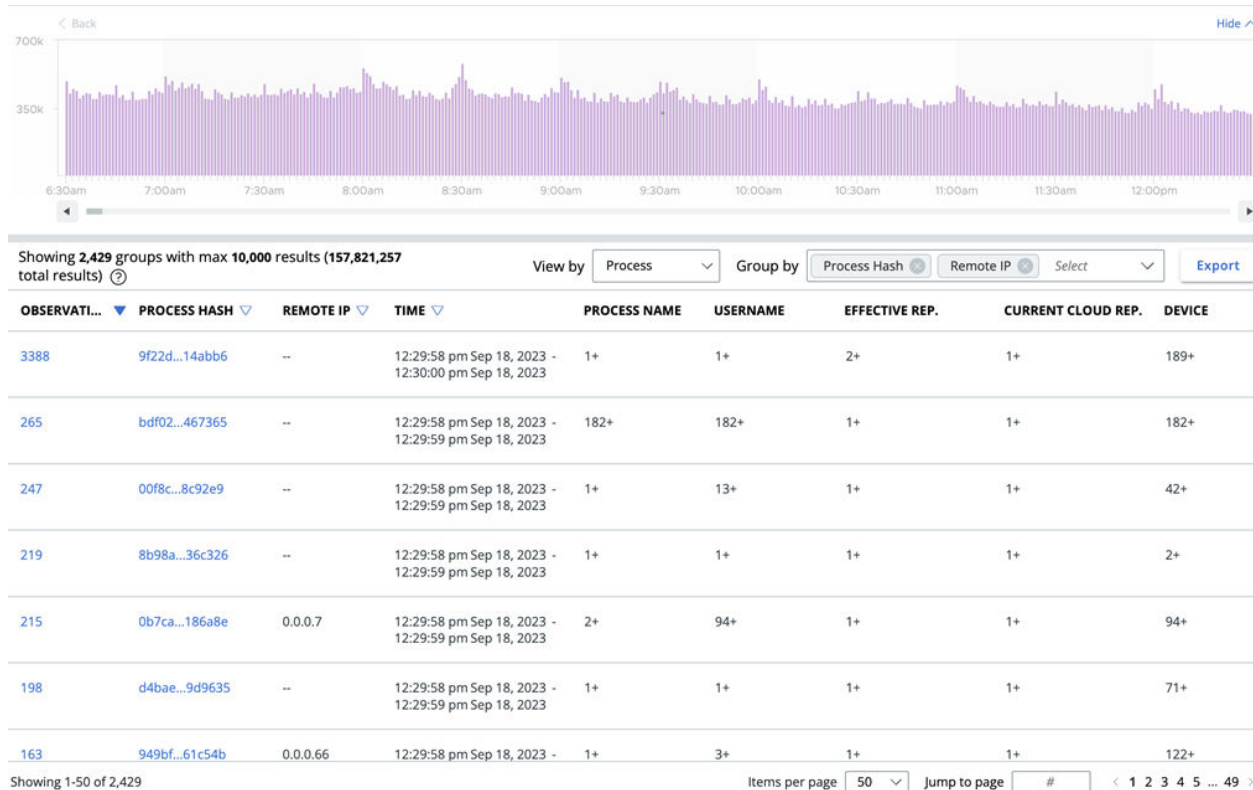
Grouped results are subject to a 10,000 result limit.

Your **View by** options are to view by **Default**, **Process**, **Devices**, or **Network**. You can configure each view to determine which columns display in each view by clicking the **Configure Table** button at the bottom of the page.

The **View by** and **Group by** features allow you to make interesting combinations of results. For example, **View by** lets you quickly change what column set you apply to the data. Selecting **View by Network** provides the IP address, port, and protocol columns. You can then **Group by Remote IP** to see how many Observations occurred for each Remote IP. You can also **View by Process + Group by Process Hash** to see all the applications running in your environment.

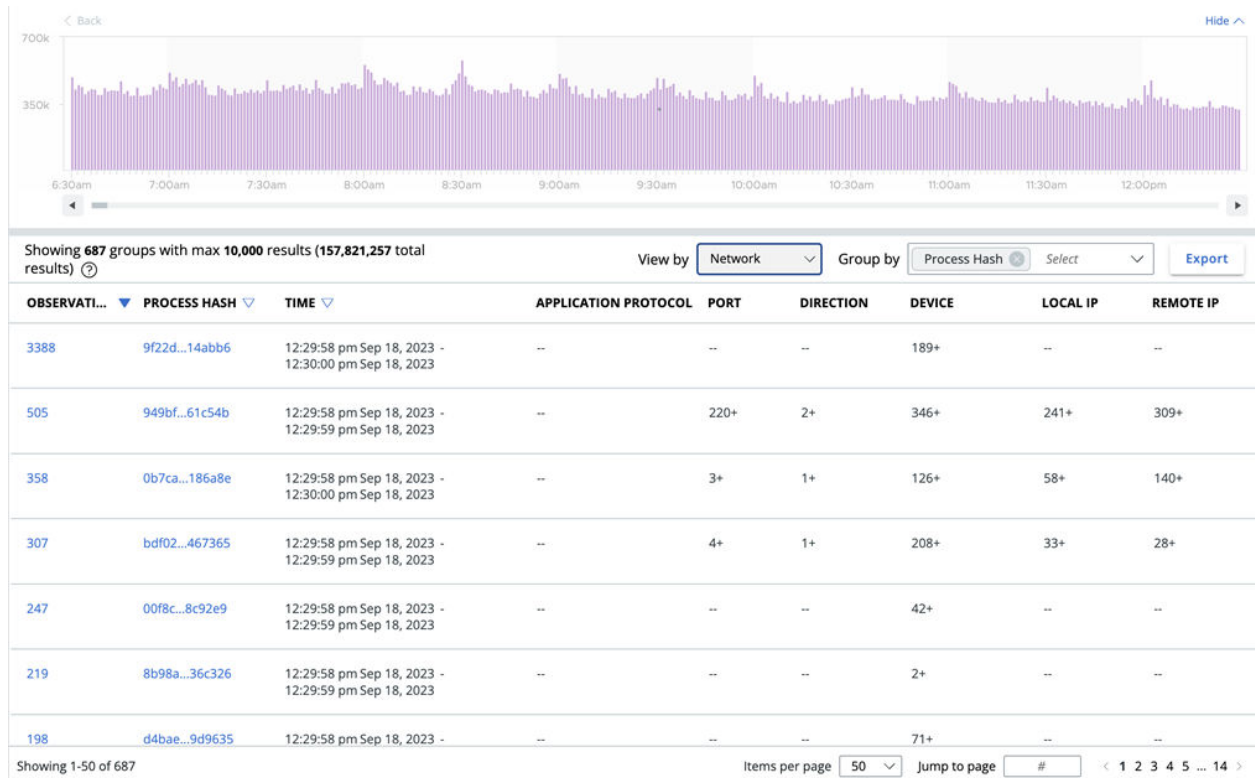


You can then add **Group by Remote IP** to see how many Remote IPs are touched by each piece of software or malware variant.



Multiple groupings lets you see the spread of an activity across multiple dimensions. Grouping first by **Process** and then by **IP** lets you view which processes contact which IP addresses most frequently. This view can help you determine who is contacting specific IP addresses.

You can view grouped data in different ways. For example, **View by Network** and then **Group by Process Hash** to see which processes touch the most IP addresses. This combination helps you see which processes create the most noise.




Observations Search Results


After you perform a search query on the Observations page and retrieve the data set in which you are interested, the results display in tabular format.

Options for viewing are as follows:

- Export this data by clicking the **Export** button at the top right of the table.
- Group and view results as described in [Group By and View By](#).
- Sort the table by using the sorting carets next to most column headers.
- Customize the columns that display by clicking the **Configure Table** button at the bottom left of the table.

Tip You can use the **Configure Table** feature to add columns that are not displayed by default, such as **ATT&CK TECHNIQUE**.

To view an Observation's process and all its events, click the **Process Analysis**  icon at the right of the row. See [Exploring XDR Data on the Process Analysis Page](#) and [Process Analysis](#).

To view additional details about an event, click the  at the right of the row. A summary of details displays. Click **Show all** in any section to view all details in that category. For example:

OBSERVATION DETAILS

Observation ID BDF0203C-A973-11EE-B6A5-3CE9F7FDC0F4:057098fe-6136-4da0

Reason The process c:\windows\system32\windowspowershell\v1.0\powershell.exe attempted to perform Process Doppelg nging by loading the image c:\users\ \desktop\doppeltest.txt within a transaction. A Terminate Policy Action was applied.

Created 10:50:29 am Jan 2, 2024

MITRE ATT&CK

Tactic [TA0005](#) - Defense Evasion 

Technique [T1055.013](#) - Process Doppelg nging 

PARENT PROCESS

cmd.exe  

CMD "C:\Windows\system32\cmd.exe" 

[Show all](#) >

PROCESS

powershell.exe  

CMD [powershell -nop -exec bypass](#) 

Effective reputation TRUSTED_WHITE_LIST

Policy action TERMINATE

Username

Signed Microsoft Windows

Techniques  --

[Show all](#) >

From this panel, you can view binary details of the event, open the Process Analysis page, or take actions on the event.

Available actions on the executable are:

- Remove hash from approved list or Remove hash from banned list
- Add hash to banned list or Add hash to approved list
- Request upload
- Find in VirusTotal
- Delete application

Available actions on the device are:

- Enable bypass

- Quarantine asset
- Go live

Note For help creating a search query, see the in-product *Search Guide*.

Investigate - Auth Events

Carbon Black Cloud Enterprise EDR features *Identity Intelligence*. Identity Intelligence provides visibility into authentication events that occur on Windows endpoints (supported by Carbon Black Cloud Windows Sensor 3.9.1+ running on Windows 10.0.15063+).

Note The **Processes** and **Auth Events** tabs are available for Carbon Black Cloud Enterprise EDR customers only.

Identity Intelligence improves the visibility that Carbon Black Cloud Enterprise EDR provides into user authentication activity. This type of endpoint telemetry is essential for identifying anomalies and threats.

With Identity Intelligence, Carbon Black Cloud Enterprise EDR collects various types of Windows authentication events, which are reported in the **Auth Events** tab on the Investigate page.

The reporting of Windows authentication events supplements the reporting of process events, which enables the correlation of authentication and process activity, and yields more context-rich threat hunting, investigations, and incident response.

Authentication event data provides insight into the following events (and more):

- Attackers' authentication-based tactics, techniques, and procedures (TTPs)
- Who was logged in to an endpoint when process activity of interest occurred
- Who attempted but failed to login to an endpoint
- Brute-force attacks
- Attempted logins outside of expected hours
- Remote authentication attempts from anomalous or suspicious sources
- Privilege escalation attempts
- Account changes
- Use of stolen credentials
- Lateral movement between endpoints
- Insider threat behavior

Some of the benefits Security Operations Center (SOC) Analysts gain from the reporting of authentication events include:

- Increased visibility into endpoint activity

- Additional context during threat hunting and incident response
- Increased potential for correlation of authentication and process events
- Reduced mean time to respond (MTTR)
- Consolidation: reduced reliance on third-party solutions for the collection of authentication events

Tip


- You can also use the [Auth Events API](#) to gain visibility into authentication events that occur on Windows endpoints.
 - You can also use the Auth Events forwarder to send all your organization's Auth Events to an external SIEM, security lake, or custom integration for further analysis. See [Add a Data Forwarder](#).
-

Enable Auth Event Collection

To enable the collection of authentication events, perform the following procedure.

Procedure

- 1 On the left navigation pane, click **Enforce > Policies**.
- 2 Select the policy to modify.
- 3 Click the **Sensor** tab.
- 4 Select the check box for **Enable Auth Events Collection**.



- 5 Click **Save**.

Viewing Auth Events in the Console

The **Auth Events** tab on the Investigate page displays user authentication events that occur on the Windows endpoints of Carbon Black Cloud Enterprise EDR customers.

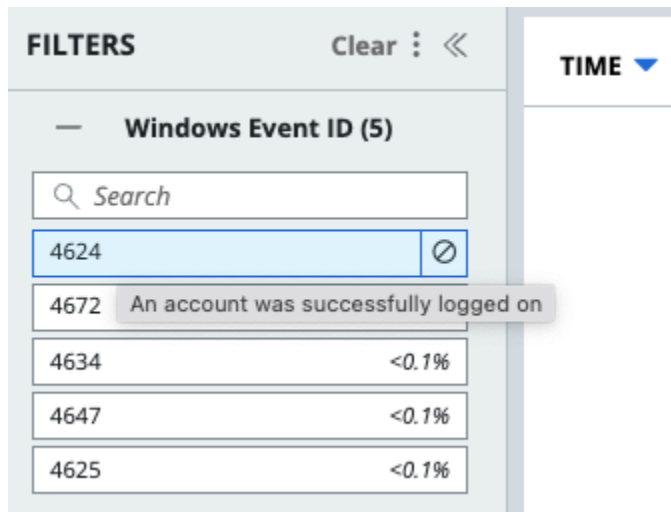
Note The collection of authentication events is disabled by default. Before you can view authentication events on the Investigate page, you must **Enable Auth Events Collection** in the policy. See [Enable Auth Event Collection](#).

On the left navigation pane, click **Investigate** and click the **Auth Events** tab. Search for events. Refer to the in-product *Search Guide* to view the available search fields for Auth Events. You can filter events by:

Container	Container Image	Device
Domain	Interactive Logon	Logon ID

Logon Type	Parent	Policy
Port	Privileges	Process
Remote Device	Remote IP	Remote Location
Remote Logon	User ID (Windows Security ID)	Username
Windows Event ID		

Note The Windows Event ID filter includes a tooltip feature that becomes visible when you hover over the filter. The tooltip describes the Windows Event ID. For example:



Windows Authentication Events

This topic describes the authentication events that Carbon Black Cloud collects and reports.

Windows Event ID	Description
4624	An account was successfully logged on
4625	An account failed to log on
4634	The account was logged off
4647	User initiated logoff
4672	Special privileges assigned to new logon (administrator equivalent)
4740	A user account was locked out
	Logon session detected

For example:

TIME ▾	EVENT ID ▾	DESCRIPTION	USERNAME ▾	DEVICE ▾
1:21:13 am Nov 9, 2022	4624	An account was successfully logged on	SYSTEM	
1:21:13 am Nov 9, 2022	4672	Special privileges assigned to new logon		
1:21:11 am Nov 9, 2022	4634	An account was logged off		
12:52:11 am Nov 9, 2022	4624	An account was successfully logged on	SYSTEM	
12:52:11 am Nov 9, 2022	4672	Special privileges assigned to new logon		

Logon session detected events occur when:

- The sensor was upgraded to a version that supports the Auth Events feature after the logon event occurred, but the sensor still detects the existence of an active logon session.
- Collection of Auth Events was enabled after the logon event occurred, but the sensor still detects the existence of an active logon session.

Logon session detected events do not have event IDs because the collection of authentication events was inactive when the authentication event occurred on the asset.

The **Auth Events** tab introduces a **Group by** feature, which allows you to group authentication events by one or more shared attributes.

You can group authentication events by:

- Windows Event ID
- Username
- Device
- Remote IP
- Time (1 minute 10 minutes, 1 hour, 1 day)

For example, to group all events that have the same Windows Event ID, select **Windows Event ID** in the **Group by** dropdown menu. The resulting table lists the number of events per grouping by Windows Event ID. You can click that number to expand the group to view all events (within the 10,000 results limit) for that particular Windows Event ID that is within the selected time range and/or number of results criteria.

To group events that have the same Windows Event ID that occurred within the same 1 hour periods, select **Windows Event ID** and **1 hour** in the **Group by** dropdown menu.

Multiple **Group by** attributes can be applied to the authentication event results at a time. The order in which the **Group by** attributes are selected does not affect how the results are grouped or displayed. The application of **Group by** attributes is not sequential.

6 groups with 774 results

EVENTS	TIME	EVENT ID	DESCRIPTION	USERNAME	DEVICE
376	1:21:13 am Nov 9, 2022	4624	An account was successfully logged on	SYSTEM	[REDACTED]
268	1:21:13 am Nov 9, 2022	4672	Special privileges assigned to new logon	[REDACTED]	[REDACTED]
111	1:21:11 am Nov 9, 2022	4634	An account was logged off	[REDACTED]	[REDACTED]
6	11:23:02 am Nov 8, 2022	4625	An account failed to log on	[REDACTED]	[REDACTED]
2	11:23:02 am Nov 8, 2022	4740	A user account was locked out	[REDACTED]	[REDACTED]
11	10:56:55 am Nov 8, 2022	--	--	[REDACTED]	[REDACTED]

Authentication Event Details

This topic describes the expanded event details that are available for authentication events.

On the left navigation pane, click **Investigate** and click the **Auth Events** tab. Search for events.

On the right of any event row, click the > to display additional event information.

Auth Event 1:22:15 am Nov 9, 2022 ✕

EVENT DETAILS ⌵

Windows event ID 4624
 Description An account was successfully logged on
 Result Access Granted

IDENTITY

User ID S-1-5-18
 Username SYSTEM
 Domain NT AUTHORITY
 Logon ID 00000000-000003E7
 Linked logon ID --
[Show all >](#)

PROCESS 📄 🗨️ ⌵

Isass.exe

CMD C:\WINDOWS\system32\lsass.exe
 Effective Reputation TRUSTED_WHITE_LIST
 Run by NT AUTHORITY\SYSTEM
 Signed Microsoft Windows Publisher
 Techniques ?
[Show all >](#)

DEVICE 🔍 Go Live ⌵

User [REDACTED]
 OS version Windows 11 x64
 Sensor version 3.9.0.2315
 Policy Standard
[Show all >](#)

When results are grouped using the **Group by** dropdown menu and you click the > for a group of authentication event results, the **Event Details** panel includes GROUP DETAILS, LAST EVENT DETAILS, PROCESS, and DEVICE sections. The GROUP DETAILS section summarizes the following:

- **Group by** criteria
- Number of events in the group
- Times of the first and last events in the group

- Additional information that is common between the events in the group

The **LAST EVENT DETAILS** section includes information about the most recent event in the group.

When you click the > for a single authentication event result, the **Event Details** panel includes **EVENT DETAILS**, **PROCESS**, and **DEVICE** sections.

The **Event Details** panel on the Auth Events page introduces a multi-attribute Investigate feature that lets you pivot to other results that have the same values for those attributes. The pivot options include:

- Username & device
- Device & remote IP (available for remote authentication events)
- Username & Windows event ID

In this example, selecting the **Username & device** option in the **Investigate** dropdown menu takes you to a search for results that have the same `Username` and `Device` values:

EVENT DETAILS

Windows event ID [4624](#)

Description An account was suc
Result Access Granted

IDENTITY

User ID

Username

Domain

Logon ID

Linked logon ID --

[Show all >](#)

Username & device

Device & remote IP

Username & Windows event ID

Single-attribute pivots are supported. Some values in the **Event Details** panel are hyperlinked to enable pivoting based on those values. In this example, `4624` is hyperlinked in the `Windows event ID` field. Clicking `4624` will take you to a search for all results that have `windows_event_id:4624` in the **Auth Events** tab.

Investigating Script-Based Attacks

Script-based attacks are commonly used to gain entry into systems and to move laterally to inflict damage. On the Investigate page, you can find information on script-based attacks and you can identify malicious code in obfuscated PowerShell scripts.

To reveal hidden threats, tools in the Carbon Black Cloud console can decode the contents of the obfuscated PowerShell scripts. You can review the decoded scripts in the right-side panel for a particular event. Syntax highlighting makes it easier to scan for string content, PowerShell commands, and function calls when you search for malicious content.

Investigate Obfuscated PowerShell Scripts



The Carbon Black Cloud console can expose specific details and the decoded version of obfuscated PowerShell scripts, which can help to provide enhanced visibility into these types of attacks.

You can use this procedure to see the decoded content of an obfuscated PowerShell script.


Procedure

- 1 On the left navigation pane, click **Investigate**.
- 2 Do one of the following, depending your product configuration:

Product	Step
Endpoint Standard	<p>On the Investigate > Observations page, find processes where the executable is <code>powershell.exe</code>.</p> <p>You can use the search facility by directly typing:</p> <pre>process_name: powershell.exe</pre> <p>You can modify the time range for the search. For further narrowing of the results, use the filters in the left pane.</p> <p>For more search fields, see the Search Guide, embedded at the top right of the page.</p>
Enterprise EDR	<p>On the Processes tab, find processes where the executable is <code>powershell.exe</code>.</p> <p>You can use the search facility by directly typing:</p> <pre>process_name: powershell.exe</pre> <p>You can modify the time range for the search. For further narrowing of the results, use the filters in the left pane.</p> <p>For more search fields, see the Search Guide, embedded at the top right of the page.</p>

- 3 Select the event or process to investigate. Click the caret  at the end of a row. The **Event Details** panel displays details of the event to the right.
- 4 In the **Process** section in the **Event Details** panel, find the **CMD** line and click the expand icon .

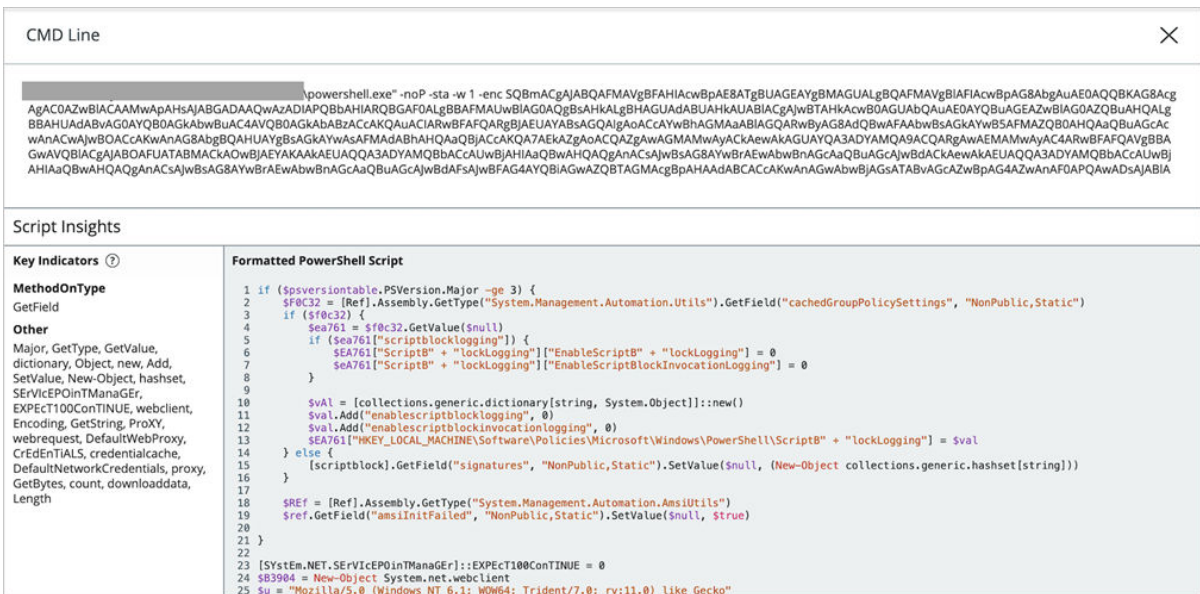
Results

After clicking  for the **Process CMD**, distinguish the difference in the output between a non-PowerShell process and a PowerShell process:

- For a non-PowerShell process, command line arguments are displayed under **CMD Line**.



- For an obfuscated PowerShell process, the decoded script code is displayed with colored text and highlighted keywords under **Key Indicators**.



What to do next

Proceed with your alert triage or threat hunting and determine whether the intent is malicious or not.

Obtain a Process Hash

You can obtain an event's process SHA-256 hash on the **Observations** tab of the Investigate page.

Procedure

- On the left navigation pane, click **Investigate** and click the **Observations** tab.
- Search for an event; you can use the **Process** filter in the left pane to narrow the search results.

3 In the **View by** dropdown menu above the search results table, select **Process**.

The process hash displays in the second column of the search results table.

4 There are three ways in which you can copy (obtain) the process hash:

- Hover over the truncated process hash. The full hash value displays: select the hash value and then press **Ctrl-C** to copy the hash.
- Click the truncated process hash and then press **Ctrl-C** to copy the hash.
- Click the carat **>** icon at the right side of the event. The **Event Details** pane opens. Scroll down to the **Process** section and click **Show all**. Select the hash value and then press **Ctrl-C** to copy the hash.

PROCESS 📄 🗑️ ▼

svchost.exe

CMD 🔗
 C:\Windows\system32\svchost.exe -k NetworkService -p -s CryptSvc

Effective reputation TRUSTED_WHITE_LIST

Run by NT AUTHORITY\NETWORK SERVICE

Signed Microsoft Windows Publisher

Techniques network_access

Hide ▼

Path c:\windows\system32\svchost.exe

MD5 f586835082f632dc8d9404d83bc16316

SHA-256 643ec58e82e0272c97c2a59f6020970d881af19c0ad5029db9c958c13b6558c7

PID 3820

Add an Investigate Query to a Watchlist Report

You can create a custom Indicator of Compromise (IOC) by adding a query to an existing or newly created threat report in an existing or newly created watchlist.

Procedure

- 1 On the left navigation bar, click **Investigate**.
- 2 Execute a query from the search text box and confirm the results.
- 3 To include this query in a watchlist's IOC, click the **Add search to Watchlist Report** link under the search text box.

The **Add Query** window displays.

- 4 Do one of the following:
 - Select an existing watchlist and threat report.
 - a Select a watchlist from the dropdown menu in the **Select a Watchlist** section.

- b Select a threat report from the dropdown menu in the **Add a query to a report** section.
 - Select an existing watchlist and create a new threat report.
 - a Select a watchlist from the dropdown menu in the **Select a Watchlist** section.
 - b Click **Add new** in the **Add query to a report** section.
 - c Enter a meaningful name for the new threat report.
 - d Optionally, include a description, level of severity to trigger the watchlist hit and related tags for the new watchlist report.
 - Create a new watchlist and threat report.
 - a Click **Add new** in the **Select a Watchlist** section.
 - b Enter a meaningful name for the new watchlist.
 - c Optionally, provide the purpose of the watchlist by populating the rest of the fields for the new watchlist.

The **Alert on hit** setting determines how (or if) you are notified when an event matches the query.
 - d Click **Add new** in the **Add query to a report** section.
 - e Enter a meaningful name for the new threat report.
 - f Optionally, include a description and level of severity to trigger the watchlist hit and related tags for the new watchlist report.
- 5 To apply the changes, click **Save**.

Results

A **Successfully created IOC** notification appears on the top of the page.

What to do next

Locate the search query and perform actions on it.

- 1 On the left navigation bar, click **Enforce > Watchlists** page and select the custom watchlist.
- 2 Select the **Reports** tab and click the name of the custom watchlist report.

You can view the newly added query that is listed under IOC and perform actions on it. You can edit, disable, delete, or investigate the query.

Alerts

9

Alerts indicate suspicious behavior and known threats in your environment. We recommend that you regularly review alerts to determine whether you need to take action or modify policies.

- On the left navigation pane, click **Alerts**.
- To expand and view alert details, **double-click** the alert row in the table.

Note Timestamps within the console are displayed in the user's local time zone. Hover over timestamps to view your local time in relation to the UTC time zone.

Tip You can also use the [Alerts API](#) to automate the retrieval and management of alerts. See also: [Alert Bulk Export](#)

Read the following topics next:

- [View Alert Details](#)
- [Group Alerts](#)
- [View Alerts in a Threat ID](#)
- [Editing the Alert Workflow](#)
- [Search Basics](#)
- [Alert Triage](#)
- [Script Host Replacement Occurrence](#)

View Alert Details

Use the following procedure to view alert details.


Note See also [Exploring XDR Data on the Alerts Page](#).

Procedure

- 1 On the left navigation pane, click **Alerts**.

Note In the table, the **Status** column displays **Policy Applied** with a red shield icon if an action was taken by a policy on a Carbon Black Analytics alert.

2 To view the details of an alert, do one of the following:

- Double-click the alert.
- Click the  icon to the right of the alert row.

The expanded, right-side pane displays. An **Alert Details** summary pane describes the type of alert, the alert ID, the reason for the alert, the policy and rule name, and the workflow status.

3 Click **Show All** under the **Determination** to view the **Anomaly Classification** pane. You can view the prevalence of an alert across all organizations and for your organization. The prevalence is categorized as very common, average, or rare. See [Anomaly Classification](#).

4 Click  to view the **Alert Details** pane in a separate tab and to open further panes.

The expanded view displays the following panes:

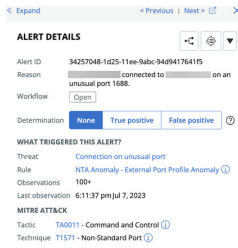
- Alert Details
- Report
- Parent Process
- Child Process
- Process
- Remediation
- Asset
- Alert ID History
- Threat ID History

5 You can:

- Click **<Previous** or **>Next** to view the alert details of the previous or subsequent alert.
- Use the respective buttons in the upper-right corner of the **Alert Details** section to further triage or investigate the alert.
- View the causes of the alert in the **What triggered this alert?** section. If the number of observations displays **100+**, you can:

- Click the **Alert triage** icon  to view 100 observations.

- Click the **Investigate** icon  to view all the data beyond the 100 observations.



- Export alert details to a PDF format.

Use the **Export** button in the upper-right corner of the **Alert Details** section. After 5-10 seconds, the PDF displays in your browser downloads.

- The file name format is `Alert_Report_{alert_id}_{date}.pdf`.
- The PDF contains a page break between each alert details section.
- You can only export a single alert at one time.
- Click the down-arrow ▼ icon to view additional actions you can take.

Note If the alert was generated by a Core Prevention policy, you will have the option to create a Core Prevention policy exception in the **Remediation** pane. See [Add a Core Prevention Process Exclusion from the Alerts Page](#).

Alert Types

Alerts can come from several sources: **Watchlists**, **USB Device Control**, **CB Analytics**, **Host-Based Firewall**, **Containers Runtime**, or **Intrusion Detection System (IDS)**. View alerts from each source by using the **Type** filter.

Watchlists Alerts

Watchlists provide custom detection and continuous monitoring of your environment for potential threats and suspicious activity.

Receiving alerts from watchlists are optional and are configurable on the **Watchlists** page when you subscribe to a watchlist or build a custom watchlist.

USB Device Control Alerts

When an end user tries to access a blocked USB device, a deny policy action is triggered, resulting in an alert. USB Device Control alerts cannot be triaged or investigated.

CB Analytics Alerts

CB Analytics alerts are detections generated by the Carbon Black Cloud analytics engine.

Host-Based Firewall Alerts

Host-Based Firewall alerts notify users when one of their defined firewall rules is violated. If the rule is set to **Block and Alert** on the **Policies** page, an associated alert is generated.

Note Host-Based Firewall alerts contain a maximum of 100 observations. Beyond 100, Carbon Black Cloud suppresses additional duplicate observations.

Containers Runtime Alerts

Containers runtime alerts indicate behavior that is suspected as malicious according to the containers runtime policy. These alerts are a result of one of the following:

- An anomaly in the workload's behavior or a result of behavior that matches a known attack pattern, such as port scanning.
- An outbound connection to IP addresses with bad reputation.

Intrusion Detection System (IDS) Alerts

IDS monitors network activity against known signatures for potential threats and suspicious activity.

We recommend only selecting the **Threat** box in the filters panel when reviewing your queue of CB Analytics alerts to help prioritize and focus your analysis.

Note IDS alerts contain a maximum of 100 observations. Beyond 100, Carbon Black Cloud suppresses additional duplicate observations.

View Specific Alert Types

Use this procedure to view specific alert types.

Procedure

- 1 On the left navigation pane, click **Alerts**.
- 2 In the **Filters** pane, under **Type**, select one of the following to display the alerts specific to that type:
 - **CB Analytics**
 - **Watchlists**
 - **USB Device Control**
 - **Host-Based Firewall**
 - **Containers Runtime**
 - **Intrusion Detection System**

Note You can select more than one type at a time.

The respective alerts display in a list to the right of the **Filters** pane.

- 3 Double-click an alert or click the > to the right of the **Actions** column to view the expanded right-side panel.
- 4 For each alert, you can use the drop-down arrow in the upper-right corner of the Alert Details section of the right-panel.

The options available depend on the alert type. See: [Take Action on Alerts](#).

Alert and Report Severity

Severity scores indicate the relative importance of an alert.

Click the **S** column to sort the alerts in your queue by severity score and identify which alerts might require immediate attention.

CB Analytics - Alert severity

Alert severity indicates the relative importance of a CB Analytics alert.

- **Severity 1-2:** Activities such as port scans, malware drops, changes to system configuration files, persistence, etc.
- **Severity 3-5:** Activities such as malware running, generic virus-like behavior, monitoring user input, potential memory scraping, password theft, etc.
- **Severity 6-10:** Activities such as reverse command shells, process hollowing, destructive malware, hidden processes and tool sets, applications that talk on the network but should not, etc.

Watchlists - Report severity

Report severity indicates the relative importance of threat report within a Watchlists alert.

The severity of a report is determined by the creator of the report. If you create your own report, you can determine the report's severity, with 1 being the least severe, and 10 being the most severe.

Target value

The target value acts as a multiplier when calculating the threat level of an alert. Target values are defined by the policy to which an endpoint belongs.

The target value is indicated by the number of filled bars under the **T** column in the alerts table.

- **Low:** One bar. Results in a lower threat level.
- **Medium:** Two bars. The baseline target value; does not add a multiplier.
- **High/Mission Critical:** Three or four bars. Both values increase the threat level under the same circumstances. You may see two or more alerts with identical descriptions but with different alert severities.

Alert ID, Event ID, and Threat ID

There are three types of IDs and it is important to understand how each is used in the application.

Event ID: A specific action that involves up to three different hashes (Parent App, Selected App, Target App) occurring on a single device at a specific time. Event IDs are found in the event details on the **Investigate** page. Every event sent from the sensor to the console is assigned a unique Event ID.

Alert ID: Similar events taking place within a similar timeframe (+/- 15m) on a single device. Event IDs are grouped into a single Alert ID by Carbon Black analytics. Each alert is assigned a unique Alert ID. This is true even if subsequent alerts have the same hash, action, or device.

Threat ID: Similar alerts tied together across multiple devices and timeframes. Threat IDs can be used to search for related Alert IDs on the **Alerts** page. If the application's hash changes, a new Threat ID is assigned.

Anomaly Classification

The **Anomaly Classification** feature detects and automatically identifies alerts that are most likely to be relevant. This feature is available for Carbon Black XDR customers and on certain watchlists.


The system determines the prevalence of the alert by looking at how many times the alert has been seen across all organizations and within your own organization. Prevalence categories include: very common, average, or rare. An alert is more likely to be marked as anomalous if its prevalence is overall rare.

You can use this feature to focus on the alerts that are anomalous and quickly respond to potential issues or threats.

Anomaly Classification Filter

On the **Alerts** page, you can use the **Anomaly Classification** filter to filter alerts into three categories:

- **Anomalous:** Displays alerts that are anomalous.
- **Not Anomalous:** Displays alerts that are not anomalous.
- **Not Classified:** Displays alerts that are not classified.



An anomalous status displays on the **Status** column if an alert is anomalous.

Group Alerts

Group alerts to view similar alerts occurring across multiple endpoints in a single row.

Note By default, alerts are automatically set to **Group by: None**.

In the **Group By: None** view, all alerts are displayed individually in a single alert row, even if an alert is seen on multiple devices.

You can identify alert prioritization and determine when actions need to be taken on an individual alert.

Use the **Group By** drop-down menu in the top right of the table to group all alerts with the same threat ID. See: [Group By: Threat ID](#).

Type/Reason Column

The **Type/Reason** column determines the threat ID of the alert and explains why the alert was created.

Threat ID groups include:

- Watchlist
- CB Analytics
- USB
- Host-Based Firewall
- Containers Runtime
- IDS

Workflow Column

The **Workflow** column indicates whether an alert is open or closed.

Click the status of the alert in the **Workflow** column to view:

- The Alert ID
- The user that updated the workflow status and the timestamp

Note The workflow column is only interactive on a single alert. You cannot click the workflow status of grouped alerts.

Group By: Threat ID

You can group alerts by threat ID to view the number of alerts with the same threat ID.

Procedure

- 1 Click the drop-down menu for **Group by**:
- 2 Select **Threat ID**.

Results

The **Group by: Threat ID** action arranges the alerts with the same threat ID in a singular alert grouping row.

You can view:

- The number of alerts with the same threat ID
- The level of severity for the group of alerts
- The type of alert and reason why the alert was created
- The first and last time an alert with that threat ID was created
- The number of devices which have alerts for a specific threat ID
- The number of alerts in the same threat ID group that are open or closed

View Alerts in a Threat ID

Use this procedure to view all alerts in a threat ID group.

Procedure

- 1 Click the drop-down menu for **Group by: Threat ID**.
- 2 Choose the alert you want to view.
- 3 In the **Actions** column, click **View Alerts**.

Results

A new window displays with the details of the alerts in the threat ID group.

View Threat ID Details

Use this procedure to view the detailed Threat ID pane.

Click  to view the **Alert Details** pane in a separate tab.

Note Click **<Previous** or **>Next** to view the alert details of the previous or subsequent alert.

The detailed Threat ID pane displays the:

- Threat ID summary pane
- Process pane
- Threat ID history pane

Threat ID Summary Pane

The Threat ID Summary pane provides further details about the alert group including:

- Type
- Threat ID
- Reason for the alert
- The number of devices the alert appeared on and the frequency
- Workflow

Click **Show All** to view the **Anomaly Classification** pane from the threat ID summary. You can view the prevalence of an alert across all organizations and for your organization. The prevalence is categorized as very common, average, or rare. See [Anomaly Classification](#).

Process Pane

The Process pane provides the following information about the threat ID alert:

- Effective reputation
- Deleted
- Signature
- Techniques

Threat ID History Pane

The Threat ID History pane displays:

- Threat ID
- Reason for the alert
- The number of devices the alert appeared on and the frequency of the alert

The pane displays any user activity and notes by users added to the threat ID alert.

Click **View All** to view all the user activity history.

Add Notes

You can add a note to the Alert ID History and Threat ID History panes. You can also delete a note but cannot delete a note added by another user.

The **View All** option opens the Alert ID History or Threat ID History pane in an expanded pop-up view that displays all notes added by users to the alert and the history of the workflow changes for the alert.

Editing the Alert Workflow

The **Workflow** column displays the status of the alert.

You can change the workflow of an alert to **Open**, **Closed**, or **In Progress**.

When closing or opening an alert an alert, you can automatically close or open the alert on all devices in the future.

Important The **Automatically close all future alerts with this threat ID** option is based on the threat ID, which is available by using the [Alerts API](#). The threat ID definition varies slightly across CB Analytics, Watchlists, USB Device Control, Host-Based Firewall, Containers Runtime, and Intrusion Detection System alert types:

- **CB Analytics:** Combination of the primary threat actor (usually the SHA-256 hash of the threat actor) and the alert reason that is derived by the Endpoint Standard Analytics engine.
- **Watchlists:** The report that triggered the Watchlist hit.
- **USB Device Control:** Represents a unique USB device.
- **Host-Based Firewall:** Alerts with the same host-based firewall rule and direction.
- **Containers Runtime:** Alerts in the same cluster and namespace with the same policy and rule.
- **IDS:** Alerts with the same process and IDS signature or rule.

If an alert is flagged for dismissal, any future alerts that contain the same threat ID are dismissed.

Note Alerts can present different SHA-256 hashes. To close or open an alert on multiple devices, the hash of the object must be the same.

Close Alerts

Use this procedure to close an alert and to close all related alerts.

Procedure

- 1 On the left navigation pane, click **Alerts**.
- 2 Choose **Group by: None**.

- 3 Click the alert you want to close.
- 4 In the **Actions** column, click the dropdown menu.
- 5 Click **Close**.

The **Close Alert** window displays.

- 6 Click the **Close As** drop-down menu to select a reason for closing the alert:
 - Resolved
 - No reason
 - Resolved - Benign/Known good
 - Duplicate/Cleanup
 - Other
- 7 In the **Manage Related Alerts** section, choose whether to:
 - Close all existing alerts with the same threat ID
 - Automatically close all future alerts with the same threat ID

Note Click **View Alerts** to view all alerts with the same threat ID.

- 8 Add a note to other users outlining the reason closing the alert, and all future alerts if applicable.
- 9 Click **Close Alert**.

Results

The workflow status of the alert changes to **Closed**. The change is recorded in the **Alert ID History** pane. Use the **Alert ID History** pane to view all previous changes to the workflow status of the alert.

Open Alerts

Use this procedure to open an alert and to open all related alerts.

Procedure

- 1 On the left navigation pane, click **Alerts**.
- 2 Choose **Group by: None**.
- 3 Click the alert you want to close.
- 4 In the **Actions** column, click the dropdown menu.
- 5 Click **Open**.

The **Open Alert** window displays.

6 In the **Manage Related Alerts** section, choose whether to:

- Open all existing alerts with the same threat ID
- Automatically open all future alerts with the same threat ID

Note Click **View Alerts** to view all alerts with the same threat ID.

7 Add a note to other users outlining the reason opening the alert, and all future alerts if applicable.

8 Click **Open Alert**.

Results

The workflow status of the alert changes to **Open**. The change is recorded in the **Alert ID History** pane. Use the **Alert ID History** pane to view all previous changes to the workflow status of the alert.

Mark Alert as In Progress

Use this procedure to mark an alert as **In Progress**.

Procedure

- 1 On the left navigation pane, click **Alerts**.
- 2 Choose **Group by: None**.
- 3 Click the alert to close.
- 4 In the **Actions** column, click the dropdown menu.
- 5 Click **Mark In Progress**.

The **Mark Alert In Progress** screen displays where you can manage all related alerts.

- 6 Choose whether to mark all alerts with the same threat ID as **In Progress**.

Note Click **View Alerts** to view all alerts with the same threat ID.

- 7 Click **Ok**.

Results

The workflow status of the alert changes to **In Progress**. The change is recorded in the **Alert ID History** pane. Use the **Alert ID History** pane to view all previous changes to the workflow status of the alert.

Search Basics

You can use the following methodologies when using the search field:

Value Search

Use complete values when searching (for example, powershell) or a trailing wildcard (for example, power*).

Search Fields

Form queries like this when including search fields: field:term

For example:

```
parent_name:powershell.exe
```

Wildcards

Expand queries using wildcards. * ? Matches a single character. For example, te?t will return results for "test" and "text" * * Matches zero or more sequential characters. For example, tes* will return results for "test," "testing," and "tester"

Leading wildcards are assumed in file extension searches.

For example: process_name:.exe

Wildcards can be used in a path if you don't quote the value and escape the following special characters with a backslash: + - && || ! () { } [] ^ " ~ * ? : /

For example: to search for (1+1):2, type: \ (1\+1\) \:2

Operators

Refine queries using operators. Operators must be uppercase.

- **AND** returns results when both terms are present
- **OR** returns results when either term is present
- **NOT** returns results when a term is not present

Escaping

Slashes, colons, and spaces must be manually escaped except when using suggestions and filters.

Date/Time Ranges

Refine queries using date/time ranges, when applicable.

For example: device_timestamp: [2018-10-25T14:00:00Z TO 2018-10-26T15:00:00Z]

Count Searches

Refine queries that include counts with ranges and wildcards.

- [3 TO *] Returns count results starting with a value of 3.
- [* TO 10] Returns counts results up to a value of 10.

Observed Alert data are no longer available on the Alerts page and are now classified as Observations

You can find the Observed Alerts data in the **Observations** page by filtering on CB Analytics.

- 1 On the left navigation pane, click **Investigate > Observations**.
- 2 Under **Filters**, select **Type > CB Analytics**.

Old Observed Alerts are not marked as alerts on the **Observations** page.

Note For help creating complex queries, see [Advanced Search Techniques](#) in the *Carbon Black Cloud User Guide*.

Alert Triage

During alert triage, you can investigate the alert and take action to address the alert.

Important If the Alert Triage page displays "no data," the system may still be gathering data in the background. Please be patient; it may take several minutes to populate the page depending on the quantity of new alerts. Refreshing the page after a brief wait may resolve the issue. (This is a known issue that will be resolved in the near future.)

- Click **Investigate** to view and analyze observations that triggered an alert on the **Investigate** page.
- Click the orange **Take Action** button to:
 - Add to approved list
 - Add to banned list
 - Request upload
 - Find in VirusTotal
 - Delete application
- View the observations that triggered an alert on the **Alert Details** pane.

Note Host-Based Firewall and IDS alerts contain a maximum of 100 observations. Beyond 100, Carbon Black Cloud suppresses additional duplicate observations.

Investigating Alerts

This section describes the best practices for investigating alerts.

Check these items:

- Priority score
- Parent path and name
- TTPs involved
- File reputation
- Network connections

- Event details
- Command lines (if there were any)

Ask these questions:

- Was another program or function successfully called?
- Is the path of the files suspicious?
- Is the process running in the “normal” path?
- What attack stage was it in?
- Was the registry modified?
- Were the file reputations worrisome?

Take other steps as needed:

- Google any application or files that you don’t recognize
- Ask a teammate to review for anything that you missed
- Review any referenced MITRE techniques or watchlist hits
- Use “custom time” to review events 15 minutes prior to occurrence for more insight
- Review observed activity for more context

Take Action on Alerts

In addition to the functions available from the **Take Action** button, there are several other actions you can take on your CB Analytics alerts.

Quarantine a device triggered by an alert

Click **Quarantine Device**, then **Request quarantine**.

Quarantining the device prevents suspicious activity and malware from affecting the rest of your network. A device remains in quarantine until it is removed from the quarantined state. It can take several minutes to place a device in quarantine.

To remove a device from quarantine, click **Unquarantine device(s)**.

Add notes

Add notes to allow for easy search and filtering of alerts, as well as a means of communication between console users. See: [Add Notes](#).

Open or close

Edit the workflow of the alert to open or close an alert. See: [Editing the Alert Workflow](#).

Use Live Response

Click **Go Live** to initiate a Live Response session. Use Live Response to perform remote investigations, contain ongoing attacks, and remediate threats. Users must be assigned a role that has Live Response permissions in the Carbon Black Cloud to use the Live Response functionality. See [Use Live Response](#) and [User Roles](#).

Live Response is available on endpoints running a version 3.0 or later sensor and which have been assigned a policy with Live Response enabled. Live Response can be used on devices in bypass mode or quarantine.

Quarantining an Asset

Quarantine mode allows both Carbon Black Support and Carbon Black Cloud administrators to continue investigating an asset from the Carbon Black Cloud console (Investigate Page, Live Response, Live Query, and so on) while reducing the risks involved with allowing a compromised asset to access the local network

You can quarantine an asset on the Alerts page or Investigate page.

When an asset is quarantined, all incoming and outgoing TCP traffic is blocked to all IP addresses or ports except for those IP addresses and ports that maintain a connection to the Carbon Black Cloud. Quarantined assets can still check in with the Carbon Black Cloud console for status changes; for example, moving from Quarantine to Active state.

During quarantine, Carbon Black Support can pull sensor logs from the asset.

The following caveats apply when an asset is quarantined:

- For assets that are running Windows or macOS operating systems, all UDP connections except for those responsible for DNS requests, UDP/53 and DHCP, UDP/67, and UDP/68, are blocked.
- For assets that are running Linux operating systems, all UDP connections except for those responsible for DNS requests are blocked. Allowed connections can use:
 - UDP/53
 - DHCP requests: UDP/67 and UDP/68 (for IPv4)
 - UDP/546 and UDP/547 (for IPv6)
- DNS and DHCP are allowed to ensure the bilateral communication between the Carbon Black Cloud and the quarantined asset.
- ARP is allowed to ensure that MAC addresses can resolve to IP addresses.
- ICMP (ping) is allowed.
- Quarantine terminates active sockets that are not exempt from Quarantine; this action effectively re-authorizes any existing connections.
- Windows Filtering Platform API is used to determine traffic type per connection on Windows.

- The types of connections, remote investigation, or remediation tools that are allowed and disallowed in quarantine mode cannot be customized

Add Determination for Alerts

You can determine if an alert is either a true or false positive.

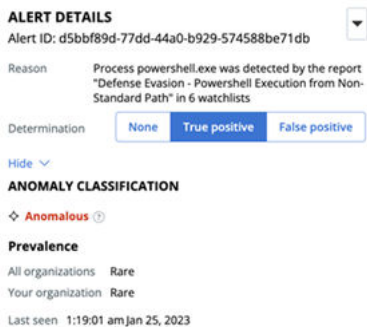
Use the **Alert Details** pane to provide a **True Positive** or **False Positive** alert determination for alerts.

By providing feedback, analysts can contribute to training the model and enhancing the accuracy of the classification system over time. By analyzing the feedback from users, the system improves its classification algorithm and becomes better trained to identify threats in the future.

Note By default, the determination is set to **None**.

Procedure

- 1 On the left navigation pane, click **Alerts**.
- 2 To view the details of an alert, do one of the following:
 - Double-click the alert.
 - Click the > to the right of the **Actions** column.



- 3 Click **True Positive** or **False Positive** to provide alert determination feedback for the alert.

Note This feedback pertains to the evaluation of the specific alert itself rather than the prediction output of the model. Providing feedback is valuable to train the alert classification system, because it performs inferences on the same input stream of alerts.

True and False Positives

This section describes true and false positives for alerts.

True Positives

True positives are alerts that are correctly labeled as malicious. They include:

- Fileless scripting attack or malicious events that might involve malware or other threats

- A file that might have a reputation of `KNOWN_MALWARE`, `SUSPECT_MALWARE`, or `PUP`, or might be `NOT_LISTED`, for example Zero-day (“0-day”)
- Observed behavior or TTPs might be suspicious based on what is “normal” for your environment
- **Detection:** Malicious activity might be detected but not prevented. Typically, this means that a policy needs to be strengthened.
- **Prevention:** Blocking might take place, but only parts of the attack may have been stopped, possibly because of different stages of the attack. Stronger policies are likely needed.

False Positives

False positives are alerts that are incorrectly labeled as malicious or flagged as one of the threat reputations (e.g., `KNOWN_MALWARE`, `SUSPECT_MALWARE`, `PUP`).


False positive can be triggered when:

- A common application is incorrectly flagged as suspicious behavior or suspicious TTPs are observed
- Software that touches canary files triggers ransomware alerts
- Unknown in-house programs are deemed suspicious
- Programs that might not have been excluded cause conflicts (that is, interoperability or unwanted blocks)

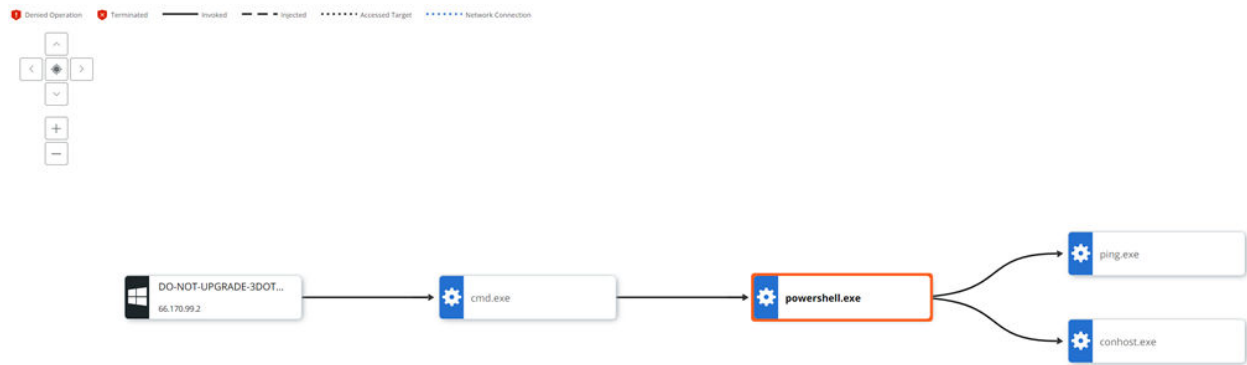
Visualizing Alerts

You can access a visualization or *process tree* of your alerts.

Note If you have Carbon Black Cloud XDR, see also [Exploring XDR Data on the Alert Triage Page](#).

On the Alerts page, click the **Alert Triage**  icon next to the item of interest. The Alert Triage page opens.

Each event in the attack stream (process, file, or network connection) is shown in the process tree as a *node*. The attack origin displays on the left and each subsequent event is shown from left to right as the attack progressed.



You can zoom in and out or pan your view of the image using the navigation buttons in the upper left of the page. You can investigate or close the alert using the respective buttons in the upper right of the page. Use the UP and DOWN arrows in the upper right to traverse between alerts.

Node Types

- **Operating System/Root Node:** The root node at the far left of the process tree represents the host device on which the original activity took place. The root node icon represents the operating system that was running on the device.
- **Gears/Processes:** Processes that have run or are still running.
- **Documents/Files:** Files that were created on disk.
- **Network Connections/IP addresses:** IP addresses are shown as network connection icons.

Note If an operation is denied, an exclamation point (!) displays next to the denied process. If a process is terminated, an X displays next to the terminated process.

Line Types




- **Invoked:** A solid line indicates that one process invoked another process, file, or network connection.
- **Injected:** A dashed line indicates that one process injected code into another process.
- **Read Memory:** A dotted and dashed line indicates that one process attempted to read the virtual memory of another process (but did not inject into the process).
- **Accessed Target:** A dotted line indicates that one process attempted to enter another process (but did not inject into the process).

Selected Node Details Panel

Click a node to view additional node information and take action in the **Selected Node Details** collapsible panel.

PROCESS

powershell.exe

CMD 🔗

powershell.exe -Command IEX([System.Text.Encoding]::A
ItemProperty -Path HKCU:\SOFTWARE\ATOMI...

Effective reputation LOCAL_WHITE

Username

Deleted Not deleted

Signed Microsoft Windows

Techniques ?

fileless

packed_code

run_system_app

run_unknown_app

mitre_t1027_obfuscate_files_or_info



mitre_t1059_001_powershell

mitre_t1059_cmd_line_or_script_inter

[Show all >](#)

DEVICE

DO-NOT-UPGRADE-3


Go Live


User

OS version Windows 10 x64

Sensor version 3.9.2.2698

Asset groups [All assets](#)

Policy [PscrTest](#)

[Show all >](#)

Each section of the panel offers actions that you can perform on the particular element. For example, the **Device** element allows you to:

- Start a **Go live** session to the device
- Enable bypass
- Quarantine the asset

Click **Show all >** to see more details within a section.

Scroll down the panel to see additional sections; for example:

▼ **THREAT ID HISTORY**
 Threat ID ⓘ 43572a2f368f93b9f1e794437b569429
 Reason The application powershell.exe is executing an encoded fileless script (ping.exe).
 Seen on 1 devices, 16 times in the last 6 months
 0 entries [View all](#) 📄

✍️

🏷️

Observations Panel

The bottom section of the Alert Triage page displays Observation details, including:

- Time of the alert
- Reason for the alert
- Process username
- Asset
- Att&ck Tactic

Alert Origin, Behaviors, and TTPs

You can access origin and behavior details about your alerts by clicking the **Alert Triage** icon.

Alert origin: Describes how the primary process for the alert was introduced onto the host, including information about how the primary process was written to disk.

Alert behaviors based on severity: Describes alert behaviors based on severity and displays an interactive TTP graph. Segments of the graph indicate the alert behavior category. Click a category label or graph segment to see a category's related TTPs, color coded by severity.

TTP color severity legend

- **Dark red:** Severe
- **Bright red:** High
- **Orange:** Medium

- **Yellow:** Low
- **Gray:** None

Tip For additional information, see: [TTPs and MITRE Techniques](#) and [TTP Reference](#) in the main section of the *User Guide*.

Alert behavior categories

- **Process Manipulation:** Behaviors with intent to modify and/or read the memory of other processes that are running on the device.
 - **Example:** Injects code into the memory of another process.
- **Generic Suspect:** Behaviors that are generic to multiple malware families, commonly exhibited by known "good" applications.
 - **Example:** Attempts to persist beyond the reboot of a device and enumerating the running processes on a system.
- **Data at Risk:** Behaviors with intent to compromise the confidentiality, availability, or integrity of data on endpoints.
 - **Example:** Ransomware-type behaviors or attempts to access user credentials.
- **Emerging Threats:** Behaviors associated with non-malware attacks.
 - **Example:** Abuse of native command line utilities such as PowerShell, and/or the exploitation of related activities such as buffer overflows.
- **Malware & Application Abuse:** TTPs that are related to files with a generally known "bad" reputation, or applications seen executing files with known bad reputations.

Note This category also represents the monitoring of the execution of system applications. However, these TTPs are given a lower priority rating because of the high likelihood of being non-malicious actions.

- **Network Threat:** Contains all TTPs that involve a process that is either communicating over the network or listening for incoming connections.

Script Host Replacement Occurrence

In Carbon Black Cloud, depending on the offering you enable, a script host replacement can occur.

In different pages of the Carbon Black Cloud console UI, you can view a different name for the same process. The name of the process calling a script is replaced with the name of the script (file) being called by that process.

For example, an event in the Carbon Black Cloud console shows `PowerShell.exe` as the process name and another event shows the `myscript.ps1` script name as the process.

The change of the name of the calling process with the name of the script being called is referred to as script host replacement.

When you enable the Enterprise EDR offering and navigate to the **Process Analysis** page, you can view the name for the calling process as `PowerShell.exe`. The sensor does not perform name replacement and the process name displays the same everywhere.

When you enable the Endpoint Standard offering and navigate to the **Alert Triage** page, you can view the name for the calling process as `myscript.ps1` due to the script host replacement. Here the sensor presents the script name as the process name when PowerShell runs a `.ps1` file to ease the security analyst in seeing the behavior without investigating the event. This is also true for the V6 Alerts API.

When both, Enterprise EDR and Endpoint Standard features are enabled, the script host replacement occurs.

You can add either of the following search terms to the watchlist IOC/search to control the name replacement visibility.

- `enhanced:true` - returns only the events that list the script (file) name as the process name.
- `enhanced:false` - returns only the events that list the process name as is.